

MAJOR REPORT • OPEN ACCESS

Searching for long-lived particles beyond the Standard Model at the Large Hadron Collider

To cite this article: Juliette Alimena *et al* 2020 *J. Phys. G: Nucl. Part. Phys.* **47** 090501

View the [article online](#) for updates and enhancements.

Recent citations

- [Collider probes of real triplet scalar dark matter](#)
Cheng-Wei Chiang *et al*
- [Displaced vertex signatures of a pseudo-Goldstone sterile neutrino](#)
Stéphane Lavignac and Anibal D. Medina
- [Searching for exotic B -meson decays enabled by the CMS L1 track trigger](#)
Jared A. Evans *et al*

Searching for long-lived particles beyond the Standard Model at the Large Hadron Collider

Juliette Alimena

(Experimental Coverage, Backgrounds, Upgrades)¹ 

James Beacham

(Document Editor, Simplified Models)² 

Martino Borsato (Backgrounds, Upgrades)³ 

Yangyang Cheng (Upgrades)⁴ 

Xabier Cid Vidal (Experimental Coverage)⁵ 

Giovanna Cottin (Simplified Models, Reinterpretations)^{6,7,8} 

David Curtin (Simplified Models)⁹ 

Albert De Roeck (Experimental Coverage)¹⁰ 

Nishita Desai (Reinterpretations)¹¹ 

Jared A Evans (Simplified Models, Experimental Coverage)¹² 

Simon Knapen

(Dark Showers)¹³ 

Sabine Kraml (Reinterpretations)¹⁴ 

Andre Lessa (Reinterpretations)¹⁵ 

Zhen Liu (Simplified Models, Backgrounds, Reinterpretations)¹⁶ 

Sascha Mehlhase

(Backgrounds)¹⁷ 

Michael J Ramsey-Musolf (Simplified Models)^{18,19} 

Heather Russell (Experimental Coverage)²⁰ 

Jessie Shelton (Simplified Models, Dark Showers)²¹ 

Brian Shuve (Document Editor, Simplified Models, Simplified Models

Library)^{22,23} 

Monica Verducci (Upgrades)²⁴ 

Jose Zurita (Experimental Coverage)^{25,26} 

Contributors & Endorsers: Todd Adams²⁷, Michael Adersberger²⁸, Cristiano Alpigiani²⁹,

Artur Apresyan³⁰, Robert John Bainbridge³¹,

Varvara Batozskaya³², Hugues Beauchesne³³, Lisa Benato³⁴,

S Berlendis³⁵, Eshwen Bhal³⁶, Freya Blekman³⁷,

Christina Borovilou³⁸, Jamie Boyd¹⁰, Benjamin P Brau³⁹,

Lene Bryngemark⁴⁰, Oliver Buchmueller³¹,

Malte Buschmann⁴¹, William Buttinger¹⁰,

Mario Campanelli⁴², Cari Ceserotti⁴³, Chunhui Chen⁴⁴,

Hsin-Chia Cheng⁴⁵, Sanha Cheong^{46,47}, Matthew Citron⁴⁸,
 Andrea Coccato⁴⁹, V Coco¹⁰, Eric Conte⁵⁰, Félix Cormier⁵¹,
 Louie D Corpe⁴², Nathaniel Craig⁴⁸, Yanou Cui²³,
 Elena Dall'Occo⁵², C Dallapiccola³⁹, M R Darwisch⁵³,
 Alessandro Davoli^{54,55}, Annapaola de Cosa⁵⁶,
 Andrea De Simone^{54,55}, Luigi Delle Rose^{57,58},
 Frank F Deppisch⁴², Biplab Dey⁵⁹, Miriam D Diamond⁹,
 Keith R Dienes^{35,60}, Sven Dildick⁶¹, Babette Döbrich¹⁰,
 Marco Drewes⁶², Melanie Eich³⁴, M ElSawy^{63,64},
 Alberto Escalante del Valle⁶⁵, Gabriel Facini⁴²,
 Marco Farina⁶⁶, Jonathan L Feng⁶⁷, Oliver Fischer²⁵,
 H U Flaecher³⁶, Patrick Foldenauer⁶⁸, Marat Freytsis^{13,69,70},
 Benjamin Fuks^{71,72}, Iftah Galon⁷³, Yuri Gershtein⁷⁴,
 Stefano Giagu⁷⁵, Andrea Giammanco⁶²,
 Vladimir V Gligorov⁷⁶, Tobias Golling⁷⁷,
 Sergio Grancagnolo⁷⁸, Giuliano Gustavino⁷⁹,
 Andrew Haas⁸⁰, Kristian Hahn⁸¹, Jan Hajer⁶²,
 Ahmed Hammad⁸², Lukas Heinrich¹⁰, Jan Heisig⁶²,
 J C Helo⁸³, Gavin Hesketh⁴², Christopher S Hill¹,
 Martin Hirsch⁸⁴, M Hohlmann⁸⁵, Tova Holmes⁸⁶,
 W Hulsbergen⁵², John Huth⁴³, Philip Ilten⁸⁷,
 Thomas Jacques⁵⁴, Bodhitha Jayatilaka³⁰,
 Geng-Yuan Jeng⁶⁰, K A Johns³⁵, Toshiaki Kaji⁸⁸,
 Gregor Kasieczka³⁴, Yevgeny Kats³³, Małgorzata Kazana⁸⁹,
 Henning Keller⁹⁰, Maxim Yu Khlopov^{91,92}, Felix Kling⁶⁷,
 Ted R Kolberg²⁷, Igor Kostiu^{52,93}, Emma Sian Kuwertz¹⁰,
 Audrey Kvam²⁹, Greg Landsberg⁹⁴, Gaia Lanfranchi^{10,95},
 Iñaki Lara⁹⁶, Alexander Ledovskoy⁹⁷, Dylan Linthorne⁹⁸,
 Jia Liu⁸⁶, Iacopo Longarini⁹⁹, Steven Lowette³⁷,
 Henry Lubatti²⁹, Margaret Lutz³⁹, Jingyu Luo¹⁰⁰,
 Judita Mamužić⁸⁴, Matthieu Marinangeli¹⁰¹,
 Alberto Mariotti³⁷, Daniel Marlow¹⁰⁰, Matthew McCullough¹⁰,
 Kevin McDermott⁴, P Mermod⁷⁷, David Milstead¹⁰²,
 Siddharth Mishra-Sharma⁸⁰, Vasiliki A Mitsou⁸⁴,
 Javier Montejo Berlingen¹⁰, Filip Moortgat^{10,103},
 Alessandro Morandini^{54,55}, Alice Polyxeni Morris⁴²,
 David Michael Morse¹⁰⁴, Stephen Mrenna³⁰,
 Benjamin Nachman¹⁰⁵, Miha Nemevšek¹⁰⁶, Fabrizio Nesti¹⁰⁷,
 Christian Ohm^{108,109}, Silvia Pascoli¹¹⁰, Kevin Pedro³⁰,
 Cristián Peña^{30,111}, Karla Josefina Pena Rodriguez³⁴,
 Jónatan Piedra¹¹², James L Pinfold¹¹², Antonio Policicchio⁷⁵,
 Goran Popara¹¹³, Jessica Prisciandaro⁶², Mason Proffitt²⁹,
 Giorgia Rauco⁵⁶, Federico Redi¹⁰¹, Matthew Reece⁴³,
 Allison Reinsvold Hall³⁰, H Rejeb Sfar⁵³, Sophie Renner¹¹⁴,
 Dean Robinson^{105,115}, Amber Roepe⁷⁹, Manfredi Ronzani⁸⁰,

Ennio Salvioni¹¹⁶, **Arka Santra**⁸⁴, **Ryu Sawada**¹¹⁷,
Jakub Scholtz¹¹⁰, **Philip Schuster**⁴⁶, **Pedro Schwaller**¹¹⁴,
Cristiano Sebastiani⁷⁵, **Sezen Sekmen**¹¹⁸,
Michele Selvaggi¹⁰, **Weinan Si**²³, **Livia Soffi**⁷⁵,
Daniel Stolarski⁹⁸, **David Stuart**⁴⁸, **John Stupak III**⁷⁹,
Kevin Sung⁸¹, **Wendy Taylor**¹¹⁹, **Sebastian Templ**⁶⁵,
Brooks Thomas¹²⁰, **Emma Torró-Pastor**²⁹,
Daniele Trocino¹⁰³, **Sebastian Trojanowski**¹²¹,
Marco Trovato¹²², **Yuhsin Tsai**¹⁶, **C G Tully**¹⁰⁰,
Tamás Álmos Vámi¹²³, **Juan Carlos Vasquez**¹²⁴,
Carlos Vázquez Sierra⁵², **K Vellidis**³⁸, **Basile Vermassen**¹⁰³,
Martina Vit¹⁰³, **Devin G E Walker**¹²⁵, **Xiao-Ping Wang**¹²²,
Gordon Watts²⁹, **Si Xie**¹¹¹, **Melissa Yexley**¹²⁶,
Charles Young⁴⁶, **Jiang-Hao Yu**^{18,127}, **Piotr Zalewski**⁸⁹ and
Yongchao Zhang¹²⁸

¹ Department of Physics, The Ohio State University, 191 W. Woodruff Ave., Columbus, OH 43210-1117, United States of America

² Department of Physics, Duke University, 120 Science Drive, Durham, NC 27710, United States of America

³ Physikalisches Institut, Ruprecht-Karls-Universität Heidelberg, Im Neuenheimer Feld 226, D-69120, Heidelberg, Germany

⁴ Cornell University, 245 East Avenue, Ithaca, NY 14853, United States of America

⁵ Instituto Galego de Física de Altas Enerxías, U. Santiago de Compostela, IGFAE, Rúa de Xoaquín Díaz de Rábago, s/n, E-15782 Santiago de Compostela, Spain

⁶ Departamento de Ciencias, Facultad de Artes Liberales, Universidad Adolfo Ibáñez, Diagonal Las Torres 2640, Santiago, Chile

⁷ Department of Physics, National Taiwan University, Taipei 10617, Taiwan

⁸ Instituto de Física, Pontificia Universidad Católica de Chile, Avenida Vicuña mackenna 4860, Santiago, Chile

⁹ Department of Physics, University of Toronto, 60 St. George Street, Toronto, ON M5S 1A7, Canada

¹⁰ CERN, Esplanade des particules 1, Geneva, Switzerland

¹¹ Department of Theoretical Physics, Tata Institute of Fundamental Physics, Homi Bhabha Road, Mumbai 400005, India

¹² Department of Physics, University of Cincinnati, 400 Geology/Physics Bldg., Cincinnati, OH 45221, United States of America

¹³ School of Natural Sciences, Institute for Advanced Study, 1 Einstein Drive, Princeton, NJ 08540, United States of America

¹⁴ Laboratoire de Physique Subatomique et de Cosmologie, Université Grenoble-Alpes, CNRS/IN2P3, 53 Avenue des Martyrs, F-38026 Grenoble, France

¹⁵ Universidade Federal do ABC, Av. dos Estados, 5001, Santo Andre, 09210-580 SP, Brazil

¹⁶ Maryland Center for Fundamental Physics, Department of Physics, University of Maryland, College Park, MD 20742, United States of America

¹⁷ Faculty of Physics, Ludwig-Maximilians-Universität München, Schellingstraße 4, D-80799 Munich, Germany

¹⁸ Amherst Center for Fundamental Interactions, Department of Physics, University of Massachusetts Amherst, Lederle Graduate Research Center 416, Amherst, MA 01003 United States of America

¹⁹ T.-D. Lee Institute and School of Physics and Astronomy, Shanghai Jiao Tong University, 800 Dongchuan Road, Shanghai, 200240, People's Republic of China

²⁰ Department of Physics, McGill University, Montreal, QC H3A 2T8, Canada

²¹ University of Illinois at Urbana-Champaign, 1110 W. Green St., Urbana, IL 61801, United States of America

²² Harvey Mudd College, 301 Platt Blvd., Claremont, CA 91711, United States of America

²³ University of California, Riverside, 900 University Avenue, Riverside, CA 92521, United States of America

²⁴ Universita' di Roma Tre and INFN, via della Vasca Navale 84, I-00146 Roma, Italy

²⁵ Institute for Nuclear Physics, Karlsruhe Inst. of Technology, Hermann-von-Helmholtz-Platz 1, D-76344 Eggenstein-Leopoldshafen, Germany

²⁶ Institute for Theoretical Particle Physics (TTP), Karlsruhe Institute of Technology, Engesserstraße 7, D-76128 Karlsruhe, Germany

²⁷ Florida State University, 77 Chieftan Way, Tallahassee, FL 32306, United States of America

²⁸ Faculty of Physics, Ludwig-Maximilians Universität München, Schellingstrasse 4, D-80799 Munich, Germany

²⁹ University of Washington, Seattle, 1410 NE Campus Parkway Seattle, WA 98195, United States of America

³⁰ Fermi National Accelerator Laboratory, PO Box 500, Batavia, IL 60510, United States of America

³¹ High Energy Physics Group, Blackett Laboratory, Imperial College, Prince Consort Road, London SW7 2AZ, United Kingdom

³² National Centre for Nuclear Research (NCBJ), Hoża 69, 00-681 Warsaw, Poland

³³ Department of Physics, Ben-Gurion University, 1 Ben-Gurion Boulevard Beer-Sheva, Beer-Sheva 8410501, Israel

³⁴ Institut für Experimentalphysik, Universität Hamburg, Luruper Chaussee 149, D-22761 Hamburg, Germany

³⁵ University of Arizona, 118 E. 4th St., Tucson AZ 85721, United States of America

³⁶ University of Bristol, HH Wills Physics Laboratory, Tyndall Avenue, Bristol, BS8 1TL, United Kingdom

³⁷ IHE, Vrije Universiteit Brussel, Pleinlaan 2, B-1050 Brussels, Belgium

³⁸ Physics Department, National and Kapodistrian University of Athens, Panepistimioupoli, Zografou, 15784, Greece

³⁹ University of Massachusetts, Amherst, 1126 LGRT, University of Massachusetts, Amherst, MA 01003-9337, United States of America

⁴⁰ Department of Physics, Division of Particle Physics, Lund University, Box 118, SE-221 00 Lund, Sweden

⁴¹ Leinweber Center for Theoretical Physics, University of Michigan, 450 Church Street, Ann Arbor MI-48109, United States of America

⁴² University College London, Gower Street WC1E 6BT London, United Kingdom

⁴³ Department of Physics, Harvard University, 17 Oxford Street, Cambridge, MA, 02138

⁴⁴ Department of Physics and Astronomy, Iowa State University, 2323 Osborn Drive, Physics 0012, Ames, IA 50011-3160, United States of America

⁴⁵ Department of Physics, University of California, Davis, One Shields Avenue, Davis, CA 95616, United States of America

⁴⁶ SLAC National Accelerator Laboratory, 2575 Sand Hill Rd, Menlo Park, CA 94025, United States of America

⁴⁷ Physics Department, Stanford University, 450 Serra Mall, Stanford, CA 94305, United States of America

⁴⁸ Department of Physics, Broida Hall, University of California, Santa Barbara, CA 93106, United States of America

⁴⁹ INFN Sezione di Genova, Via Dodecaneso, 33, I-16146—Genova, Italy

⁵⁰ Institut Pluridisciplinaire Hubert Curien (IPHC), Département Recherches Subatomiques, Université de Strasbourg/CNRS-IN2P3, 23 Rue du Loess, F-67037 Strasbourg, France

⁵¹ Department of Physics & Astronomy, The University of British Columbia, 2329 West Mall, Vancouver, BC V6T 1Z4, Canada

⁵² Nikhef National Institute for Subatomic Physics, Science Park 105, 1098 XG Amsterdam, The Netherlands

⁵³ Antwerp University, Prinsstraat 13, B-2000 Antwerpen, Belgium

⁵⁴ SISSA, Via Bonomea 265, Trieste, I-34136, Italy

⁵⁵ INFN Sezione di Trieste, via Bonomea 265, I-34136 Trieste, Italy

⁵⁶ University of Zürich, Winterthurerstrasse 190, 8057 Zurich, Switzerland

⁵⁷ Department of Physics and Astronomy, University of Florence, Via G. Sansone 1, 50019 Sesto Fiorentino, Italy

⁵⁸ University of Southampton, Highfield, Southampton SO17 1BJ, United Kingdom

⁵⁹ Institute of Particle Physics, Central China Normal University, Wuhan, Hubei, People's Republic of China

⁶⁰ Department of Physics, University of Maryland, 4296 Stadium Drive, College Park, MD 20742, United States of America

⁶¹ Texas A&M University, 4242 TAMU College Station, TX 77843-4242, United States of America

⁶² Centre for Cosmology, Particle Physics and Phenomenology (CP3), Université catholique de Louvain, Chemin du Cyclotron 2, B-1348, Louvain-la-Neuve, Belgium

⁶³ Basic Science Department, Faculty of Engineering, The British University in Egypt, El Sherouk City, Misr-Ismalia Road, Postal No 11837, PO Box 43, Cairo, Egypt

⁶⁴ Physics Department, Faculty of Science, Beni Suef University, Qism Bani Sweif, Bani Sweif, Beni Suef Governorate, Egypt

⁶⁵ Institute of High Energy Physics, Austrian Academy of Sciences, Nikolsdorfer Gasse 18, A-1050 Vienna, Austria

⁶⁶ Stony Brook University, Stony Brook, NY 11794-3840, United States of America

⁶⁷ Department of Physics and Astronomy, University of California, Irvine, CA 92697-4575, United States of America

⁶⁸ Institut für Theoretische Physik, Universität Heidelberg, Philosophenweg 16, D-69120 Heidelberg, Germany

⁶⁹ Raymond and Beverly Sackler School of Physics and Astronomy, Tel Aviv University, Tel-Aviv 69978, Israel

⁷⁰ Institute of Theoretical Science, University of Oregon, Eugene, OR 97403, United States of America

⁷¹ LPTHE, UMR 7589, Sorbonne Université et CNRS, 4 Place Jussieu, F-75252 Paris Cedex 05, France

⁷² Institut Universitaire de France, 103 boulevard Saint-Michel, F-75005 Paris, France

⁷³ New High Energy Theory Center, Rutgers, The State University of New Jersey, Piscataway, NJ 08854-8019, United States of America

⁷⁴ Rutgers University, 136 Frelinghuysen Rd, Piscataway, NJ 08854 United States of America

⁷⁵ Sapienza Università di Roma and INFN Roma1, P.le A. Moro 5, Roma, I-00185, Italy

⁷⁶ LPNHE, Sorbonne Université, Paris Diderot Sorbonne Paris Cité, CNRS/IN2P3, Barre 12-22, 4 Place Jussieu, F-75252 Paris Cedex 05, France

⁷⁷ Département de Physique Nucléaire et Corpusculaire, Université de Genève, 24, quai Ernest-Ansermet CH-1211 Genève 4, Switzerland

⁷⁸ Humboldt-Universität, Newtonstraße 15, D-12489 Berlin, Germany

⁷⁹ The University of Oklahoma, 440 W. Brooks St. Norman, OK 73019, United States of America

⁸⁰ New York University, 726 Broadway, New York, NY 10003, United States of America

⁸¹ Northwestern University, 2145 Sheridan Rd, Evanston, IL 60208, United States of America

⁸² Departement of Physics, University of Basel, Klingelbergstrasse 82, 4056, Basel, Switzerland

⁸³ Departamento de Física y Astronomía, Facultad de Ciencias, Universidad de La Serena, Avenida Cisternas 1200, La Serena, Chile

⁸⁴ Instituto de Física Corpuscular/Consejo Superior de Investigaciones Científicas—University of Valencia (IFIC/CSIC—UV), Carrer del Catedrático José Beltrán Martínez, 2, E-46980 Paterna, València, Spain

⁸⁵ Dept. of Aerospace, Physics & Space Sciences, Florida Institute of Technology, 150 W. University Blvd., Melbourne, FL 32901, United States of America

⁸⁶ Enrico Fermi Institute, University of Chicago, 5640 S. Ellis Avenue, RI-183, Chicago, IL 60637, United States of America

⁸⁷ University of Birmingham, Edgbaston, Birmingham B15 2TT, United Kingdom

⁸⁸ Waseda University, Ookubo 3-4-1, Shinjuku-ku, Tokyo 169-8555, Japan

⁸⁹ National Centre for Nuclear Research (NCBJ), Andrzej Sołtana 7, 05-400 Otwock-Świerk, Poland

⁹⁰ RWTH Aachen University, Otto-Blumenthal-Straße, D-52074 Aachen, Germany

⁹¹ National Research Nuclear University MEPhI, Kashirskoe chaussee 31, Moscow 115409, Russia

⁹² APC Laboratory, 10, rue Alice Domon et Léonie Duquet, F-75205, Paris Cedex 13, France

⁹³ Institute for Nuclear Research of the National Academy of Sciences (KINR), 47 Nauky Avenue, 03680, Kyiv, Ukraine

⁹⁴ Brown University, 182 Hope St, Providence, RI 02912, United States of America

⁹⁵ Laboratori Nazionali di Frascati, INFN, via E. Fermi 40, I-00044 Frascati RM, Italy

⁹⁶ Instituto de Física Teórica (UAM-CSIC), C/Nicolás Cabrera 13-15 Campus de Cantoblanco, E-28049 Madrid, Spain

⁹⁷ Physics Department, University of Virginia, PO Box 400714, 382 McCormick Rd, Charlottesville, VA, 22904-4714, United States of America

⁹⁸ Carleton University, 1125 Colonel By Dr, Ottawa, ON K1S 5B6, Canada

⁹⁹ Università degli studi di Roma ‘La Sapienza’ and INFN Roma, Dipartimento di Fisica G. Marconi. Piazzale Aldo Moro 5, I-00185, Roma, Italy

¹⁰⁰ Physics Department, Washington Road, Princeton University, Princeton, NJ 08544, United States of America

¹⁰¹ École Polytechnique Fédérale de Lausanne, Route Cantonale, 1015 Lausanne, Switzerland

¹⁰² Fysikum, Stockholms Universitet, Roslagstullsbacken 21, SE-114 21 Stockholm, Sweden

¹⁰³ Department of Physics and Astronomy, University of Ghent, Proeftuinstraat 86, B-9000 Ghent, Belgium

¹⁰⁴ Northeastern University, 360 Huntington Ave., Boston, MA 02115, United States of America

¹⁰⁵ Physics Division, Lawrence Berkeley National Laboratory, Berkeley, 1 Cyclotron Road, CA 94720, United States of America

¹⁰⁶ Jožef Stefan Institute, Jamova 39, Ljubljana 1000, Slovenia

¹⁰⁷ Università dell’Aquila, via Vetoio, L’Aquila, I-67100, Italy

¹⁰⁸ KTH Royal Institute of Technology, AlbaNova Universitetscentrum, KTH, Physics Department, SE-106 91 Stockholm, Sweden

¹⁰⁹ Oskar Klein Centre for Cosmoparticle Physics, Stockholm University, SE-106 91 Stockholm, Sweden

¹¹⁰ IPPP, Department of Physics, Durham University, South Road, Durham, DH1 3LE, United Kingdom

¹¹¹ California Institute of Technology, 1200 E California Blvd, MC 256-48, Pasadena, CA, 91125, United States of America

¹¹² Physics Department, University of Alberta, Edmonton, AB T6G 2E1, Canada

¹¹³ Ruder Bošković Institute, Bijenička cesta 54, 10000 Zagreb, Croatia

¹¹⁴ PRISMA Cluster of Excellence & Mainz Inst. for Theoretical Physics, Johannes Gutenberg Univ., Staudingerweg 7, D-55128 Mainz, Germany

¹¹⁵ Santa Cruz Institute for Particle Physics and Department of Physics, University of California Santa Cruz, Santa Cruz, CA 95064, United States of America

¹¹⁶ Technical University of Munich, Physics Department, James-Franck-Strasse 1, D-85748 Garching, Germany

¹¹⁷ International Center for Elementary Particle Physics, The University of Tokyo, Hongo 7-3-1, Bunkyo-ku, Tokyo 113-0033, Japan

¹¹⁸ Department of Physics, Kyungpook National University, 80 Daehakro, Bukgu, Daegu 41566, Republic of Korea

¹¹⁹ York University, 4700 Keele St, Toronto, ON M3J 1P3, Canada

¹²⁰ Lafayette College, 730 High St. Easton, PA 18042, United States of America

¹²¹ Consortium for Fundamental Physics, School of Mathematics and Statistics, University of Sheffield, Hounsfield Road, Sheffield S3 7RH, United Kingdom

¹²² Argonne National Laboratory, 9700 Cass Avenue, Lemont, IL 60439, United States of America

¹²³ Wigner Research Centre for Physics, Konkoly-Thege Miklós út 29-33, Budapest 1121, Hungary

¹²⁴ Universidad Técnica Federico Santa María, Avenida España 1680, Valparaíso, 2340000, Chile

¹²⁵ Department of Physics and Astronomy, Dartmouth College, Hanover, NH 03755 United States of America

¹²⁶ Lancaster University, Lancaster LA1 4YW, United Kingdom

¹²⁷ CAS Key Laboratory of Theoretical Physics, Institute of Theoretical Physics, Chinese Academy of Sciences, Beijing 100190, People's Republic of China

¹²⁸ Department of Physics and McDonnell Center for the Space Sciences, Washington University, St. Louis, MO 63130, United States of America

E-mail: juliette.alimena@cern.ch, j.beacham@cern.ch, martino.borsato@cern.ch, yangyang.cheng@cornell.edu, xabier.cid.vidal@cern.ch, gcottin@phys.ntu.edu.tw, dcurtin@physics.utoronto.ca, deroeck@mail.cern.ch, desai@theory.tifr.res.in, jaredaevans@gmail.com, knapen@ias.edu, sabine.kraml@lpsc.in2p3.fr, andre.lessa@ufabc.edu.br, zliuphys@umd.edu, sascha.mehlhase@cern.ch, mjrm@physics.umass.edu, hrussell@cern.ch, sheltonj@illinois.edu, bshuve@g.hmc.edu, monica.verducci@cern.ch, jose.zurita@kit.edu, tadams@hep.fsu.edu, michael.adersberger@cern.ch, Cristiano.Alpigiani@cern.ch, apresyan@fnal.gov, robert.bainbridge@cern.ch, varvara.batozskaya@cern.ch, beauches@post.bgu.ac.il, lisa.benato@cern.ch, simon.berlendis@cern.ch, eshwen.bhal@cern.ch, freya.blekman@cern.ch, Christiclip@Hotmail.Com, Jamie.Boyd@cern.ch, bbrau@physics.umass.edu, lene.bryngemark@cern.ch, oliver.buchmueller@cern.ch, buschman@umich.edu, william.buttinger@cern.ch, mario.campanelli@cern.ch, ccesarotti@g.harvard.edu, cchen23@iastate.edu, cheng@physics.ucdavis.edu, sanha@slac.stanford.edu, mcitron@cern.ch, andrea.coccato@cern.ch, victor.coco@cern.ch, eric.conte@iphe.cnrs.fr, fcormier@cern.ch, l.corpe@cern.ch, nraig@physics.ucsb.edu, yanou.cui@ucr.edu, elena.dall'occo@cern.ch, carlod@physics.umass.edu, mohamed.anwar@cern.ch, alessandro.davoli@sissa.it, decosa@cern.ch, andrea.desimone@sissa.it, luigi.dellerose@fi.infn.it, f.deppisch@ucl.ac.uk, biplab.dey@cern.ch, mdiamond@physics.utoronto.ca, dienes@email.arizona.edu, sven.dildick@cern.ch, babette.dobrich@cern.ch, marco.drewes@uclouvain.be, melanie.eich@desy.de, mai.el.sawy@cern.ch, a.escalante.del.valle@cern.ch, [7](mailto:gabriel.</p>
</div>
<div data-bbox=)

facini@cern.ch, farina.phys@gmail.com, jlf@uci.edu, oliver.fischer@kit.edu, henning.flacher@cern.ch, foldenauer@thphys.uni-heidelberg.de, freytsis@ias.edu, fuks@lpthe.jussieu.fr, iftah.galon@physics.rutgers.edu, gershtein@physics.rutgers.edu, stefano.giagu@cern.ch, andrea.giammanco@cern.ch, vgligor@lpnhe.in2p3.fr, Tobias.Golling@unige.ch, sergio.grancagnolo@cern.ch, sergio.giuliano.gustavino@cern.ch, andy.haas@nyu.edu, kristian.hahn@northwestern.edu, jan.hajer@uclouvain.be, ahmed.hammad@unibas.ch, lukas.heinrich@cern.ch, jan.heisig@uclouvain.be, jchelo@userena.cl, gavin.hesketh@ucl.ac.uk, chill@physics.osu.edu, mahirsch@ific.uv.es, hohlmann@fit.edu, wouter@nikhef.nl, huth@physics.harvard.edu, philten@cern.ch, t.d.jacques@gmail.com, boj@fnal.gov, geng-yuan.jeng@cern.ch, johns@physics.arizona.edu, toshiaki.kaji@cern.ch, gregor.kasieczka@cern.ch, katsye@bgu.ac.il, malgorzata.kazana@cern.ch, henning.keller@cern.ch, khlopov@apc.in2p3.fr, fkling@uci.edu, tkolberg@hep.fsu.edu, igor.kostiuk@cern.ch, emma.sian.kuwertz@cern.ch, akvam@cern.ch, landsberg@hep.brown.edu, Gaia.Lanfranchi@lnf.infn.it, inaki.lara@csic.es, alexander.ledovskoy@cern.ch, dylan.linthorne@carleton.ca, liuj1@uchicago.edu, iacopo.longarini@cern.ch, steven.lowette@cern.ch, lubatti@u.washington.edu, margaret.susan.lutz@cern.ch, jingyu@princeton.edu, judita.mamuzic@cern.ch, matthieu.marinangeli@cern.ch, alberto.mariotti@vub.be, marlow@princeton.edu, matthew.mccullough@cern.ch, kpm82@cornell.edu, philippe.mermod@cern.ch, milstead@fysik.su.se, sm8383@nyu.edu, vasiliki.mitsou@cern.ch, jmontejo@cern.ch, filip.moortgat@cern.ch, amorandi@sissa.it, alice.morris@cern.ch, david.michael.morse@cern.ch, mrenna@fnal.gov, bpnachman@lbl.gov, miha.nemevsek@ijs.si, nesti@aquila.infn.it, ohm@cern.ch, silvia.pascoli@durham.ac.uk, pedrok@cern.ch, cmorgoth@fnal.gov, karla.pena@cern.ch, piedra@cern.ch, jpinfo@ualberta.ca, antonio.policicchio@cern.ch, gpopara@irb.hr, jessica.prisciandaro@cern.ch, mproffit@cern.ch, giorgia.rauco@cern.ch, federico.redi@cern.ch, mreece@g.harvard.edu, ahall@fnal.gov, haifa.rejeb.sfar@cern.ch, sorenner@uni-mainz.de, drobinson@lbl.gov, amber.rebecca.roepe@cern.ch, manfredi.ronzani@cern.ch, ennio.salvioni@tum.de, arka.santra@cern.ch, sawada@icepp.s.u-tokyo.ac.jp, jakubscholtz@gmail.com, schuster@slac.stanford.edu, pedro.schwaller@uni-mainz.de, cristiano.sebastiani@cern.ch, ssekmen@cern.ch, michele.selvaggi@cern.ch, weinan.si@cern.ch, livia.soffi@cern.ch, stolar@physics.carleton.ca, stuart@physics.ucs.edu, john.stupak@ou.edu, Kevin.Kai.Hong.Sung@cern.ch, taylorw@yorku.ca, sebastian.templ@cern.ch, thomasbd@lafayette.edu, Emma.Torro.Pastor@cern.ch, daniele.trocino@cern.ch, Sebastian.Trojanowski@ncbj.gov.pl, marco.trovato@cern.ch, yhtsai@umd.edu, cgtully@princeton.edu, tamás.almos.vami@cern.ch, juan.vasquezcar@usm.cl, carlos.vazquez@cern.ch, konstantinos.vellidis@cern.ch, basile.vermassen@cern.ch, martina.vit@cern.ch, devin.g.walker@dartmouth.edu, xia.wang@anl.gov, gwatts@uw.edu, sixie@caltech.edu, m.yexley@cern.ch, young@slac.stanford.edu, jhyu@itp.ac.cn, piotr.zalewski@ncbj.gov.pl and yongchao2228@gmail.com

Received 23 May 2019

Accepted for publication 17 September 2019

Published 2 September 2020



Abstract

Particles beyond the Standard Model (SM) can generically have lifetimes that are long compared to SM particles at the weak scale. When produced at experiments such as the Large Hadron Collider (LHC) at CERN, these long-lived particles (LLPs) can decay far from the interaction vertex of the primary proton–proton collision. Such LLP signatures are distinct from those of promptly decaying particles that are targeted by the majority of searches for new physics at the LHC, often requiring customized techniques to identify, for example, significantly displaced decay vertices, tracks with atypical properties, and short track segments. Given their non-standard nature, a comprehensive overview of LLP signatures at the LHC is beneficial to ensure that possible avenues of the discovery of new physics are not overlooked. Here we report on the joint work of a community of theorists and experimentalists with the ATLAS, CMS, and LHCb experiments—as well as those working on dedicated experiments such as MoEDAL, milliQan, MATHUSLA, CODEX-b, and FASER—to survey the current state of LLP searches at the LHC, and to chart a path for the development of LLP searches into the future, both in the upcoming Run 3 and at the high-luminosity LHC. The work is organized around the current and future potential capabilities of LHC experiments to generally discover new LLPs, and takes a signature-based approach to surveying classes of models that give rise to LLPs rather than emphasizing any particular theory motivation. We develop a set of simplified models; assess the coverage of current searches; document known, often unexpected backgrounds; explore the capabilities of proposed detector upgrades; provide recommendations for the presentation of search results; and look towards the newest frontiers, namely high-multiplicity ‘dark showers’, highlighting opportunities for expanding the LHC reach for these signals.

Keywords: beyond the Standard Model, long-lived particles, Large Hadron Collider, high-luminosity LHC, collider phenomenology, high-energy collider experiments

1. Introduction

Document editors: James Beacham, Brian Shuve

Particles in the Standard Model (SM) have lifetimes spanning an enormous range of magnitudes, from the Z boson ($\tau \sim 2 \times 10^{-25}$ s) through to the proton ($\tau \gtrsim 10^{34}$ years) and electron (stable), as shown in figure 1.

Similarly, models beyond the SM (BSM) typically predict new particles with a variety of lifetimes. In particular, new weak-scale particles can easily have long lifetimes for several reasons, including approximate symmetries that stabilize the long-lived particle (LLP), small couplings between the LLP and lighter states, and suppressed phase space available for decays. For particles moving close to the speed of light, this can lead to macroscopic,




Figure 1. Particle lifetime $c\tau$, expressed in meters, as a function of particle mass, expressed in GeV, for a variety of particles in the Standard Model [1].

detectable displacements between the production and decay points of an unstable particle for $c\tau \gtrsim 10 \mu\text{m}$ ¹²⁹.

The experimental signatures of LLPs at the Large Hadron Collider (LHC) are varied and, by nature, are often very different from signals of SM processes. For example, LLP signatures can include tracks with unusual ionization and propagation properties; small, localized deposits of energy inside of the calorimeters without associated tracks; stopped particles (SPs) that decay out of time with collisions; displaced vertices (DV) in the inner detector (ID) or muon spectrometer (MS); and disappearing, appearing, and kinked tracks. A schematic of a sampling of the variety of LLP signatures that are different from both SM signatures and the majority of BSM searches at the LHC is shown in figure 2.

Because the LLPs of the SM have masses $\lesssim 5 \text{ GeV}$ and have well-understood experimental signatures, the unusual signatures of BSM LLPs offer excellent prospects for the discovery of new physics at particle colliders. At the same time, standard reconstruction algorithms may reject events or objects containing LLPs precisely because of their unusual nature, and dedicated searches are needed to uncover LLP signals. These atypical signatures can also resemble noise, pile-up (PU), or mis-reconstructed objects in the detector; due to the rarity of such mis-reconstructions, Monte Carlo (MC) simulations may not accurately model backgrounds for LLP searches, and dedicated methods are needed to do so.

Although small compared to the large number of searches for prompt decays of new particles, many searches for LLPs at the ATLAS, CMS, and LHCb experiments at the LHC have already been performed; we refer the reader to chapter 3 for descriptions of and references to these searches. Existing LLP searches have necessitated the development of novel methods for identifying signals of LLPs, and measuring and suppressing the relevant backgrounds. Indeed, in several scenarios searches for LLPs have sensitivities that greatly

¹²⁹ Recently, a comprehensive collection of the vast array of theoretical frameworks within which LLPs naturally arise has been assembled as part of the physics case document for the proposed MATHUSLA experiment [2]. Because the focus of the current document is on the experimental signatures of LLPs and explicitly not the theories that predict them, the combination of the MATHUSLA physics case document (and the large number of references therein) and the present document can be considered, together, a comprehensive view of the present status of theoretical motivation and experimental possibilities for the potential discovery of LLPs produced at the interaction points (IPs) of the LHC.




Figure 2. Schematic of the variety of challenging, atypical experimental signatures that can result from BSM LLPs in the detectors at the LHC. Shown is a cross-sectional plane in azimuthal angle, ϕ , of a general purpose detector such as ATLAS or CMS. From [3].

exceed the search for similar, promptly decaying new particles (as is true, for example, for directly produced staus in supersymmetry [4]). The excellent sensitivity of these searches, together with the lack of a definitive signal in any prompt channels at the LHC, have focused attention on other types of LLP signatures that are not currently covered. These include low-mass LLPs that do not pass trigger or selection thresholds of current searches, high multiplicities of LLPs produced in dark-sector showers, or unusual LLP production and decay modes that are not covered by current methods. Given the excellent sensitivity of LHC detectors to LLPs, along with the potentially large production cross sections of LLPs and the enormous amount of data expected to be collected when the LHC switches to high-luminosity running in the 2020s, it is imperative that the space of LLP signatures be explored as thoroughly as possible to ensure that no signals are missed. This is particularly important now, with the recent conclusion of LHC Run 2, as new triggering strategies for LLP signatures in the upcoming Run 3 can be investigated with urgency. Moreover, decisions are currently being made about detector upgrades for Phase 2 of the LHC, and design choices should be made to ensure that sensitivity to LLPs is retained or possibly improved through high-luminosity running, as may indeed be the case for many of the plans under consideration by the main experiments.

The increased interest in LLP signatures at the LHC is naturally complementary to the recognition that there are several BSM scenarios that give rise to particles difficult to optimally detect at the LHC—either promptly decaying particles or those with naturally long lifetimes—and that are searched for or planned to be searched for in fixed-target experiments, B-factories, and beam dump experiments [5–7].

The growing theoretical and experimental interest in LLPs has been mirrored by an increased activity in proposals for searches for LLPs produced at the main LHC IPs—either within the existing ATLAS, CMS, and LHCb collaborations or with new, dedicated detectors—new experimental analyses, and meetings to communicate results and discuss new ideas.

Workshops focused on LLPs at the University of Massachusetts, Amherst¹³⁰, in November of 2015; Fermilab; and KITP (UCSB)¹³¹ in May of 2016, among others, highlighted the need for a community-wide effort to map the current space of both theoretical models for LLPs and the atypical experimental signatures that could be evidence of LLPs, assess the coverage of current experimental methods to these models, and identify areas where new searches are required. Additionally, the work presented in these meetings underscored the importance of presenting the results of experimental searches in a manner that allows for their application to different models, and generated new ideas for designing analyses with the goal of minimizing model dependence. Such largely model-independent presentation makes current searches more powerful by increasing their applicability to new scenarios, while reducing redundancies in searches and ensuring that gaps in coverage are identified and addressed. This task extends beyond the purview of any particular theoretical model or experiment, and requires an effort across collaborations to address the needs of the LLP community and illuminate a path forward.

This flurry of activity eventually coalesced in the establishment of a more central and regular platform—the LHC LLP Community—for experimentalists at the LHC and those in the theoretical and phenomenological communities to exchange ideas about LLP searches to ensure the full discovery potential of the LHC. This began with a mini-workshop at CERN in May of 2016¹³² and has continued with workshops in April of 2017 at CERN¹³³, October of 2017 at ICTP Trieste¹³⁴, May of 2018 again at CERN¹³⁵, at Nikhef, in Amsterdam, in October of 2018¹³⁶, May of 2019 at CERN¹³⁷ and November of 2019 at Ghent University¹³⁸.

This is the work undertaken by the LHC LLP Community and presented in this document. Based on the most pressing needs identified by the community, we organize the work of this initiative into a few key realms:

- *Simplified models*: We seek to identify a minimal (but expandable) set of simplified models that capture, with a very limited number of free parameters, the most important LLP signatures motivated by theory and accessible at the LHC. The simplified models approach has been successfully applied to models such as supersymmetry (SUSY) and dark matter (DM), and proposals exist for LLP simplified models in particular contexts. We aim to provide a basis of models that serves as a focal point for the other studies performed by the community, as well as a library that can be used in simulating LLP signal events, to allow for a common grammar to better understand how current and future searches cover LLP signature space.

¹³⁰ ‘LHC Searches for Long-Lived BSM Particles: Theory Meets Experiment’, <https://physics.umass.edu/acf1/seminars-and-workshops/lhc-searches-for-long-lived-bsm%2Dparticles-theory-meets-experiment>.

¹³¹ ‘Experimental Challenges for the LHC Run II’, <http://online.kitp.ucsb.edu/online/experlhc16/>.

¹³² ‘LHC Long-Lived Particle Mini-Workshop’, https://indico.cern.ch/e/LHC LLP_2016.

¹³³ ‘Searches for LLPs at the LHC: First workshop of the LHC LLP Community’, https://indico.cern.ch/e/LHC LLP_April_2017.

¹³⁴ ‘Searches for LLPs at the LHC: Second workshop of the LHC LLP Community’, https://indico.cern.ch/e/LHC LLP_October_2017.

¹³⁵ ‘Searching for LLPs at the LHC: Third workshop of the LHC LLP Community’, https://indico.cern.ch/e/LHC LLP_May_2018.

¹³⁶ ‘Searching for LLPs at the LHC: Fourth workshop of the LHC LLP Community’, https://indico.cern.ch/e/LHC LLP_October_2018.

¹³⁷ ‘Searching for LLPs at the LHC: Fifth Workshop of the LHC LLP Community’, https://indico.cern.ch/e/LHC LLP_May_2019.

¹³⁸ ‘Searching for LLPs at the LHC: Sixth Workshop of the LHC LLP Community’, https://indico.cern.ch/e/LHC LLP_November_2019.

- *Experimental coverage:* In spite of the many successful LLP searches undertaken by the ATLAS, CMS, and LHCb experiments, there remains a need for a systematic study of the complementary coverage of LLP searches to the parameter spaces of LLP models. Having developed a simplified model basis, we provide a comprehensive overview of the sensitivity of existing searches, highlighting gaps in coverage and high-priority searches to be undertaken in the future.
- *Backgrounds to LLP searches:* We provide a summary and analysis of backgrounds for LLP signals at the LHC, sources of which can be rare, unexpected, and largely irrelevant for searches for prompt BSM particles, and thus not fully well understood. We assemble the collected knowledge and experience of backgrounds to prior searches with the intention of providing insight into the opportunities and challenges of searching for LLP signatures.
- *Upgrades and triggering strategies:* We discuss the prospects for LLP searches with upgraded detectors for Phase 2 of LHC running, with a focus on how upgrades can offer new sensitivity to LLPs as well as mitigate the effects of PU. New opportunities for improving sensitivity of triggers and searches to LLPs are additionally presented for the upgrades planned for Run 3. This is tied to the crucial question of triggers for LLPs; we discuss the performance of current triggers for LLPs, as well as the effects of future upgrades to the trigger system. Most importantly, we identify a few concrete upgrade studies that should be performed by the experiments that are of prime interest to the community.
- *Reinterpretation of LLP searches:* Due to the non-standard nature of the objects used in analyses, LLP searches are notoriously hard to reinterpret for models beyond those considered by the experimental collaborations. Designing searches and presenting search results in a way that is broadly applicable to current and yet-to-be-developed LLP models is crucial to the impact and legacy of the LLP search program. We discuss the reinterpretation of the LLP searches by means of concrete examples to illustrate specific challenges and, based on the lessons learned from this procedure, we provide recommendations on the presentation of LLP experimental search results.
- *Dark showers:* Current LLP search strategies have limited sensitivity to models where the LLPs are very soft, highly collimated, and come in large multiplicities, as can occur in models of dark-sector showers. We report on recent progress in theoretically parameterizing the space of dark-shower models and signatures, as well as experimental searches to uncover these signals.

Finally, we provide information about current and proposed experiments to search for LLPs at the LHC via dedicated detectors. These include the MoEDAL monopole search, the milliQan milli-charged particle (mCP) experiment, the MATHUSLA surface detector for ultra-LLPs, the CODEX-b proposal for a new detector near LHCb, and the FASER proposal for a long, narrow detector located in the forward direction well downstream one of the collision points.

This is the first report of the LHC LLP Community initiative, and is expected to be followed by future reports as our collective understanding of these signatures as a means of discovering new physics at the LHC evolves.

2. Simplified models yielding LLPs

Chapter editors: James Beacham, Giovanna Cottin, David Curtin, Jared Evans, Zhen Liu, Michael Ramsey-Musolf, Jessie Shelton, Brian Shuve

Contributors: Oliver Buchmueller, Alessandro Davoli, Andrea De Simone, Kristian Hahn, Jan Heisig, Thomas Jacques, Matthew McCullough, Stephen Mrenna, Marco Trovato, Jiang-Hao Yu

Long-lived particles (LLPs) arise in many well-motivated theories of physics BSM, ranging from heavily studied scenarios such as the minimal supersymmetric SM (MSSM) [8–12] to newer theoretical frameworks such as neutral naturalness [13–15] and hidden sector DM [16–21]. Macroscopic decay lengths of new particles naturally arise from the presence and breaking of symmetries, which can be motivated by cosmology (such as DM and baryogenesis) [22–35], neutrino masses [36–52], as well as solutions to the hierarchy problem [13–15, 53–61] or in asymptotically safe completions of the SM [62]; indeed, LLPs are generically a prediction of new hidden sectors at and below the weak scale [63–71]. An extensive and encyclopedic compilation of theoretical motivations for LLPs has already been performed for the physics case of the proposed MATHUSLA experiment [2], and we refer the reader to this document and the references therein for an in-depth discussion of theoretical motivations for LLPs. Given the large number of theories predicting LLPs, however, it is clear that a comprehensive search program for LLPs is critical to fully leverage the LHC’s immense capability to illuminate the physics of the weak scale and beyond.

The simplified model framework has proven to be a highly successful approach to characterizing signals of BSM physics. Simplified models have driven the development of searches for new signatures at the LHC and allowed existing searches to be reinterpreted for many models beyond the one(s) initially targeted in the analysis. Comprehensive simplified model programs exist for scenarios featuring prompt decays of new particles [72–78] or DM produced at colliders [79–90]. Simplified models are so successful because the majority of search sensitivity is driven by only a few broad aspects of a given BSM signature, such as the production process, overall production rate, and decay topology. Meanwhile, the sensitivity of searches is typically insensitive to other properties such as the spin of the particles involved [91–94].

To extend the simplified model approach to LLP signatures in a systematic way, we develop a proposal for a set of simplified models which aims to ensure that experimental results can be characterized as follows: (i) *powerful*, covering as much territory in model space as possible; (ii) *efficient*, reducing unnecessary redundancy among searches; (iii) *flexible*, so that they are broadly applicable to different types of models; and (iv) *durable*, providing a common framework for MC simulation of signals and facilitating the communication of results of LLP searches so that they may be applied to new models for years to come. We elaborate on these goals in section 2.1. This framework helps illuminate gaps in coverage and highlight areas where new searches are needed, and we undertake such a study in chapter 3. Our efforts build on earlier work proposing simplified model programs for LLPs motivated by particular considerations such as SUSY or DM [95–100].

In our work, we concentrate on establishing an initial basis of simplified models representative of theories giving rise to final states with one or two LLPs¹³⁹. The simplified model approach is very powerful for LLP signatures: the typically lower backgrounds for displaced signatures allow searches to be highly inclusive with respect to other objects in the

¹³⁹ Some models predict moderately higher LLP multiplicities, but the coverage of such signatures from 1-2 LLP searches is good provided the LLPs do not overlap in the detector. Our proposed simplified models are not, however, representative of high-multiplicity signatures such as dark showers (see section 2.6 and chapter 7).

event or the identification of objects originating from the decay of an LLP. This enables a single analysis to have sensitivity to a wide variety of models for LLP production and decay.

We organize our simplified models in terms of LLP channels characterized by a combination of a particular LLP production mode with a particular decay mode. Because the production and decay positions of LLPs are physically distinct¹⁴⁰, it is often possible to factorize and consider separately their production and decay¹⁴¹. For each LLP channel, the lifetime of the LLP is taken to be a free parameter. We emphasize that the LLP channel defined here is *not* the same as an experimental signature that manifests in the detector: a single channel can give rise to many different signatures depending on where (or whether)¹⁴² the LLP decays occur inside the detector, while a single experimental search for a particular signature could potentially cover many simplified model channels. In this chapter, we focus on the construction and simulation of a concrete basis of LLP simplified model channels; a partial mapping of existing searches into our basis of simplified models is discussed in chapter 3, along with the highest-priority gaps in current coverage and proposals for new searches.

As discussed in the existing simplified model literature, simplified models have their own limited range of applicability [78, 86–88, 106]. For example, the presentation of search results in terms of simplified models often assume 100% branching fractions into particular final states. In a UV model where the LLP decays in a very large number of ways, none of the individual simplified model searches may be sufficient to constrain it. Similarly, if the LLP is produced in a UV model with other associated objects that spoil the signal efficiency (for example, the production of energetic, prompt objects collimated with the LLP such that the signal fails isolation or displacement criteria; this is particularly important for high-multiplicity or dark-shower scenarios, as discussed in chapter 7), then the simplified model result does not apply and a more targeted analysis is required to cover the model. Nevertheless, the simplified models framework allows us to organize possible production modes and signatures in a systematic way and identify if there are any interesting signals or parts of parameter space that are missed by current searches. Therefore, we present a proposal for simplified models here with the understanding that there exist scenarios where UV models remain important for developing searches and presenting results.

The basis of simplified models presented here is a starting point, rather than a final statement. The present goal is to provide a set of simplified models that covers the majority of the best-motivated and simplest UV models predicting LLPs, which we outline in section 2.2. Many of these contain singly and doubly produced LLPs (or in some cases, three-to-four relatively isolated LLPs, which are typically covered well by searches for 1–2 LLPs) and so we restrict our simplified model proposal to cover these multiplicities. By design, simplified models do not include all of the specific details and subtle features that may be found in a given complete model. Therefore, the provided list is meant to be

¹⁴⁰ Indeed, the decay position may be so far from the collision point that external detectors can also be used to search for ultra-long-lived neutral or mCPs [101–105].

¹⁴¹ In addition to production and decay, a third consideration is the propagation of particles through the detector. While neutral LLPs undergo straightforward propagation, states with electric or color charge (e.g. SUSY R -hadrons), or particles with exotic charges such as magnetic monopoles or quirks, typically engage in a more complicated and often very uncertain traverse through the detector. This spoils the factorization of LLP production and decay. The subtleties related to LLPs with electric or color charge is discussed more in section 2.4.3. A trickier question is how to best simulate such states: since LLPs with electric or color charge interact with the detector material, there must be an interface between the detector simulation software and the program implementing decay. This is discussed further in section 2.5.2.

¹⁴² The case of detector-stable particles is understood to be included in the simplified models by setting $c\tau \rightarrow \infty$. In this case there is manifestly no dependence on the decay mode. See section 2.3.2 for further details.

expanded to cover new or more refined models as the LLP-search program develops. For instance, extending the simplified model framework to separately treat final states with heavy-flavor particles is of great interest (in analogy with the prompt case [107–109]); see section 2.6 for a discussion of this and other limitations of the current framework along with future opportunities for expansion. High-multiplicity signatures such as dark showers or emerging jets present different experimental and theoretical issues, which are discussed in chapter 7. Finally, a broader set of simplified models may be needed to present the results of experimental searches and to allow ready application of experimental results to UV models of interest (see chapter 6).

2.1. Goals of the present simplified model framework

The purpose of the simplified model framework is to provide a simple, common language that experimentalists and theorists can use to describe theories of LLPs and the corresponding mapping between models and experimental signatures. We therefore want our simplified model space to:

1. Use a minimal but sufficient set of models to cover a wide range of the best-motivated theories of LLPs.
2. Furnish a simple map between models and signatures to enable a clear assessment of existing search coverage and possible gaps.
3. Expand flexibly when needed to incorporate theories and signatures not yet proposed.
4. Provide a concrete MC signal event generation framework for signals.
5. Facilitate the reinterpretation of searches by supplying a sufficiently varied set of standard benchmark models for which experimental efficiencies can be provided for validation purposes.

Note that points #1 and #5 are somewhat in tension with one another: we wish to have a compact set of models that can be the subject of systematic study in terms of experimental signatures, but expressing experimental results in terms of only this set of simplified models may make it challenging to reinterpret experimental searches for UV models that are not precisely described by one of the simplified models. In this section, we prioritize having a minimal set of simplified models for the purpose of studying experimental coverages and generating new search ideas, while we defer a discussion of simplified models in the presentation and reinterpretation of search results to chapter 6¹⁴³.

In the remainder of this chapter, we construct a proposal for a minimal basis of simplified models for events with one or two LLPs. We begin with a discussion of the well-motivated UV theories that predict the existence of LLPs, and identify a set of umbrella models that yield LLPs in section 2.2. We next identify the relevant (simplified) production and decay modes for LLPs in section 2.3, emphasizing that each channel for production and decay has a characteristic set of predictions for the number and nature of *prompt* accompanying objects (AOs) producing along with the LLP. In section 2.4, we combine these production and decay modes into our simplified model basis set and highlight how different umbrella models naturally populate the various LLP channels. Section 2.5 and the appendix present a framework and instructions for how the best-motivated simplified model channels can be

¹⁴³ We note that, in general, more benchmark models may be needed for enabling reliable reinterpretation than the minimal set discussed here. An example where an extended set of simplified models is used can be seen in the heavy stable charged particle (HSCP) reinterpretation in section 6.3.2 (table 5).

simulated in MC using a new model library provided in the [appendix](#). Finally, limitations of the existing framework, along with opportunities for its further development are outlined in section [2.6](#).

2.2. Existing well-motivated theories for LLPs

Here we provide a brief distillation of many of the best-motivated theories with LLPs into five over-arching categories, focusing in particular on those that give rise to single and double production of LLPs at colliders. We emphasize that each of these categories is a broad umbrella containing many different individual models containing LLPs; in many cases, the motivations and model details among theories within a particular category may be very different, but tend to predict similar types of LLPs. Additionally, the categories are not mutually exclusive, with several examples of UV models falling into one or more category. In all cases, long lifetimes typically arise from some combination of hierarchies of scales in interactions that mediate decays; small couplings; and phase space considerations (such as small mass splittings between particles or large multiplicities of final-state particles in a decay). Many of the broad theoretical motivations for LLPs have recently been summarized in the literature [\[2\]](#).

The UV umbrella models we consider are:

- *Supersymmetry-like theories (SUSY)*. This category contains models with multiple new particles carrying SM gauge charges and a variety of allowed cascade decays. Here LLPs can arise as a result of approximate symmetries (such as R -parity [\[55, 110, 111\]](#) or indeed SUSY itself in the case of gauge mediation [\[112\]](#)) or through a hierarchy of mass scales (such as highly off-shell intermediaries in split SUSY [\[113\]](#), or nearly-degenerate multiplets [\[63, 64, 114\]](#), as in anomaly-mediated SUSY breaking [\[65\]](#)). Finally, models of SUSY hidden sectors such as stealth SUSY [\[54\]](#) generically lead to LLPs. Our terminology classifies any non-SUSY models with new SM gauge-charged particles, such as composite Higgs or extra-dimensional models, under the SUSY-like umbrella because of the prediction of new particles above the weak scale with SM gauge charges. In this category, LLP production is typically dominated by SM gauge interactions, whether of the LLP itself or of a heavy parent (HP) particle that decays to LLPs.
- *Higgs-portal theories (Higgs)*. In this category, LLPs couple predominantly to the SM-like Higgs boson. This possibility is well motivated because the SM Higgs field provides one of the leading renormalizable portals for new gauge-singlet particles to couple to the SM, and the experimental characterization of the Higgs boson leaves much scope for couplings of the Higgs to BSM physics [\[115, 116\]](#). The most striking signatures here are exotic Higgs decays to low-mass particles [\[117\]](#) (as in many Hidden Valley (HV) scenarios [\[66, 67\]](#)), which can arise in models of neutral naturalness [\[13, 14, 118\]](#) and DM [\[119\]](#), as well as in more exotic scenarios such as relaxion models [\[120\]](#). The Higgs is also special in that it comes with a rich set of associated production modes in addition to the dominant gluon-fusion process, with vector-boson fusion (VBF) and Higgs-strahlung (VH) production modes allowing novel opportunities for triggering on and suppressing backgrounds to Higgs-portal LLP signatures. Indeed, in many scenarios where LLPs are produced in exotic Higgs decays, associated-production modes can be the only way of triggering on the event.
- *Gauge-portal theories (ZP)*. This category contains scenarios where new vector mediators can produce LLPs. These are similar to Higgs models, although here the vector mediator is predominantly produced from $q\bar{q}$ initial states without other associated

objects except for gluon initial-state radiation (ISR). Examples include models where both SM fermions and LLPs carry a charge associated with a new Z' (for a review, see [121]), as well as either Abelian or non-Abelian dark photon or dark Z models [122] in which the couplings of new vector bosons to the SM are mediated by kinetic mixing. Scenarios with LLPs coupled to new gauge bosons are well motivated by theories of DM, particularly models with significant self-interactions [123–125] and/or sub-weak mass scales [17, 18, 20, 126, 127].

- *DM theories.* Non-SUSY and hidden-sector DM scenarios are collected in this category, which encompasses models where the cosmological DM is produced as a final state in the collider process. Examples of multi-component DM theories include models of new electroweak multiplets [64, 128–130], strongly interacting massive particles (SIMPs) [131], iDM [132–135], models with DM coannihilation partners [98, 136–141] (including scenarios where the coannihilation partners are out of chemical equilibrium, giving distinctly predictions for the relic abundance [142–145]), and non-thermal ‘freeze-in’ scenarios [146–153]. In many of these models, the collider phenomenology and LLP lifetime can be tied to the DM relic abundance [135, 146, 154]. For LLPs decaying inside the detector, an important feature distinguishing this category from the Higgs and gauge scenarios above is that an explicit detector-level signature of a DM candidate, i.e. missing energy (\cancel{E}_T), is a necessary and irreducible component [30, 66, 67, 100, 133, 135, 155–157].
- *Heavy neutrino theories ($RH\nu$).* The see-saw mechanism of SM neutrino mass generation predicts new right-handed neutrino (RHN) states [158–162]. If the RHNs have masses in the GeV to TeV range, they typically have a long lifetime and can be probed at the LHC [38, 43–45, 48, 49, 163–174]. Examples of well-motivated, UV-complete models with RHNs include the neutrino minimal SM (ν MSM) [175, 176] and the left–right symmetric model [177–180]. Characteristic features of models in this category are LLPs produced singly via SM neutral- and charged-current (CC) interactions, and lepton-rich signatures in terms of prompt and displaced objects (often in association with quarks). For example, in extended scenarios like left–right symmetric models, production through new right-handed W and Z bosons can result in between one and four LLPs, and cascade decays between RHNs can lead to phenomena such as doubly displaced decays. Additionally, RHNs can be produced via Higgs decays [40, 41, 44, 170, 172, 181].

It is possible for a given model to fit into two or more of the umbrella UV model categories. For example, a SUSY theory with a stable lightest SUSY particle (LSP) could have the LSP serve as a DM candidate, while alternatively DM could be a new electroweak multiplet, giving rise to SUSY-like signatures [64, 128–130]. In other models featuring particles charged under a confining gauge group (such as ‘quirks’ [182]), there can exist many production possibilities for the LLPs, including via the Higgs portal and the annihilation of new TeV-scale states (see, for example, [183]). Thus, the umbrella models should not be considered as exclusive categories, but rather as over-arching scenarios that motivate particular classes of signatures (such as new SM gauge-charged particles in the SUSY-like category, or presence of \cancel{E}_T in DM models).

In developing our simplified model framework below, we construct maps between these UV model categories and the simplified model channels to illuminate some of the best-motivated combinations of production and decay modes for LLPs. This allows us to focus on the most interesting channels and assess their coverage in chapter 3.




Figure 3. Schematic illustrations of LLP production modes in our simplified model framework. From top to bottom and left to right: direct pair production (DPP); HP; Higgs modes (HIG), including gluon fusion and VBF production (not shown here is VH production); heavy resonance (RES); charged current (CC).

2.3. The simplified model building blocks

As discussed above, production and decay can largely be factorized in LLP searches¹⁴⁴. This allows us to specify the relevant production and decay modes for LLP models separately; we then put them together and map the space of models into the umbrella categories of motivated theories.

2.3.1. Production modes. Motivated by our over-arching UV frameworks, we can identify a minimal set of interesting production modes for LLPs. Schematic diagrams for each production mode are shown in figure 3. These production modes determine LLP signal rates both by relating the LLP production cross section to meaningful theory parameters such as gauge charges or Higgs couplings, and by determining the kinematic distribution of the LLP. Additionally, a given production mechanism also makes clear predictions for the number and type of *prompt* objects accompanying the LLP(s). These prompt AOs can be important for both triggering on events with LLPs and for background rejection, particularly when the LLP has a low mass or decays purely hadronically, and they can be either SM states (leptons, \cancel{E}_T , tagging jets) or BSM objects such as Z' or dark photons [157, 184, 185].

¹⁴⁴ Once again, we comment that non-factorization of production and decay due to LLP interaction with the detector material and non-trivial propagation effects arise in models with LLPs with electric or color charge, and we discuss these subtleties further in section 2.5.2.

- *Direct-pair production (DPP):* Here the LLP is dominantly pair-produced non-resonantly from SM initial states. This is most straightforwardly obtained when the LLP is charged under a SM gauge interaction. In this case, an irreducible production cross section is then specified by the LLP gauge charge and mass. Such continuum DPP can also occur in the presence of a (heavy, virtual) mediator (e.g. an initial quark–antiquark pair may exchange a virtual squark to pair produce bino-like neutralinos); in this case the production cross section is essentially a free parameter, as it is determined by the unknown heavy mediator masses and couplings.
- *Heavy parent (HP):* The LLP is produced in the decays of on-shell HP particles that are themselves pair produced from the pp initial state. The production cross section is essentially a free parameter, and is indirectly specified by the gauge charges and masses of the HP particles. Heavy-parent production gives very different kinematics for the LLP than DPP, and often produces additional prompt AOs in the rapid cascade decays of the parents.
- *Higgs (HIG):* Here the LLP is produced through its couplings to the SM-like Higgs boson. This case has an interesting interplay of possible production modes. The dominant production is via gluon fusion, which features no AOs beyond gluon ISR. Owing to its role in electroweak symmetry breaking, however, the Higgs has associated production modes (VBF, VH), each with its own characteristic features. The best prospects for discovery are for LLP masses below $m_h/2$, in which case the LLPs can be in decays of the on-shell SM-like Higgs boson. Higher-mass LLPs can still be produced via an off-shell Higgs, albeit at substantially lower rates [26, 186]. The LLP can be pair produced or singly produced through the Higgs portal depending on the model; an LLP X can also be produced in association with \cancel{E}_T via $h \rightarrow XX + \cancel{E}_T$ or $h \rightarrow X + \cancel{E}_T$. The cross section (or, equivalently, the Higgs branching fraction into the LLP) is a free parameter of the model. The Higgs mass can also be taken as a free parameter: there exist many theories that predict new exotic scalar states (such as the singlet-scalar extension to the SM [119]), and these new scalars can be produced in the same manner as the SM Higgs.
- *Heavy resonance (RES):* Here the LLP is produced in the decay of an on-shell resonance, such as a heavy Z' gauge boson initiated by $q\bar{q}$ initial state. Note that production via an off-shell resonance is kinematically similar to the DPP mode. As with HIG, the LLP can be pair produced or singly produced (potentially in association with \cancel{E}_T). In RES models, ISR is the dominant source of prompt AOs. Models with new heavy scalars could conceivably fall into either RES or HIG; the main determining factor according to our organizational scheme is whether the scalar possesses Higgs-like production modes such as VBF and VH. Note that heavy resonance decays to SM particles also occur in these models, and searches for such resonances [187–193] may complement the sensitivity for decays to LLPs.
- *Charged current:* In models with weak-scale RHNs, the LLP can be produced in the leptonic decays of W/W' . Single production is favored. Prompt charged leptons from the CC interaction are typical prompt AOs.

It is important to note that each of the above production mechanisms has its own ‘natural’ set of triggers to record the signal. For example, HIG production can be accompanied by forward jets or leptons that are characteristic of VBF or VH production. Similarly, CC production often results in prompt charged leptons, while HP production comes with AOs from the HP decay. However, the reader should be cautioned that this does not necessarily mean that the ‘natural’ trigger is *optimal* for a particular signal. For example, the HIG modes suggest the use of VBF- or VH-based triggers, but if the LLP decays leptonically, it might be

more efficient to trigger on the lepton from the LLP decay. Thus, the final word on which trigger is most effective for a given simplified model depends on the production mode as well as the nature and kinematics of the LLP decay. The prompt AOs of each production mode could still, however, be used to extend sensitivity to the model (see section 6.3).

We also comment that some models may span several production modes. For example, a charged LLP that is part of an electroweak multiplet and nearly degenerate with a stable, neutral component [63–65, 99, 128–130, 194, 195] gives both DPP signatures (via $pp \rightarrow \chi^+ \chi^-$) and CC production (via associated production $pp \rightarrow \chi^\pm \chi^0$). Comprehensive coverage of each of the above production modes will allow for a conservative determination of sensitivity for models that span many production modes.

2.3.2. Decay modes. We now list a characteristic set of LLP decay modes. As we attempt to construct a minimal, manageable set of decay-mode building blocks, it is important to bear in mind that a given experimental search for LLPs can frequently be sensitive to a variety of possible LLP decay modes. As a result, it is not always necessary to perform separate searches for each possible decay mode as might otherwise be needed for prompt signatures.

The fact that LLP searches can be sensitive to many LLP decay modes is, in part, because LLPs that decay far from the collision point offer fewer avenues for particle identification. For example, for an LLP decaying inside of the calorimeter, most decay products are reconstructed as missing energy, or an energy deposition in the calorimeter. Consequently, particle identification criteria are typically relaxed in comparison to requirements on searches without displaced objects. Indeed, these ‘loose’ collider objects can differ significantly from the corresponding ‘tight’ prompt objects. This leads to more inclusive analyses that can cover a wider range of signatures with a single search.

Additionally, backgrounds for LLP searches are often small; for a comprehensive discussion of backgrounds to LLP searches, see chapter 4. As a result, tight identification and/or reconstruction criteria typically found in exclusive prompt analyses are no longer needed to suppress backgrounds. For example, ATLAS has a displaced vertex (DV) search sensitive to di-lepton and multi-track vertices that is relatively inclusive with respect to other objects originating from near the DV [196]. Similarly, CMS has an analysis sensitive to events with one each of a high-impact-parameter muon and electron without reconstructing a vertex or any other objects [197]. For these examples, the backgrounds are sufficiently low that other requirements may be relaxed and the specific decay mode of the LLP may not be too important so long as certain objects (such as muons) are present or the decay occurs in a specific location. An even more extreme example in this regard is the search for highly-ionizing tracks sensitive to electrically and color-charged LLPs. While the searches are primarily targeted to detector-stable particles (HSCPs or R -hadrons) they can also be used to probe intermediate lifetimes for which only a certain fraction of LLPs traverse the tracker before decaying (see e.g. [142]). Both because of low backgrounds as well as modified particle identification criteria compared to prompt searches, LLP searches can often be inclusive and therefore covered by a more limited range of simplified models.

In some cases, however, the topology of a decay does matter. One potentially important factor that influences the sensitivity of a search to a particular model is whether the LLP decays into two SM objects versus three, because the kinematics of multi-body decay are distinct from two-body decay and this may affect the acceptance of particular search strategies. An additional simplified model featuring a three-body decay of the LLP may consequently be needed to span the space of signatures.

Below, we describe an irreducible set of decay modes that can be used to characterize LLP signatures for various LLP charges (including neutral, electrically charged, and color

charged). For each, we also provide an explicit example for how the decay would appear in a particular UV model. We emphasize that the following decay modes are loosely defined with the understanding that their signatures are also representative of similar, related decay modes; for example $2j$ or $2j + \cancel{E}_T$ can also be proxies for $3j$ because searches for multi-body hadronic LLP decays can be sensitive to both and typically do not require reconstruction of a third jet. It should also be noted that we are not recommending searches to be optimized to the exact, exclusive decay mode because that could suppress sensitivity to related but slightly more complicated LLP decays.

- *Di-photon decays:* The LLP can decay resonantly to $\gamma\gamma$ (like in Higgs-portal models or left-right symmetric models [198]) or to $\gamma\gamma + \text{invisible}$ (in DM models). This latter mode stands as a proxy for other $\gamma\gamma + X$ decays where the third object is not explicitly reconstructed, although whether X is truly invisible can influence the triggers used. *Example: a singlino decaying to a singlet (which decays to $\gamma\gamma$) and a gravitino in Stealth SUSY* [54].
- *Single-photon decays:* The LLP decays to $\gamma + \text{invisible}$ (like in SUSY models). The SUSY model mandates a near-massless invisible particle, while other models (such as DM theories [134, 156]) allow for a heavy invisible particle. *Example: a bino decaying to photon plus gravitino in gauge-mediated models of SUSY breaking* [199].
- *Hadronic decays:* The LLP can decay into two jets (jj) (like in Higgs and gauge-portal models, or RPV SUSY), $jj + \text{invisible}$ (SUSY, DM, or neutrino models), or $j + \text{invisible}$ (SUSY). Here, ‘jet’ (j) means either a light-quark parton, gluon, or b -quark. This category also encompasses decays directly into hadrons (for example, LLP decay into π^+ plus an invisible particle [63–65]). *Example: a scalar LLP decaying to $b\bar{b}$ due to mixing with the SM Higgs boson, as in models of neutral naturalness* [13, 14, 118].
- *Semi-leptonic decays:* The LLP can decay into a lepton + 1 jet (such as in leptoquark (LQs) models) or 2 jets (like in SUSY or neutrino models). *Example: a RHN decaying to a left-handed lepton and an on- or off-shell hadronically decaying W boson (or W' boson in a left-right symmetric model)* [163].
- *Leptonic decays:* The LLP can decay into $\ell^+\ell^- (+\text{invisible})$, or $\ell^\pm + \text{invisible}$ (as in Higgs-portal, gauge-portal, SUSY, or neutrino models). Here the symbol ℓ may be any flavor of charged lepton, but the decays are lepton flavor-universal and (for $\ell^+\ell^-$ decays) flavor-conserving. *Example: a wino decaying to a neutralino and an on- or off-shell leptonic Z boson in SUSY* [55].
- *Flavored leptonic decays:* The LLP can decay into $\ell_\alpha + \text{invisible}$, $\ell_\alpha^+\ell_\beta^-$ or $\ell_\alpha^+\ell_\beta^- + \text{invisible}$ where flavors $\alpha \neq \beta$ (as in SUSY or neutrino models). *Example: a neutralino decaying to two leptons and a neutrino in R-parity-violating SUSY* [55]; or a RHN decaying to two leptons and a neutrino [200].

In all cases, both the LLP mass and proper lifetime are free parameters. Therefore, the case of detector-stable particles is automatically included by taking any of the above decay modes and taking the lifetime to infinity¹⁴⁵. We emphasize that, depending on the location of the LLP within the detector, these decay modes may or may not be individually distinguishable: a displaced di-jet decay will look very different from a displaced di-photon decay in the tracker, but nearly identical if the decay occurs in the calorimeter. The goal here is to identify promising channels (as distinct from detector signatures).

¹⁴⁵ As mentioned earlier, in the $c\tau \rightarrow \infty$ limit the decay mode becomes irrelevant. However, an exception is the search for particles that are stopped inside the detector material and decay out of time, which are discussed in section 3.5.3.

As an example of how the above-listed decay modes cover the most important experimental signatures, we consider a scenario of an LLP decaying to top quarks. This scenario is very well motivated (for instance, with long-lived stops in SUSY) and might appear to merit its own decay category of an LLP decaying to one or more top quarks. However, the top quark immediately decays to final states that *are* covered in the above list, giving an effective semileptonic decay mode ($t \rightarrow b\ell^+\nu$) and a hadronic decay mode ($t \rightarrow bjj$) of the LLP. Similarly, LLP decays to four or more final states are typically covered by the above inclusive definitions of decay modes; this provides motivation not to over-optimize experimental searches to the specific, exclusive features of a particular decay mode.

While it would be ideal to have separate experimental searches for each of the above decay modes (when distinguishable), it is rare for specific models to allow the LLP to decay in only one manner; as in the example of an LLP decaying to a top quark, a number of decay modes typically occur with specific predictions for the branching fractions. As another example, if the LLP couples to the SM via mixing with the SM Higgs boson, then the LLP decays via mass-proportional couplings giving rise to b - and τ -rich signatures. If, instead, the LLP decays through a kinetic mixing as in the case of dark photons or Z bosons, then the LLP can decay to any particle charged under the weak interactions, giving rise to a relatively large leptonic branching fraction in addition to hadronic decay modes. This allows some level of prioritization of decay modes based on motivated UV-complete models; for example, the Higgs-portal model prioritizes searches for heavy-flavor quarks and leptons in LLP decay, while the gauge-portal model prioritizes searches for electrons and muons in LLP decay. Ultimately, however, it is desirable to retain independent sensitivity to each individual decay mode as much as possible. Indeed, for each decay mode listed above, models exist for which the given decay mode would be the main discovery channel.

Invisible final-state particles: Where invisible particles appear as products of LLP decays, additional model dependence arises from the unknown mass of the invisible particle. The invisible particle could be a SM neutrino, DM, an LSP in SUSY, or another BSM particle. The phenomenology depends strongly on the mass splitting, $\Delta \equiv M_{\text{LLP}} - M_{\text{invisible}}$. If $\Delta \ll M_{\text{LLP}}$ (i.e. $M_{\text{LLP}} \sim M_{\text{invisible}}$), the spectrum is compressed and the visible decay products of the LLP are soft. This could, for instance, lead to signatures such as disappearing tracks (DTs) or necessitate the use of ISR jets to trigger on the LLP signature. If the mass splitting is large, $M_{\text{invisible}} \ll M_{\text{LLP}}$, then the signatures lose their dependence on the invisible particle mass.

We suggest three possible benchmarks: a compressed spectrum with $\Delta \ll M_{\text{LLP}}$ (example: a nearly degenerate chargino-neutralino pair, giving rise to soft leptons or DTs [63–65, 99, 128–130, 194, 195]); a massless invisible state, $\Delta = M_{\text{LLP}}$ (example: a next-to-lightest SUSY particle (NLSP) decaying to SM particles and a massless gravitino in gauge-mediated SUSY breaking (GMSB) [112, 201–206]); and an intermediate splitting corresponding to a democratic mass hierarchy, $\Delta \approx M_{\text{LLP}}/2$ (example: NLSPs in mini-split SUSY [57, 58, 96]).

2.4. A simplified model proposal

In this section, we present a compact set of simplified model channels that, broadly speaking, covers the space of theoretical models in order to motivate new experimental searches. Such a minimal, compact set may not be optimal for reinterpretation of results (where variations on our listed production and decay modes may influence signal efficiencies and cross section sensitivities), but rather provides a convenient characterization of possible signals to ensure that no major discovery mode is missed. These models may therefore serve as a starting point

for systematically understanding experimental coverage of LLP signatures and devising new searches, but may need to be extended in future for the purposes of facilitating reinterpretation. We undertake an in-depth discussion of these topics in section 6.

We classify LLPs according to their SM gauge charges, as these dictate the dominant or allowed LLP production and decay modes, and can give rise to different signatures (for example, DTs and hadronized LLPs). We separately consider LLPs that are: (a) neutral; (b) electrically charged but color neutral; and (c) color charged. In the latter case, it is important to distinguish between the long-lived parton (which carries a charge under quantum chromodynamics, QCD) that hadronizes prior to decay, and the physical LLP, which is a color-singlet ‘*R*-hadron’ (using the standard nomenclature inspired by SUSY). The decays of the *R*-hadron are still dominated by the parton-level processes.

All of the following models have the LLP mass and lifetime as free parameters. For HP production, the parent mass is an additional parameter, while for invisible decays, several different benchmarks for mass splittings between LLP and invisible final state may have to be separately considered as described in section 2.3.2. The cross section may have a theoretically well-motivated target value depending on UV-model parameters, but phenomenologically can generally be taken as a free parameter.

We emphasize that in spite of the many simplified model channels proposed below, a small number of experimental LLP searches can have excellent coverage over a wide range of channels (at least for certain lifetime ranges). The list is intended to be comprehensive in order to identify whether there are new searches that could have a similarly high impact on the space of simplified models, and identify where the gaps in coverage are.

2.4.1. Neutral LLPs. The simplified model channels for neutral LLPs are shown in table 1, where *X* indicates the LLP.

In our initial proposal, which is the first iteration of the simplified model framework, it is sufficient to consider as ‘jets’ all of the following: $j = u, d, s, c, b, g$. It is worth commenting that *b*-quarks pose unique challenges and opportunities. Since *b*-quarks are themselves LLPs, they appear with an additional displacement relative to the LLP decay location. They also often give rise to soft muons in their decays, which could in principle lead to additional trigger or selection possibilities. However, these subtleties can be addressed in further refinements of the simplified models; we discuss this further in section 2.6. Similarly, we consider e , μ , and τ to be included in the broad category of ‘leptons’, with the proviso that searches should be designed where possible with sensitivity to each.

When multiple production modes are specified in one row of the table, this means that multiple especially well-motivated production channels give rise to similar signatures. Typically only one of these simplified model production modes will actually need to be included when developing and assessing sensitivity of an experimental search, but we sometimes include multiple different production modes as individuals may variously prefer one over the other.

In each entry of the table, we indicate which umbrella category of well-motivated UV models (section 2.2) can predict a particular (production) \times (decay) mode. An asterisk (*) on the umbrella model indicates that \cancel{E}_T is required in the decay. A dagger (†) indicates that this particle production \times decay scenario is not present in the *simplest and most minimal* implementations or spectra of the umbrella model, but could be present in extensions of the minimal models. While the HIG production signatures are best-motivated for the SM-like 125 GeV Higgs, exotic Higgses of other masses can still have the same production modes and so m_H can be taken as a free parameter.

Table 1. Simplified model channels for neutral LLPs. The LLP is indicated by X . Each row shows a separate production mode and each column shows a separate possible decay mode, and therefore every cell in the table corresponds to a different simplified model channel of (production) \times (decay). We have cross-referenced the UV models from section 2.2 with cells in the table to show how the most common signatures of complete models populate the simplified model space. The asterisk (*) shows that the model definitively predicts missing energy in the LLP decay. A dagger (\dagger) indicates that this particle production \times decay scenario is not present in the *simplest and most minimal* implementations or spectra of the umbrella model, but could be present in extensions of the minimal models. When two production modes are provided (with an ‘or’), either simplified model can be used to simulate the same simplified model channel.

Production	Decay					
	$\gamma\gamma(+\text{inv.})$	$\gamma + \text{inv.}$	$jj(+\text{inv.})$	$jj\ell$	$\ell^+\ell^-(+\text{inv.})$	$\ell_\alpha^+\ell_{\beta\neq\alpha}^-(+\text{inv.})$
25 DPP: sneutrino pair or neutralino pair HP: squark pair, $\tilde{q} \rightarrow jjX$ or gluino pair $\tilde{g} \rightarrow jjX$ HP: slepton pair, $\tilde{\ell} \rightarrow \ell X$ or chargino pair, $\tilde{\chi} \rightarrow WX$ HIG: $h \rightarrow XX$ or $\rightarrow XX + \text{inv.}$ HIG: $h \rightarrow X + \text{inv.}$ RES: $Z(Z') \rightarrow XX$ or $\rightarrow XX + \text{inv.}$ RES: $Z(Z') \rightarrow X + \text{inv.}$ CC: $W(W') \rightarrow \ell X$	\dagger	SUSY	SUSY	SUSY	SUSY	SUSY
	\dagger	SUSY	SUSY	SUSY	SUSY	SUSY
	\dagger	SUSY	SUSY	SUSY	SUSY	SUSY
	Higgs, DM*	\dagger	Higgs, DM*	RH ν	Higgs, DM* RH ν^*	RH ν^*
	DM*, RH ν	\dagger	DM*	RH ν	DM*	\dagger
	Z', DM*	\dagger	Z', DM*	RH ν	Z', DM*	\dagger
	DM	\dagger	DM	RH ν	DM	\dagger
	\dagger	\dagger	RH ν^*	RH ν	RH ν^*	RH ν^*

We remind the reader that the production modes listed in table 1 encompass also the associated production of characteristic prompt objects. For example, the Higgs production modes not only proceed through gluon fusion, but also through VBF and VH production, each of which results in associated prompt objects such as forward jets in VBF, and leptons or \not{E}_T in VH. All of the production modes listed in table 1 could be accompanied by ISR jets that aid in triggering or identifying signal events. It is therefore important that searches are designed to exploit such prompt AOs whenever they can improve signal sensitivity, especially with regard to triggering.

To demonstrate how to map full models onto the list of simplified models (and vice-versa), we consider a few concrete cases. For instance, if we consider a model of neutral naturalness where X is a long-lived scalar that decays via Higgs mixing (for instance, X could be the lightest quasi-stable glueball), then the process where the SM Higgs h decays via $h \rightarrow XX, X \rightarrow b\bar{b}$ would be covered with the HIG production mechanism and a di-jet decay. Entirely unrelated models, such as the case where X is a bino-like neutralino with RPV decays $h \rightarrow XX, X \rightarrow jjj$ could be covered with the same simplified model because most hadronic LLP searches do not have exclusive requirements on jet multiplicity. Similarly, a hidden-sector model with a dark photon, A' , produced in $h \rightarrow A'A', A' \rightarrow f\bar{f}$ would also give rise to the di-jet signature when f is a quark, whereas it would populate the $\ell^+\ell^-$ column if f is a lepton. Finally, a scenario with multiple hidden-sector states X_1 and X_2 , in which X_2 is an LLP and X_1 is a stable, invisible particle, could give rise to signatures like $h \rightarrow X_2 X_2, X_2 \rightarrow X_1 jj$ that would be covered by the same HIG production, hadronic-decay simplified model; however, we see how \not{E}_T can easily appear in the final state, and that the LLP decay products may not be entirely hadronic. Therefore, the simplified models in table 1 can cover an incredibly broad range of signatures, but only if searches are not overly optimized to particular features such as \not{E}_T and LLPs decaying entirely visibly (which would allow reconstruction of the LLP mass)¹⁴⁶.

2.4.2. Electrically charged LLPs: $|Q| = 1$. For an electrically charged LLP, we need to consider far fewer production modes because of the irreducible gauge production associated with the electric charge. We still consider the additional possibility of a HP scenario where the parent has a QCD charge, as this could potentially dominate the production cross section, see e.g. [95]. We summarize our proposals in table 2.

Note that we group all resonant production into the Z' simplified model. The reason is that the SM Higgs cannot decay into two on-shell charged particles due to the model-independent limits from LEP on charged particle masses, $M \gtrsim 75\text{--}90\text{ GeV}$ (see, for example, [207]); because of this lower limit on the LLP mass, it is less important to use AOs for triggering and reconstructing charged LLP signatures than for neutral LLPs. Additionally, there are fewer allowed decay modes because of the requirement of charge conservation.

For concreteness, we recommend using $|Q| = 1$ as a benchmark for charged LLPs for the purpose of determining allowed decay modes. Although other values of Q are possible, these often result in cosmologically stable charged relics or necessitate different decay modes than those listed here. Additionally, LLPs with $|Q| = 1$ are motivated within SUSY [63–65, 208–210] and within Type-III seesaw models of neutrino masses [211–214]. We note that there exist already dedicated searches for heavy quasi-stable charged particles with non-

¹⁴⁶ This should not, of course, be interpreted as saying that searches shouldn't be done that exploit these features. Instead, our position is that experiments should bear in mind the range of topologies and models covered by each cell in table 1 when designing searches, and that some more inclusive signal regions should be established where possible.

Table 2. Simplified model channels for electrically charged LLPs such that $|Q| = 1$. The LLP is indicated by X . Each row shows a separate production mode and each column shows a separate possible decay mode, and therefore every cell in the table corresponds to a different simplified model channel of (production) \times (decay). We have cross-referenced the UV models from section 2.2 with cells in the table to show how the most common signatures of complete models populate the simplified model space. The asterisk (*) shows that the model definitively predicts missing energy in the LLP decay. A dagger (\dagger) indicates that this particle production \times decay scenario is not present in the *simplest and most minimal* implementations or spectra of the umbrella model, but could be present in extensions of the minimal models. When two production modes are provided (with an ‘or’), both production simplified models can be used to cover the same experimental signatures.

Production	Decay			
	$\ell + \text{inv.}$	$jj(+\text{inv.})$	$jj\ell$	$\ell\gamma$
DPP: chargino pair or slepton pair	SUSY DM*	SUSY DM*	SUSY	\dagger
HP: $\tilde{q} \rightarrow jX$	SUSY DM*	SUSY DM*	SUSY	\dagger
RES: $Z' \rightarrow XX$	Z', DM^*	Z', DM^*	Z'	\dagger
CC: $W' \rightarrow X + \text{inv.}$	DM*	DM*	$\text{RH}\nu$	\dagger

standard charges [215, 216]. Because such searches are by construction not intended to be sensitive to the decays of the LLP, the existing models are sufficient for characterizing these signatures and they do not need to be additionally included in our framework.

For massive particles with $|Q| = 1$ with intermediate or large lifetimes such that the LLP traverses a significant part (or all) of the tracker, the highly ionizing track of the LLP provides a prominent signature. This can be exploited for an efficient suppression of backgrounds while keeping identification and/or reconstruction criteria as loose and, hence, as inclusive as possible. In particular, for decay-lengths of the order of or larger than the detector size, the signature of highly ionizing tracks and anomalous time of flight (TOF) (i.e. searches for HSCPs; see sections 3.5 and 6.4.1) constitute an important search strategy covering a large range of lifetimes present in the parameter space of theoretically motivated models. While the searches for HSCPs are largely inclusive with respect to additional objects in the event, they depend strongly on the velocity of the LLP. For $\beta \rightarrow 1$ one loses the discriminating power against minimally ionizing particles, while for small velocities, $\beta \lesssim 0.5$, the reconstruction becomes increasingly difficult due to timing issues. It is therefore important to include the HP production scenario which covers a much larger kinematic range than direct production alone and which may feature a much wider range of signal efficiencies than the DPP scenario [97].

While the signatures in table 2 form a minimal set, they also encompass some scenarios that merit special comment. One of these is the DT signature [63–65, 99, 128–130, 194, 195], in which a charged LLP decays to a nearly degenerate neutral particle. The lifetime is long in this scenario due to the tiny mass splitting between the two states. Formally, these are included in the chargino or slepton DPP modes in table 2 with decays to $\ell + \text{inv.}$ or $q\bar{q}' + \text{inv.}$ taken in the limit where the splitting between the charged LLP and the invisible final state is of $\mathcal{O}(200 \text{ MeV})$. In the case of a hadronic decay, X decays to a soft pion that is very challenging to reconstruct and so the track simply disappears. This is an important scenario that is already the topic of existing searches [217, 218]. As the degeneracy between

Table 3. Simplified model channels for LLPs with color charge. The LLP is indicated by X . Each row shows a separate production mode and each column shows a separate possible decay mode, and therefore every cell in the table corresponds to a different simplified model channel of (production) \times (decay). We have cross-referenced the UV models from section 2.2 with cells in the table to show how the most common signatures of complete models populate the simplified model space. A dagger (†) indicates that this particle production \times decay scenario is not present in the *simplest and most minimal* implementations or spectra of the umbrella model, but could be present in extensions of the minimal models. When two production modes are provided (with an ‘or’), both production simplified models can be used to cover the same experimental signatures.

		Decay			
Production		$j + \text{inv.}$	$jj (+\text{inv.})$	$j\ell$	$j\gamma$
DPP: squark pair or gluino pair		SUSY	SUSY	SUSY	†

the charged LLP and the neutral state is relaxed, other signatures are possible; this parameter range is well motivated both by SUSY and DM models with coannihilation [98, 136, 137].

Finally, we comment on the challenges of simulating the charged LLP simplified models. Because the LLP bends and interacts with detector material prior to its decay, the simulation of the LLP propagation is important in correctly modeling the experimental signature. The subsequent decay of the LLP must either be hard-coded into the detector simulation, or allow for an interface with programs such as Pythia 8 to implement the decays. We discuss the challenges of simulating signals for LLPs with electric or color charge in section 2.5.2.

2.4.3. LLPs with color charge. An LLP charged under QCD is more constrained than even electrically charged LLPs. Because of the non-Abelian nature of the strong interactions, the gauge pair-production cross section of the LLP is specified by the LLP mass and its representation under the color group, $SU(3)_C$. We do not consider LLP production via a HP particle because that cross section is unlikely to dominate the total production rate at the LHC relative to DPP. The simplified model channels are provided in table 3.

A complication of the QCD-charged LLP is that the LLP hadronizes prior to its decay, forming an R -hadron bound state. The modeling of hadronization and subsequent propagation is directly related to many properties of the long-lived parton, such as electric charge, flavor, and spin. Event generators such as Pythia 8 have routines [219, 220] to simulate LLP hadronization, although it is unclear how precise these predictions are. For a point of comparison, using the default settings of Pythia 8 yields an estimate of the neutral R -hadron fraction from a gluino (color-octet fermion, \tilde{g}) of approximately 54%, while the neutral R -hadron fraction for a stop (scalar top partner) is estimated to be 44% [96]. After hadronization, the charge of the R -hadron may change as it passes through the detector. For instance, some estimates [221, 222] suggest that heavy, color-octet gluinos \tilde{g} would predominantly form mesons (e.g. $(u\tilde{g}\bar{d})$) at first. They eventually drop to the lower-energy neutral singlet baryon $\tilde{\Lambda} = (\tilde{g}uds)$ state when interacting with the protons and neutrons within the calorimeters.

The modeling of LLP hadronization and propagation is crucial to designing searches for color-charged LLPs and assessing their sensitivity. For example, only the charged R -hadrons can be found in HSCP search; if the LLP charge changes as it passes through the detector, HSCP searches may have limited sensitivity. To take this into account, the experimental searches include both tracker-only or tracker + calorimeter signal regions [4, 223], which

enhances sensitivity to the scenario in which R -hadrons lose their charge by the time they reach the calorimeters.

Because no R -hadrons have been discovered to date and hence their properties cannot be directly measured, R -hadron modeling in detector simulations is challenging. We discuss the challenges of simulating the propagation and decays of LLPs with color charge in section 2.5.2.

2.5. *Proposal for a simplified model library*

The simplified models outlined in the above sections provide a common language for theorists and experimentalists to study the sensitivity of existing searches, propose new search ideas, and interpret results in terms of UV models. Each of these activities demands a simple framework for the simulation of signal events that can be used to evaluate signal efficiencies of different search strategies and map these back onto model parameters. Requiring individual users to create their own MC models for each simplified model is impractical, redundant, and invites the introduction of errors into the analysis process.

In this section, we propose and provide a draft version of a *simplified model library* consisting of model files and MC generator cards that can be used to generate events for various simplified models in a straightforward fashion. Because each experiment uses slightly different MC generators and settings, this allows each collaboration (as well as theorists) to generate events for each simplified model based on the provided files. Depending on how the LLP program expands and develops over the next few years, it may become expedient to expand the simplified model library to include sets of events in a standard format (such as the Les Houches format [224]) that can be directly fed into event-generator and detector-simulation programs. Given the factorization of production and decay of LLPs that is valid for neutral LLPs, this could involve two mini-libraries: a set of production events for LLPs and a set of decay events for LLPs, along with a protocol for ‘stitching’ the events together.

The current version of the library is available at the LHC LLP Community website¹⁴⁷, hosted at CERN. In appendix, we also provide tables that list how to simulate each LLP simplified model channel with one of the specified base models. These proposals are based on the models outlined in section 2.5.1 and often match the best-motivated simplified models from section 2.4, and also building on the DM-inspired LLP simplified models proposed and detailed in [100]. The library currently focuses on models of neutral LLPs; simulating the propagation of charged LLPs along with the full range of decays listed in sections 2.4.2–2.4.3 requires more careful collaboration with detector simulation and other MC programs to ensure that they can practically be used in experimental studies.

We provide model files in the popular universal FeynRules output (UFO) format [225], which is designed to interface easily with parton-level simulation programs such as MadGraph5_aMC@NLO [226]. The goal is to cover as many of the simplified models of section 2.4 with as few UFO models as possible; this limits the amount of upkeep needed to maintain the library and develops familiarity with the few UFO models needed to simulate the LLP simplified models. We provide specific instructions for how to simulate each simplified LLP channel along with the UFO models.

2.5.1. *Base models for library.* In order to reproduce the simplified model channels of section 2.3, we need a collection of models that:

¹⁴⁷ <http://cern.ch/longlivedparticles>

- Includes additional gauge bosons and scalars to allow vector- and scalar-portal production of LLPs (RES and HIG).
- Includes new gauge-charged fermions and scalars to cover direct and simple cascade production modes of LLPs (DPP and HP).
- Includes a RHN-like state with couplings to SM neutrinos and leptons (CC).
- *Recommended, but optional:* Allows for the decays of the LLP particle through all of the decay modes listed in section 2.3, either through renormalizable or higher-dimensional couplings. If couplings that allow LLP decay are included in the UFO model, then the decays can be performed directly at the matrix-element level in programs such as MadGraph5_aMC@NLO [226] and accompanying packages such as MadSpin [227]. Alternatively, it is possible for neutral LLPs to simulate the production and decay as a single process; in such cases, numerical instabilities sometimes arise, for which dedicated event generators are needed [49]. If the couplings needed for LLP are not in the UFO model, then LLPs can be left stable at the matrix-element level and decays implemented via Pythia 8 [219, 220], which allows for the straightforward implementation of decays according to a phase-space model, but does not correctly model the angular distribution of decay products. Instructions for implementing decays in Pythia are included with the model library files.

Fortunately, an extensive set of UFO models is already available for simulating the production of BSM particles. We note that extensions or generalizations of only three already-available UFO models are needed at the present time; the SUSY models in particular can cover many of the simplified models since they contain an enormous collection of new fermions and scalars. We also provide an optional fourth model, the Hidden Abelian Higgs Model, that can be helpful to simulate HIG and ZP theories.

1. *The MSSM:* The use of this model is motivated by and allows for the simulation of SUSY-like theories. The model contains a whole host of new particles with various gauge charges and spins. Therefore, an MSSM-based model allows for the simulation of many of the simplified model channels. In particular, we note that existing UFO variants of the MSSM that include gauge-mediated supersymmetry breaking (GMSB) couplings (including decays to light gravitinos), R -parity violation (making unstable the otherwise stable LSPs [55, 110, 111]), and the phenomenological MSSM (pMSSM) [228, 229] already cover most of the SUSY-motivated LLP scenarios. In some cases, the model is modified to give direct couplings between the Higgs states and gluons/photons.
2. *The left-right symmetric model (LRSM):* This UFO model is best for simulating UV theories with right-handed neutrinos ($R\nu$). The UFO model supplements the SM by an additional $SU(2)_R$ symmetry, which gives additional charged and neutral gauge bosons. The model is available in the simplified models library and contains a RHN which is the typical LLP candidate. The LLP can be produced via SM W , Z , or via the new gauge and Higgs bosons (both charged and neutral) present in the theory¹⁴⁸. The LRSM therefore contains many of the charged and neutral current LLP production processes outlined in section 2.4.1.
3. *Dark-matter simplified models (DMSM):* These UFO models are best for simulating UV theories in the DM class. These UFO models have been created by the LHC DM working group [86]. They typically consist of a new BSM mediator particle (such as a scalar or a Z') coupled to invisible DM particles. The UFO models can either be modified to include an unstable LLP, or else the otherwise stable ‘DM’ particle can be decayed via Pythia.

¹⁴⁸ Additional LRSM tools are available at <https://sites.google.com/site/leftrighthep/>.

The utility and applicability of the DM simplified model framework to LLPs has already been demonstrated with a detailed proposal and study of classes of DM simplified models for LLPs [100]. These models are particularly good for simulating LLP production via a heavy resonance (RES), and can also simulate continuum production of LLPs in the limit where the mediator is taken to be light and off-shell (DPP).

4. *(Optional) The hidden abelian Higgs model (HAHM):* This UFO model contains new scalars and gauge bosons and so can be used to simulate both Higgs-portal and gauge-portal (ZP) theories. The model consists of the SM supplemented by a ‘hidden sector’ consisting of a new U(1) gauge boson and a corresponding Higgs field. The physical gauge and Higgs bosons couple to the SM via kinetic and mass mixing, respectively. The HAHM allows for straightforward simulation of Higgs-portal production of LLPs, as well as Z' models and many hidden sector scenarios. The UFO implementation is from [230].

If additional decay modes are needed beyond those in the specified simplified models, then the library can be updated to include the new couplings mediating the decay. Alternatively, the LLPs can be left stable at parton level and decayed in event generators such as Pythia.

A detailed list of processes that can be used to simulate each simplified model channel is provided in the [appendix](#). The primary purpose of the library is to be used to simulate events for determining acceptances, and, as a result, the signal cross section is not important. Thus, for example, SM gauge interactions can be used as proxies for much weaker exotic interactions. Similarly, the spins of the particles are generally of subdominant importance: replacing the direct production of a fermion with the direct production of a scalar will not fundamentally alter the signature. As long as results are expressed in terms of sensitivity to cross sections and not couplings, the results can be qualitatively (and in many cases, quantitatively) applied to any similar production mode regardless of spin. However, we caution the reader that changing the spin of the LLP (or its parent) can change the angular distribution, and since in some cases LLP searches are typically more sensitive to aspects of event geometry than prompt searches, the second-order effects of spin could have more of an effect than for prompt simplified models.

2.5.2. LLP propagation and interaction with detector material. Long-lived particles with electric or QCD charges interact with the detector material prior to decay, and their propagation through the detector must be correctly modeled. The propagation of both LLPs with color charge (in the form of R -hadrons) and electrically charged LLPs can be implemented in the Geant4 (G4) toolkit [231]. For example, routines exist to simulate the propagation of color-charged LLPs [232, 233]. G4 also includes routines that can implement N -body decays of LLPs using a phase-space model. This works fine for decays of LLPs to leptons, photons, invisible particles such as neutrinos, as well as exclusive hadronic decays.

However, G4 cannot implement decays to partons that subsequently shower and hadronize. One solution to this limitation is employed by CMS [234, 235] and ATLAS [236] in their searches for stopped LLPs. In these analyses, the signal simulation proceeds in two stages. During the first stage, the production of the LLP and its subsequent interactions with the detector are simulated. Once the stopping point of the LLP is determined, a new event is simulated including the LLP decay; the LLP decay products are then manually moved to the stopping point from the first stage. G4 is then run a second time to determine the efficiency for reconstructing the LLP decay signal.

It would be preferable to fully automate the simulation of decays of charged LLPs after propagation in G4. There exists in G4 a class called G4ExtDecayer, which can be used to

implement decays by interfacing with an external generator. This class has been used to interface G4 with Pythia 6¹⁴⁹. The interface with Pythia 6 has been used most recently to model LLP gluino propagation and decay in a search for DV and missing energy in ATLAS [237]. Work is ongoing to extend this functionality to Pythia 8 and to simplify the interface.

An additional challenge of simulating LLP decays is that, if the LLP undergoes a multi-body decay, generators such as Pythia use a phase-space model to implement the decays. If more accuracy is required, it may be preferable to use the full matrix element via generators such as MadGraph5 [226, 238]. If the matrix element is important for computing the decay of the LLP, then either an interface with MadGraph is needed to implement the decay prior to passing the vertex back to Pythia 8 for showering and hadronization, or matrix-element-based methods within the event generator itself must be used.

Because of the need to interface with G4 in simulating the decays of LLPs with electric or color charges, we do not at this point include such decay modes in our simplified model library. The decays of such LLPs will be most easily simulated via an interface with Pythia 8 once it is finalized.

Finally, we comment that LLPs can have even stranger propagation properties than LLPs with electric or color charges. For example, quirks are LLPs that are charged under a hidden-sector gauge interaction that confines at macroscopic scales [182]. Because the confinement scale can be just about any distance, quirks can have very unusual properties; as a specific example, if electrically charged quirk–antiquark pairs are bound on the millimeter or centimeter level, they behave as an electric dipole and therefore do not leave conventional tracks that bend in the magnetic field. Other confinement scales give rise to different behaviors, such as meta-stable heavy charged particles and non-helical tracks [239, 240]. In scenarios where the quirks carry color charge, the quirks hadronize and can undergo charge-flipping interactions as they move through the detector. These quirk scenarios can be challenging to model, and no public code exists that allows for the propagation and interaction of quirks with the detector material; we encourage the collaborations to validate and release any internal software they may have to study the propagation of quirks (for more discussion, see the discussion of quirks in section 3.5)¹⁵⁰.

2.6. Limitations of simplified models and future opportunities

We conclude our discussion of simplified models with a more extensive discussion of the limitations of the current simplified model proposal in its application to models of various types, along with opportunities for future development. The presented framework is only the first step of a simplified model program that is comprehensive in terms of generating LHC signatures and allowing straightforward reinterpretation of experimental results for UV models. The framework we have developed with separate, modular components for LLP production and decay is amenable to expansion, and we encourage members of the theory and experimental communities to continue to do so over the coming years to ensure maximal utility of the simplified models framework.

One significant simplification we have undertaken in our framework is to define a ‘jet’ as any of $j = u, c, d, s, b, g$. In reality, different partons give rise to different signatures, especially when one of the ‘jets’ is a heavy-flavor quark. Jets initiated by b and c quarks have some useful distinguishing features, such as the fact that the underlying heavy-flavor meson

¹⁴⁹ See http://geant4-userdoc.web.cern.ch/geant4-userdoc/Doxygen/examples_doc/html/Exampledecayer6.html

¹⁵⁰ Ideally, this software would be well-documented to facilitate sharing between experiments. A successful example of readily shareable software between experiments is the G4 package for R -hadrons and other particles’ interaction with matter, found at <http://r-hadrons.web.cern.ch/r-hadrons/>.

decays at a distance slightly displaced from the proton interaction vertex and that there are often associated soft leptons resulting from meson decays. In particular, it is possible that the soft muons associated with B -meson decays could be used to enhance trigger and reconstruction prospects for LLPs decaying to b -jets [241]. However, heavy quarks also constitute an important backgrounds for LLP searches, and so LLPs decaying to b - and c -jets may necessitate dedicated treatment in future. Similarly, LLP decays to τ leptons may merit further specialized studies.

Another property of the current framework is that it is restricted to LLP signatures of low multiplicity. By ‘low multiplicity’, we mean collider signatures with one or two LLPs. Searches inspired by these models are also suitable for many scenarios with three or four LLPs per event (which include models with dark-Higgs decays into lepton-jets (LJs) [155], or left-right symmetric models [169]), since the LLP signatures are generally extremely rare and so only one or two typically need to be identified in a given event to greatly suppress backgrounds. Thus, as long as the search is inclusive with respect to possible additional displaced objects, the signature can be covered with low-multiplicity strategies. As the LLP multiplicity grows, however, the simplified model space we have presented requires modification. This is both because the individual LLPs grow softer, making them harder to reconstruct on an individual level, and they become less separated in the detector, which makes isolation and identification of signal a challenge. On the other hand, the high LLP multiplicity may allow for new handles for further rejecting backgrounds, and the kinematics can vary widely based on the model (for example, in some ‘quirky’ scenarios, LLPs can be produced in a variety of ways with different kinematic distributions [120]). In extreme cases, signals can even mimic PU [242]. High-multiplicity signatures therefore require dedicated modeling, and we defer the study of these signatures to chapter 7.

Finally, we conclude by noting that simplified models are intended to provide a general framework to cover a broad swath of models. Any simplified model set-up, however, cannot cover every single UV model without becoming as complex as the UV model space itself. As with the case of promptly decaying new particles, care must also be taken in the interpretation of simplified models [78, 86–88, 106]: for example, constraints on simplified models assuming 100% branching fractions of LLPs to a particular final state may not accurately represent the constraint on a full model due to the large multiplicity of possible decay modes. There will additionally always be very well-motivated models that predict specific signatures that are challenging to incorporate into the simplified model framework outlined here. Experimental searches for these signatures should still be done where possible, but we encourage theorists and experimentalists alike to think carefully about how to design such searches so as to retain maximal sensitivity to simplified models that may give rise to similar signatures.

3. Experimental coverage of LLP signatures

Chapter editors: Juliette Alimena, Xabier Cid Vidal, Albert de Roeck, Jared Evans, Heather Russell, Jose Zurita

Contributors: David Curtin, Alberto Escalante del Valle, Philippe Mermod, Antonio Policicchio, Brian Shuve

A critical component of any discussion of LLP searches at the LHC is the comprehensive review of the existing searches from ATLAS, CMS, and LHCb, and an assessment of their coverage and any gaps therein. This is an inherently challenging task, given the varied and atypical objects often defined and utilized in LLP analyses and the differences among the

experiments. As such, the following discussion assumes little-to-no background on LLP search strategies and includes a high level of detail regarding the current analyses. The focus of the discussion is on the existing studies, while acknowledging that the landscape for new physics models and LLP signatures can be broader than the ones described here.

Backgrounds to most of these studies are typically small, as most LLP signatures are not naturally mimicked by any irreducible SM processes. Backgrounds for LLP searches typically include peripheral or machine effects, those rarely important for searches for prompt physics, including cosmic muons, beam halo, detector noise, and cavern backgrounds. Such backgrounds are discussed in detail in chapter 4. As rare as these backgrounds typically are, their rates are not completely negligible, and particular, model-dependent selection requirements (based on, for instance, the LLP mass range or specific decay modes) must be made to reduce backgrounds as much as possible and, in some cases, make the searches ‘background-free’. Additionally, many default object reconstruction algorithms are not designed to detect particles originating from decays of LLPs, and so dedicated reconstruction of tracks, jets, leptons, or other objects may be required for LLP searches. Taken altogether, these factors make LLP searches very different from searches for prompt objects, and the following discussion additionally aims to collate and summarize the current techniques for LLP reconstruction at the LHC.

A particular challenge for many LLP signatures is the trigger. With the exception of certain dedicated ATLAS triggers in the calorimeters or MS, there are no Level-1 (L1) triggers that directly exploit the displaced nature of LLP decays, and L1 trigger thresholds must be surpassed by standard objects (such as leptons or high-energy jets) for the event to be recorded¹⁵¹. Throughout this chapter, we highlight the role and limitations of the trigger(s) employed in current searches, and the design of customized LLP triggers is to be encouraged to probe new and otherwise inaccessible regions of parameter space.

A detailed review of all existing searches is presented in sections 3.1 through section 3.5. This survey of the current experimental coverage aims to highlight the highest-priority searches still yet to be performed, which we summarize in section 3.6. In all cases, we focus on the latest version of each analysis. Notably we will typically present searches based on data taken at a center-of-mass energy $\sqrt{s} = 13$ TeV, and discuss searches using Run 1 data only when the newer version is not yet available, or when there are conceptual differences between two versions of the same analysis.

Because LLPs travel macroscopic distances in the detectors, many of the search strategies rely on the identification of displaced objects, namely SM particles (charged leptons, photons, hadrons, jets) that are produced at a location away from the primary vertex (PV) where the hard pp collision takes place. The secondary vertex at which the decay of the LLP occurs is referred to as a DV. As far as possible, our classification of searches is linked to the parton-level objects produced in LLP decays, which allows a relatively straightforward linkage to LLP models (as well as simplified models; see chapter 2). Borrowing the terminology from prompt searches, we consider the following categories for the analogous displaced objects produced in LLP decays: all-hadronic (jets), leptonic, semi-leptonic, and photonic. However, we caution the reader that these ‘jets’ or ‘photons’ may not be of the standard type, and so other objects may pass the selections of these analyses. The remaining searches fall in the ‘other long-lived exotics’ category, mostly consisting of non-standard tracks (DTs, HSCPs, quirks, etc), but also including some trackless signals, such as SPs and SIMPs. These

¹⁵¹ It is true that, depending on the signature, some of these caveats can be circumvented by a sensible use of existing prompt triggers. For example, photon triggers will collect displaced electrons, calorimeter/jet triggers will record displaced hadronic vertices, etc.

categories are not to be interpreted as exclusive; many models and searches could fit into several categories¹⁵². For example, [243–245] show how searches for different signatures and LLP lifetimes can be combined to cover large parts of the parameter spaces of particular UV models.

3.1. All-hadronic decays

ATLAS has several searches for displaced decays with hadronic objects, including searches for two objects decaying in the hadronic calorimeter (HCAL) [246, 247]; decays within the muon system (MS) or ID [248]; ID decays in association with large \mathcal{E}_T [237]; and ID decays in association with large \mathcal{E}_T , jets, or leptons [196]. CMS has inclusive searches for displaced jets using 13 (8) TeV data [249, 250] [251]. Moreover, the CMS displaced jets searches are relatively inclusive and so also cover LLPs with semi-leptonic decays despite having no specific lepton requirements. CMS also has a search for DV in multijet events [252]. LHCb has searches for both one [253] and two [254] all-hadronic DVs in their detector. Here the discussion is restricted to summarizing the hadronic channels, while those studies including leptons [196, 249] will be revisited in sections 3.2 and 3.3 for the fully-leptonic and semi-leptonic cases, respectively.

3.1.1. ATLAS searches. The reconstruction of displaced tracks in the ATLAS ID [255] follows a two-step procedure. In the first iteration, the default track identification algorithm is applied, which uses hits in the pixel system, semiconductor tracker, and transition radiation tracker (TRT) to reconstruct tracks with a small impact parameter. The hits not associated to a track during the first pass are used in a second run of the track finder, with loose requirements on the transverse and longitudinal impact parameters (d_0 and z_0) and the number of silicon hits that are shared (or not shared) with another track. This two-step procedure is referred to as the *large radius tracking* (LRT) algorithm by the ATLAS collaboration. Applying the LRT procedure is CPU-intensive, and thus it is only run once per data-processing campaign, on a subset of specially-requested events [255].

In searches where the LLPs decay exclusively in the ID, standard triggers are used to select events with high- p_T jets, \mathcal{E}_T , or high- p_T leptons [196, 237]. An ATLAS 13 TeV search [237] uses a standard \mathcal{E}_T trigger and an offline requirement of $\mathcal{E}_T > 250$ GeV. The 8 TeV search [196] covers a larger range of topologies, and the event must have either $\mathcal{E}_T > 180$ GeV or contain four, five, or six jets with $p_T > 90, 65$, or 55 GeV to pass the trigger. In both searches, the ID vertex is required to have at least 5 tracks and the invariant mass of the DV tracks to fullfil $m_{\text{DV}} > 10$ GeV. These searches are interpreted in the context of various SUSY scenarios involving gluinos or squarks decaying into leptons, jets and missing energy, namely *R*-parity-violating (RPV), general gauge mediation, and split SUSY. In the latter case *R*-hadrons¹⁵³ are considered. The particular LLP decay topology determines which trigger and analysis mode (specified by jet and lepton multiplicity, small/large \mathcal{E}_T , etc) has the best sensitivity. The LLPs covered by these searches are typically high mass ($\gtrsim 100$ GeV), and correspond to the direct-pair-production and HP production modes with hadronic decays (in the language of the simplified models presented in section 2). However,

¹⁵² Another important ingredient of the LLP searches is the possibility to reinterpret their results to a large variety of models, namely be able to *recast* them. While we refer the interested reader to chapter 6 it is worth mentioning here that many existing searches publicly provide useful recasting information, such as efficiency maps or model-independent bounds on production cross sections.

¹⁵³ *R*-hadrons form when BSM colored particles hadronize due to a lifetime larger than the hadronization scale. In split SUSY the *R*-hadrons are typically long-lived due to their decays being mediated by heavy squarks.

these searches do not have sensitivity to low-mass LLPs, especially those resulting from the Higgs, Z' , or CC production portals and then decaying hadronically.

For LLPs decaying in the HCAL or MS, dedicated *CalRatio* and *MuonRoI* triggers are employed [246–248, 256], allowing the searches to place limited requirements on the non-displaced portion of the event. We describe these triggers in more detail shortly. The efficiency of these triggers is 50%–70% for decays within the relevant geometric detector region, and negligible outside of them (see figure 3 of [248]). The results of these analyses are interpreted in terms of a $\Phi \rightarrow ss$ model, where Φ is a heavy scalar boson with $100 \text{ GeV} < m_\Phi < 1000 \text{ GeV}$ and s is a long-lived, neutral scalar decaying to hadrons with branching fractions dictated by the Yukawa coupling. This can map to Higgs or Z' production modes and hadronic decay mode in the simplified models.

The *CalRatio* trigger selects events with at least one trackless jet that has a very low fraction of energy deposited in the ECAL¹⁵⁴. These *CalRatio* jets are characteristic of an LLP that decays within or just before the HCAL. The 13 TeV analysis [246] requires two *CalRatio* jets, where the exact *CalRatio* criteria are determined using a series of machine learning techniques to optimally discriminate the displaced decay signature from QCD jets and beam-induced background. Using the simplified $\Phi \rightarrow ss$ model with $125 \text{ GeV} < m_\Phi < 1000 \text{ GeV}$ and $5 \text{ GeV} < m_s < 400 \text{ GeV}$, good sensitivity is observed for $c\tau$ between 0.05 and 35 m, depending on the Φ and LLP masses. Notably, SM-like Higgs boson decays to LLP pairs are constrained below 10% branching ratio in the most sensitive lifetime ranges, with exact limits dependent on the LLP mass [246]. The 8 TeV result also requires two *CalRatio* jets, and shows sensitivity for $100 \text{ GeV} < m_\Phi < 900 \text{ GeV}$ and $10 \text{ GeV} < m_s < 150 \text{ GeV}$ [247].

The *MuonRoI* trigger selects events with clusters of L1 regions of interest (RoIs) in the MS that are isolated from activity in the ID and calorimeters. It is efficient for LLPs that decay between 3 and 7 m transversely or 5–13 m longitudinally from the PV, for LLP masses greater than 10 GeV. After trigger selection, the ATLAS analysis in question requires either two reconstructed DVs in the MS [257] or one ID vertex and one MS vertex [248]. This ID–MS combination provides increased sensitivity to shorter lifetimes than an analysis only considering MS vertices, and shows good sensitivity to $100 \text{ GeV} < m_\Phi < 900 \text{ GeV}$ and $10 \text{ GeV} < m_s < 150 \text{ GeV}$. Decays of a SM-like Higgs boson to LLP pairs are constrained below 1% in the most sensitive $c\tau$ regions (with cross section limits as low as 50 fb). The efficiency degrades for benchmarks with higher LLP boosts or very low mass LLPs, as fewer tracks are reconstructed. Another ATLAS search includes signal regions with only 1 DV in the MS, with sensitivity to SM-like Higgs decays to LLPs extending down to branching fractions of 0.1% [258]; this search also presents constraints on a wide range of models that helps facilitate reinterpretation for other BSM scenarios. In addition, a combination of the results from this search with the results from the 13 TeV *CalRatio* search was performed in [246] for the models common to both, and provides a summary of the ATLAS results for pair-produced neutral LLPs.

Recently, ATLAS presented a new study for hadronically decaying LLPs produced in association with a leptonically decaying Z boson [259]. In this analysis, the LLP decays inside of the HCAL. The use of lepton triggers on the associated Z decay products gives sensitivity to production of a *single* low-mass LLP, whereas other searches typically require 2 DVs; it is therefore an excellent example of the utility of prompt associated objects in obtaining

¹⁵⁴ The variable used to discriminate between *CalRatio* jets and standard jets is $\log_{10}(E_{\text{HAD}}/E_{\text{EM}})$, where E_{HAD} and E_{EM} are the fractions of the measured energies of the jets appearing in the HCAL and ECAL, respectively. The trigger selects trackless jets with $\log_{10}(E_{\text{HAD}}/E_{\text{EM}}) > 1.2$, which corresponds to an electromagnetic fraction of 0.067.

sensitivity to low-mass LLPs. The model constraints are expressed in terms of a Higgs portal model where the SM-like Higgs decays to a bosonic LLP.

With the exceptions of [196, 237, 259], which require prompt activity in addition to the DV and have comparatively high trigger thresholds, the ATLAS all-hadronic analyses require two DVs, and thus are insensitive to models that produce a single DV inside the detector¹⁵⁵.

3.1.2. CMS searches. The CMS analyses [249–251] are based on a dedicated offline *displaced jet tagging* algorithm using tracker information to identify pairs of displaced jets. The triggers used here are based on large values of $H_T = \sum |p_{T,j}| = 350(500)$ GeV for 8 (13) TeV, where the H_T sum runs over all jets with $p_{T,j} > 40$ GeV and $|\eta_j| < 3.0$. The trigger for the 13 TeV analysis based on 2015 data additionally requires either two jets with $p_T > 40$ GeV and no more than two associated prompt tracks ($d_0 < 1$ mm) and the H_T threshold is lowered to 350 GeV if the two jets each have at least one track that originates far from the PV¹⁵⁶. Only events with two or more displaced jets are kept in the analysis, while those with only one are used as a control sample to estimate the prompt jet misidentification rate. For $c\tau < 3$ mm¹⁵⁷, the algorithm is inefficient as more than two tracks tend to have impact parameters less than 1 mm; for $c\tau > 1$ m the search is inefficient as most decays occur too far from the PV to form reconstructable tracks. A key difference among these searches is that the 8 TeV [251] and 13 TeV (2016 data) [250] analyses explicitly reconstruct the DV, while the 13 TeV (2015 data) [249] analysis does not.

CMS interprets the signal in several benchmark models that can be mapped to the direct pair production simplified model production mode, including a neutral LLP decaying hadronically and a color-charged LLP decaying into a jet plus a lepton. For neutral LLP pair production decaying democratically into light jets, the trigger efficiencies for $c\tau = 30$ mm are reported to be 2%, 41%, 81%, and 92% for 50, 100, 300, 1000 GeV masses, respectively. It is evident that the requirements on H_T and on $p_{T,j}$ make the search inefficient for low LLP masses. Indeed, a phenomenological recast of the 8 TeV analysis [251] in terms of rare decays of a SM-like Higgs boson with a mass of 125 GeV Higgs sets very mild bounds for LLP masses below $m_h/2$ [260]. Thus, the CMS search has limited sensitivity to low-mass, hadronically decaying LLPs through the Higgs, Z' , or CC simplified production modes.

As mentioned above, CMS also has a search for DV in multijet events [252], which was released near the time of the final editing of this manuscript.

3.1.3. LHCb search. The LHCb searches [253, 254] trigger directly on DVs with a transverse distance of $L_{xy} > 4$ mm) with four or more tracks, vetoing dense material regions in which hadronic interactions with the detector can mimic LLP decays. The trigger thresholds are, however, low. For example, the invariant mass of particles associated with the vertex must exceed 2 GeV and the scalar sum p_T of tracks at the vertex must exceed 3 GeV. Jet reconstruction is then performed offline with standard algorithms. The benchmark model used by these searches is a scalar particle decaying to two neutral LLPs, π_v (dark or ‘valley’ pions), which corresponds to the Higgs simplified model with hadronic decay modes. The parent particle can be either a SM-like 125 GeV Higgs [253] or a Higgs-like scalar with mass in the 80–140 GeV range [254]. The search is performed for π_v masses between 25 and

¹⁵⁵ This may be the result of a signal that produces two DVs, but the lifetime is sufficiently long that only one DV appears inside the detector.

¹⁵⁶ In this case this is defined by requiring that the transverse impact parameter significance $|d_0|/\sigma_{d_0}$ have a value greater than 5.

¹⁵⁷ Note that in principle the low $c\tau$ regime can be covered with standard triggers, however they require higher H_T thresholds.




Figure 4. Comparison of the ATLAS [196], CMS [251] and LHCb [253] reaches for dark pions π_V decaying into jets. The CMS result is taken from the recast done in [260] of the 8 TeV analysis [251]. In the shaded regions $B(H \rightarrow \pi_V \pi_V)$ is constrained to be below 50%. Note that the ATLAS reach extends to higher masses as well; the plot was produced using the benchmark scenarios presented in [196], hence the meaningful bound is on the lifetimes. Taken from [253].

50 GeV and decay lengths between 0.6 and 15 mm. It is expected that LHCb will extend their coverage to shorter lifetimes by improving the understanding of the material and to lower masses by using fat-jets and jet-substructure to access larger boosts [261]. In principle, the search is also sensitive to direct pair production of LLPs.

Because of the low thresholds, the LHCb search focuses on low-mass LLPs with short lifetimes, for which it has excellent sensitivity. However, its sensitivity for other signatures is limited by the geometry of the detector and the LHCb luminosity compared to ATLAS and CMS. A model-dependent direct comparison among the LHCb, ATLAS and CMS reaches for the Higgs production mode decaying into dark pion LLPs can be seen in figure 4.

3.1.4. Summary. Searches in hadronic final states do not currently cover LLP parent masses below ~ 100 GeV in a comprehensive way. This is typically due to the large $p_{T,j}$ requirements at the trigger level, with an exception being the DV reconstruction at LHCb. Additionally, the powerful ATLAS searches for LLPs decaying in the HCAL or MS require two LLP decays in the detector, meaning that as of this writing there is no sensitivity to singly produced long-lifetime LLPs with hadronic decays¹⁵⁸. While the existing searches are typically sensitive to both direct pair production and HP production of LLPs, not all of the searches provide benchmarks with a variety of LLP production kinematics and boost.

A potential way to extend the sensitivity of current analyses is to use other existing triggers exploiting such things as VBF production modes, leptons, E_T , etc, to trigger on associated *prompt* objects and perform the hadronic DV reconstruction offline. The ATLAS 8 TeV study [196] does employ multiple triggers (such as lepton triggers), but in each case the

¹⁵⁸ Highly-inclusive searches for single LLPs decaying in the ATLAS MS have been proposed [262], finding that backgrounds are appreciable and need to be controlled using data-driven methods.

triggered object must be associated with the DV (for a lepton trigger, the lepton must originate from the DV). If, instead, a prompt lepton or VBF trigger were used with the offline reconstruction of a separate displaced object, sensitivity could be recovered to low-mass hadronic DVs in a variety of simplified models, including Higgs production (via VBF or VH associated production modes) [260, 263] or CC production (in association with a prompt lepton) [264]. In particular, triggering on associated prompt objects would improve the efficiency of reconstructing low-mass hadronic LLPs produced in the decays of a SM-like 125 GeV Higgs. As there is no theoretical lower limit on the masses of light neutral LLPs, it is imperative to lower the LLP mass coverage as much as possible. If at all possible, a dedicated *online* reconstruction of DVs would allow for a further reduction on the p_T threshold, giving sensitivity to light LLP masses.

3.2. Leptonic decays

All three experiments have searches for a pair of leptons coming from a DV [196, 265–269]. CMS also has a search requiring exactly one isolated muon and one isolated electron (i.e. events with additional isolated leptons are discarded) with large transverse impact parameters ($0.2 \text{ cm} < |d_0| < 10 \text{ cm}$), but without any other additional requirement including, for example, that the reconstructed tracks do not need to point to a common vertex [197]. This loose selection makes the search sensitive to a variety of new physics scenarios. Light and boosted LLPs can decay into collimated light leptons, dubbed *LJs* [127], which are searched for at both CMS [270, 271] and ATLAS. ATLAS has searches for both displaced [272, 273] and prompt LJs [274]. The LHCb collaboration also looks for light, neutral LLPs decaying into $\mu^+ \mu^-$ pairs by studying *B*-meson decays to kaons, for exclusive decay channels for both neutral [267] and charged [268] *B*-mesons, as well as dark photons that decay into muon pairs [269].

3.2.1. CMS searches. The CMS searches trigger on leptons reconstructed using information from either the tracker [265] or the muon chambers [266], where the latter search uses only muons. In the tracker-based analysis, the LLP is reconstructed by forming pairs of charged leptons (where muons are required to have opposite signs), with p_T cuts of 26 GeV for muons, and 36 (21) GeV for the leading (subleading) electron. This yields slightly larger efficiencies in the muon channel. The transverse impact parameter $|d_0|$ needs to be 12 times larger than its uncertainty σ_d (approximately corresponding to a distance $\gtrsim 200 \mu\text{m}$) to reject prompt backgrounds. In the MS-based analysis, muon candidates are reconstructed using hits in the muon chambers, and no information from the silicon tracker is used. In order to avoid biases from a loose beamspot constraint in the seeding step, these muons undergo an additional refit step. These candidates are referred to as *re-fitted stand-alone* muons, and they need to fulfill $p_T > 26 \text{ GeV}$, $|\eta| < 2$, and to be separated by $\Delta R > 0.2$. More importantly, these candidates are rejected if they can be matched to a $p_T > 10 \text{ GeV}$ track in the inner tracker, which efficiently excludes prompt muons and also renders this study fully complementary to the tracker-based one. Both these searches are interpreted in terms of decays of an SM-like Higgs H ($H \rightarrow XX, X \rightarrow l^+l^-$) and RPV squarks, covering proper lifetimes of $0.01\text{--}10^5 \text{ cm}$ for the Higgs scenario, and $0.1\text{--}10^4 \text{ cm}$ for the SUSY case. The difference in the lower reach of $c\tau$ is due to the larger boost factor of the Higgs. These benchmarks map to the direct pair production, HP and Higgs production simplified models, with flavor-conserving leptonic decays of the LLP. There is good sensitivity down to relatively low masses (LLPs of masses $\gtrsim 20 \text{ GeV}$ produced in Higgs decays) due to the low lepton trigger thresholds.

Additionally, CMS has a search for one electron and one muon, each with large transverse impact parameter ($200\ \mu\text{m} < |d_0| < 10\ \text{cm}$) [197]. Events are selected using a dedicated trigger for $e\mu$ pairs that applies a p_{T} cut on the leptons (42 GeV for electrons, 40 GeV for muons) but, unlike standard triggers, places no restriction on the maximum d_0 or distance from the PV. Events with exactly one muon and exactly one electron are kept, and then separated into ‘prompt’, ‘displaced control’ and ‘signal’ regions, defined as $|d_0| < 100\ \mu\text{m}$, $100\ \mu\text{m} < |d_0| < 200\ \mu\text{m}$, and $|d_0| > 200\ \mu\text{m}$, respectively. This selection makes the signal region almost free of leptons coming from SM processes, with rare tau-leptons, B -mesons or D -mesons as the largest remaining background.

Although in the original search the results are interpreted in the context of long-lived RPV stops (excluding masses below 870 GeV for $c\tau = 2\ \text{cm}$)¹⁵⁹, this search has been shown to be sensitive to many scenarios, including long-lived staus in gauge mediated SUSY breaking [205] and RHNs [45]. Indeed, this search has sensitivity to LLPs produced via any of the simplified model production modes and (semi)-leptonic decays that give exactly one electron and one muon. On the other hand, models where LLPs decay only to *either* muons *or* electrons (e.g. $\tilde{\mu} \rightarrow \mu\tilde{G}$) are unconstrained by this search. Furthermore, same-sign lepton signatures and signatures with additional leptons are not constrained by the current search but could be covered by extensions of the search [45, 205]. Due to the generality of tau-specific models, searches for hadronic tau channels is also desirable. This search has sensitivity to relatively low-mass LLPs; however, the 8 TeV analysis [275] has lower thresholds ($p_{\text{T}} > 22\ \text{GeV}$ on both leptons) albeit with a requirement for shorter decay distances ($|d_0| < 2\ \text{cm}$), and so has superior sensitivity to very low- p_{T} displaced signals. Maintaining low trigger thresholds is necessary to obtain sensitivity to the lowest-mass leptonic LLP signals.

3.2.2. ATLAS searches. The primary ATLAS search for displaced leptons [196] triggers on muons without an ID track, electrons, or photons¹⁶⁰. The trigger and offline p_{T} criteria are relatively high, requiring one of the following: one muon of at least 50 GeV; one electron of at least 110 GeV; one photon of at least 130 GeV; or two electrons, photons, or an electron and a photon with minimum p_{T} requirements for both objects in the 38–48 GeV range. The DV is formed from opposite-sign leptons, irrespective of flavor, and needs to be located more than 4 mm away from the PV in the transverse plane. DVs in regions with dense detector material are vetoed to suppression backgrounds from converted photons (e.g. $\gamma p \rightarrow e^+e^-p$). This search is in principle sensitive to events with a reconstructed DV mass $m_{\text{DV}} > 10\ \text{GeV}$, but the high p_{T} requirements for the leptons restrict the sensitivity to low-mass LLPs.

ATLAS has also recently released a search for pairs of muons that correspond to a DV [276]. This search is sensitive to LLP decays that occur sufficiently far from the IP that the muons are reconstructed only in the MS. The analysis has four separate trigger pathways: $E_{\text{T}} > 110\ \text{GeV}$; one muon with $p_{\text{T}} > 60\ \text{GeV}$ and $|\eta| < 1.05$; two muons with $p_{\text{T}} \gtrsim 15\ \text{GeV}$ and $\Delta R_{\mu\mu} < 0.5$; or three muons with $p_{\text{T}} > 6\ \text{GeV}$. Thus, the search has sensitivity to final states with high and low masses (down to $m_{\mu\mu} = 15\ \text{GeV}$), as well as with various lepton multiplicities. Offline selections require the muons to have $p_{\text{T}} > 10\ \text{GeV}$ and opposite charge, and the search is efficient at reconstructing muons

¹⁵⁹ We note that a CMS prompt search for LQs has been recasted using the same model, finding stringent constraints for lifetimes below a few millimeters. This reinterpretation is discussed in detail in section 6.7.

¹⁶⁰ Electrons with large transverse impact parameters d_0 tend to be missing a track at trigger level and are reconstructed as photons.

for transverse impact parameters up to $|d_0| = 200$ cm. Muons are also required to satisfy isolation requirements from jets as well as from nearby tracks. The search puts constraints on a SUSY model and a model of dark photon production in Higgs decays; this can be used to determine sensitivity to all of the simplified production modes from chapter 2 in the $\mu\mu$ decay mode (or $\mu\mu$ in association with other objects). It also demonstrates how combining many trigger pathways and loose selection requirements can enhance sensitivity to a wide range of LLP models.

3.2.3. LHCb searches. LHCb has a search that looks for the direct production of both promptly decaying and long-lived dark photons [269]. As a result of the direct production, dark photons do not tend to be highly boosted in the transverse direction. Events¹⁶¹ are required to have a single muon with $p_T > 1.8$ GeV, or two muons with a product of transverse momenta $\gtrsim (1.5 \text{ GeV})^2$. The displaced search constrains previously uncovered dark photon parameter space around masses of ~ 300 MeV.

The LHCb searches for displaced leptons in rare B meson decays [267, 268] rely on standard techniques to identify the B^\pm decay vertex and the kaons and pions in the event, and the di-muon invariant mass $m(\mu^+\mu^-)$ variable is scanned for excesses. The $X \rightarrow \mu^+\mu^-$ vertex is not required to be displaced from the B^\pm vertex, and thus the constraints apply to both prompt and LLPs. The analysis probes LLP masses of 214 (250) MeV $< m_X < 4350$ (4700) MeV for the $B^0 \rightarrow K^*\mu^+\mu^-$ ($B^+ \rightarrow K^+X$, $X \rightarrow \mu^+\mu^-$) process, with the mass range being limited by kinematics.

3.2.4. Lepton-jet searches. Searches for LJs are focused on $\mathcal{O}(\text{GeV})$ LLP masses and distinctly boosted signatures, and thus we treat them separately.

The ATLAS 8 TeV search [272] considers three types of LJs: those containing only muons, only electrons/pions, or a mixture of the two. The muon and electron/pion LJs can contain either two or four leptons, while the mixed LJ must contain two muons and a jet consistent with a displaced electron/pion pair. As these signatures contain relatively soft leptons, the ATLAS 8 TeV analysis uses a trigger that requires three muon tracks in the MS with $p_T > 6$ GeV. There is a built-in limitation to this trigger, which is that the L1 requirement of three separate muon RoIs makes it only sensitive to topologies with two LJs in which one LJ has a wide enough opening angle between two muons to create two level-one RoIs. For the electron LJs, when the electrons are produced in the HCAL they are indistinguishable from a hadronic decay and thus the CalRatio trigger is used.

In the 13 TeV ATLAS analysis [273], a *narrow-scan* muon trigger is additionally used. This trigger starts off by selecting events with one muon with $p_T > 20$ GeV, then requires a second muon with $p_T > 6$ GeV within $\Delta R = 0.5$ of the leading muon.

Both the 8 and 13 TeV ATLAS searches are interpreted for Higgs-like scalar particles (with masses of 125 and 800 GeV) that decay effectively into either two or four lepton pairs, with each lepton pair assumed to come from a low-mass ‘dark’ photon, γ_D . The ATLAS result excludes exotic Higgs branching ratios below 10% for dark photon lifetimes $2 < c\tau < 100$ mm. Note that here γ_D is also allowed to decay to pions and so the results can also be interpreted for hadronically and semi-leptonically decaying LLPs. This corresponds to

¹⁶¹ For the prompt dark photon search, events are reconstructed at trigger level so that all online reconstructed particles are recorded, while the rest of event information is discarded [277]. The prompt search constrains entirely new territory above 10 GeV.




Figure 5. Comparison of the lepton-jet searches at ATLAS [272] and CMS [271] with respect to a dark photon scenario [155] vis-a-vis dark photon limits coming from low-energy experiments. Figure taken from [271].

the Higgs production mode in the simplified models proposal with an admixture of flavor-conserving leptonic and hadronic LLP decays.

The CMS LJ search focuses on fully muonic LJs and has been performed with both the 8 TeV dataset [270] and part of the 13 TeV dataset [271]. The 13 TeV search is sensitive to di-muon parent particle masses up to 8.5 GeV. Events are selected with a di-muon trigger with standard isolation requirements. Further selection requires at least four muons, forming a minimum of two opposite-charged pairs. CMS uses a benchmark model with scalars decaying into either lighter scalars or dark photons, with varying scalar and dark photon mass. For the case of a 125 GeV Higgs they can exclude an exotic branching ratio of below 0.1% for some parameter points. The CMS search can be compared with the ATLAS results, as can be seen in figure 5. We note that this study includes sensitivity to both prompt and displaced muonic lepton jets.

3.2.5. Summary. To summarize, the lepton searches rely on fairly standard lepton identification, with vertex reconstruction being performed mostly offline. Searches for leptonically decaying LLPs typically enjoy low trigger thresholds and good sensitivity to LLP production rates. Extending the success of the leptonic LLP program to future LHC running will necessitate maintaining low-threshold triggers for displaced leptons in a high-luminosity environment; this is a major challenge but one that must be overcome. Another outstanding challenge is coverage of LLP decays to τ leptons, which lie at the interface between hadronic and leptonic searches. Such decays are very well motivated from the theoretical point of view, as a Higgs-like scalar can typically decay about 300 times more often into $\tau^+\tau^-$ if

kinematically allowed, and also one could have large rates into mixed decay modes such as $\tau^+\mu^-$. A displaced hadronic τ is a striking object, and most likely will have few backgrounds. Hence, limits on exclusive displaced τ s would be of utmost importance¹⁶².

As the leptonic searches explicitly require opposite-sign leptons, the same-sign lepton signature (motivated from Majorana neutrinos; see the LHCb search in section 3.3, or heavy, doubly-charged LLPs) is currently neglected. Furthermore, the sensitivity of the CMS search for two high- $|d_0|$ leptons is currently only sensitive to opposite sign $e\mu$ pairs, with a veto on additional leptons. Relaxing these requirements would greatly enhance sensitivity to certain scenarios, especially the simplified models with flavor-conserving leptonic LLP decays.

Lepton-jet searches currently cover only final states with at least two LLPs and some muons in the final state¹⁶³, and the same statement currently applies to both the LHCb searches for dark photons [269, 277] and for LLPs produced in B -meson decays. The existing ATLAS LJ studies express their results in terms of specific dark photon models¹⁶⁴, which makes it complicated to apply the results to other models. We refer the reader to section 6 for a further discussion of this topic. In extending LJ searches, it would be beneficial to have additional searches with a single LJ or low-mass, leptonically-decaying LLP (which are motivated in models with hidden sectors and Majorana neutrinos, for example in [43, 135]). In addition, the status of coverage in the intermediate mass-transition region between ‘standard’ displaced lepton pairs and LJs is unclear, and may potentially harbor a gap; adopting the simplified model approach for leptonic LLP decays with masses varying between the GeV and weak scale would ensure comprehensive coverage of low-mass leptonically decaying LLPs.

Finally, we comment on gaps in sensitivity to very low-mass leptonically decaying LLPs. A benchmark LLP model is the heavy neutral lepton (HNL), which corresponds to the CC production mode with (semi-)leptonic LLP decays in our simplified model framework. HNLs constitute an important physics case that leads to multi-lepton displaced signatures from W decays, with nice prospects at ATLAS and CMS (see, e.g. [43, 48, 49]). While previous searches were not sensitive to this scenario due to either high- p_T requirements or the requirement of two DVs in the same event, the presence of a prompt lepton from the W allows the relaxation of these requirements in a dedicated analysis. Moreover, the prospects of triggering on a prompt lepton in such searches was studied recently in [48] and demonstrated in a prompt search in [278]¹⁶⁵. The identification of two leptons from the vertex is a powerful discriminant against backgrounds from meta-stable particle decays and hadronic interactions in material. This permits a potentially cleaner exploration of the lower HNL mass range (3–6 GeV) than in the semi-leptonic channel (see section 3.3) despite the lower branching ratio. It should be noted that HNL models can predict LLP decays to all three lepton flavors (either democratically or hierarchically), necessitating the capability to reconstruct displaced leptons of all flavors, including taus.

¹⁶² We note that if the τ originates from outside the tracker, the hadronically decaying taus are indistinguishable from other displaced hadrons. For instance, in the ATLAS search utilizing muon RoIs [248] the results are interpreted for a model with a scalar particle with Higgs-like couplings to SM fermions, which includes a branching fraction into $\tau^+\tau^-$. However, if the τ originates from the ID, the low number of tracks associated to it (one to three) will not fulfill the requirements of the ATLAS study of five or more tracks associated to a DV.

¹⁶³ The ATLAS 8 TeV search [272] included a search channel with two electron-only LJs, but the performance was poor and it was excluded from the final result.

¹⁶⁴ Recall that the LJ studies also consider the $\gamma_D \rightarrow \pi^+\pi^-$ decay mode.

¹⁶⁵ Note that the displaced large transverse impact parameter $e\mu$ CMS search [197] fails to cover this scenario due to the aforementioned lepton veto, which eliminates sensitivity to the tri-lepton signals discussed in [43], as well as relatively high lepton p_T trigger thresholds compared to the kinematics of 4-body W decay.




Figure 6. LHCb reach for displaced semi-leptonic decays. Taken from [279].

3.3. Semi-leptonic decays

As semi-leptonic signatures include aspects of both hadronic and leptonic LLP decays, many of the previously discussed searches can partially cover these cases, and some do so explicitly. For instance, the ATLAS search for electrons and muons accompanied by tracks [196], the inclusive CMS search for DVs [249] (which contains a specific model interpretation called ‘*B*-lepton’ addressing precisely this channel), and the search for a large impact parameter $e\mu$ pair by CMS [197] are all inclusive with respect to other hadrons produced in the LLP decay, provided the leptons are sufficiently isolated¹⁶⁶. In addition, LHCb has dedicated searches for semi-leptonically decaying LLPs [279] and semi-leptonic decays of long-lived Majorana neutrinos coming from B^- mesons [280]. The two CMS searches [197, 249] need no further explanation for how they cover semi-leptonic LLPs because of their very inclusive nature (see section 3.2.1), but we now describe the rest of these searches in some detail.

3.3.1. LHCb searches. The LHCb search for semi-leptonic LLP decays selects events with a muon trigger, then reconstructs a DV offline [279]. The results are interpreted in terms of four distinctive topologies: single LLP production in association with a new particle (in this case a gluino), double LLP pair production via direct pair production, Higgs decay, or via squark pair production (HP). The LLP then decays to two quarks and a muon (which maps to the jjl simplified model decay). Material regions are vetoed for the DV, which results in the dominant background arising from heavy flavor production either directly or from W/Z decays. The signal discrimination is obtained from a multivariate analysis based on the muon p_T and impact parameter, and subsequently the search is optimized based on the LLP reconstructed mass and the muon isolation. This study is sensitive to low-mass LLPs with lifetimes between 1.5 and 30 mm, as can be seen in figure 6.

The LHCb search for Majorana neutrinos [280], N , probes Majorana neutrinos produced in leptonic B decays, $B^\pm \rightarrow \mu^\pm N$. The Majorana neutrino subsequently decays exclusively to

¹⁶⁶ Note that the lepton isolation affects most of the semi-leptonic searches.

$N \rightarrow \mu^\pm \pi^\mp$; both prompt and displaced decays are considered¹⁶⁷. A same-sign muon requirement, along with the reconstruction of the N and B meson masses, greatly reduces backgrounds to the search. The sensitivity of the search is limited by the restriction to muons in the final state (so models that predominantly decay to e or τ are not constrained), and the same-sign lepton requirement gives sensitivity only to lepton-number-violating processes. More inclusive searches looking for additional N production modes [282], decays with opposite-sign leptons, or searches targeting decays of heavier mesons like B_c [283] could also improve the sensitivity to semi-leptonically decaying LLPs.

3.3.2. ATLAS search. The ATLAS search for semi-leptonic LLP decays [196] looks for a vertex with a lepton accompanied by tracks. This search uses the same trigger as the dilepton vertex search described in section 3.2. The DV is required to have a lepton as well as at least four additional associated displaced tracks, and the invariant mass of the tracks must exceed 10 GeV. Thus, the search in principle can have sensitivity down to masses $\gtrsim 10$ GeV, although the high p_T threshold for the displaced electron/muon typically limits sensitivity to low-mass LLPs; the vertex must contain a muon with $p_T > 55$ GeV or an electron with $p_T > 125$ GeV.

3.3.3. Summary. When considering the application of inclusive hadronic or leptonic searches to semi-leptonic LLP decays, it is important to understand how the simultaneous presence of leptons and jets in the signal can degrade the sensitivity. For instance, prompt jet searches explicitly can remove non-standard jets through jet cleaning cuts. Lepton isolation criteria can severely reduce the signal acceptance for a highly-boosted LLP decaying into a lepton and a jet, and they might also veto extra tracks in the events. Thus, boosted semi-leptonic decays (as might be found in the displaced decay of a low-mass, RHN produced via W decay) may not be covered by existing searches.

One of the major gaps in semi-leptonic LLP searches is at the smallest LLP masses. In this case, it can be challenging for the leptons from LLP decays to pass trigger thresholds and/or isolation criteria; backgrounds from heavy-flavor and other processes are also higher for semi-leptonic processes than fully leptonic ones. However, there are very good motivations for low-mass semi-leptonic LLPs from the HNL benchmark model¹⁶⁸ introduced in section 3.2.5, which predicts LLPs for HNLs of masses below 30 GeV. The signature of HNLs from W decays with displaced semi-leptonic HNL decays is an important item on the search agenda of ATLAS, CMS and LHCb [38, 43, 48, 49, 284, 285]. Recently, there has also been a proposal to search for low-mass HNLs in heavy ion collisions [286]. The semi-leptonic channel has the highest branching ratio (about 50% in the relevant mass range [287]) and can therefore offer the best discovery prospects at LHC experiments for HNL masses up to 30 GeV as long as a DV mass cut of around 6 GeV is made to mitigate backgrounds from B -mesons, $m_B \sim 5$ GeV, and backgrounds from random-crossings can be suppressed. The lower end of the 6–30 GeV mass range corresponds to a non-perturbative regime for the hadronization of the HNL decay products. As the number of charged hadrons significantly affects the DV reconstruction efficiency, the validation of event generator outputs for this process is an important issue currently being addressed by the community (see e.g. [48]). The ability of LHCb to trigger directly on the HNL decay products and better reconstruct

¹⁶⁷ Some care is required in interpreting the results of the search on a model with a Majorana neutrino, as the original theory interpretation is problematic [281].

¹⁶⁸ In the language of simplified models, this corresponds to the CC production mode, where the HNL LLP is produced in association with a prompt lepton. The HNL then decays semi-leptonically via the $jj\ell$ channel.




Figure 7. Summary of projected experimental sensitivities to HNLs in various experiments, in the coupling strength (U^2 for dominant mixing to ν_μ) versus mass plane. The projections labelled ‘LHC Run 3’ and ‘HL-LHC’ are for HNLs in W decays in general-purpose experiments, and the one labelled ‘Mathusla’ assumes the full HL-LHC MATHUSLA dataset. We also show the recent CMS result for the prompt tri-lepton signature [278]. Prospects at proton beam-dump experiments are also shown for an already existing experiment, NA62 [288], and for the planned experiment SHiP [289]. Existing constraints from direct searches are indicated as coloured solid lines [290–294]. The lines labelled ‘Seesaw’ and ‘BBN’ show lower theoretical constraints from the observed neutrino masses (assuming a normal hierarchy) and primordial nucleosynthesis, respectively [295]. The line labelled ‘BAU’ is an upper theoretical constraint in the ν MSM model for accounting for baryon asymmetry in the Universe while the lightest HNL is a dark-matter candidate [295].

displaced tracks can in some cases compensate for its lower acceptance and luminosity, as exemplified by a recent search for DVs composed of one muon and several tracks [279, 285]; similar arguments can be made in the case of heavy ion collisions [286]. This possibly enables LHCb to better probe the more challenging tau channel. Figure 7 shows the overall expected reach of LHC searches in the HNL coupling strength (for the muon channel) versus mass plane, using assumptions detailed in [284], similar to those in [38, 43].

The sensitivity of LHC experiments to HNLs is complementary to that of fixed-target experiments which can probe lower couplings thanks to high-intensity beams albeit at lower mass ranges (i.e. targeting HNLs from c and b decays). The CERN SPS provides great opportunities with the running NA62 experiment [296] and the planned SHiP experiment [289], which comprise a vacuum decay vessel and spectrometer tracker downstream of the target to reconstruct vertices of long-lived neutral particles¹⁶⁹. These provide the best sensitivity to HNL masses up to 2 GeV, where they probe a region of the parameter space favored in models which simultaneously explain neutrino masses, matter–antimatter asymmetry and DM [295, 298–300] (see figure 7).

3.4. Photonic decays

There are two ways in which photons coming from LLP decays do not resemble standard photons. First, they cannot be traced back to the PV, thus giving rise to *non-pointing* photons. Second, they can arrive at the ECAL at a time slightly later than expected because the LLP

¹⁶⁹ The proposed detector FASER would also have the capability to reconstruct such vertices [297].

moves slower than the speed of light; these are referred to as *delayed* photons. We note that both kinds of unusual photons can be vetoed in searches for photons that originate from the PV, and thus prompt searches typically provide weaker bounds on LLP scenarios than for promptly decaying signals. ATLAS has a search for non-pointing and delayed photons [301] using the full 8 TeV dataset, which supersedes the 7 TeV analysis [302]. In CMS there are studies for delayed photons in the ECAL [303] and for non-pointing photons detected via their conversion to e^+e^- pairs [304]. The underlying topology in all these models is the neutralino decay into a gravitino and a photon ($\chi_1^0 \rightarrow \gamma \tilde{G}$), ubiquitous in GMSB scenarios [53, 305]. Hence all these studies require large \mathcal{E}_T in the final state. This corresponds to HP production of LLPs with decays to a single photon and \mathcal{E}_T in the simplified model framework; all searches described below use the Snowmass Slopes Point 8 benchmark, which is not straightforward to map to a physical spectrum of HP masses.

3.4.1. ATLAS search. The ATLAS study [301] benefits from the capability of the liquid-argon electromagnetic calorimeters (ECALs) to measure the flight direction and the photons' TOF. Resolutions on Δz_γ , the separation between the PV of the event and the extrapolated origin of the photon, and $|\Delta t_\gamma|$, difference of the arrival time of the photon with respect to the prompt case, are as low as 15 mm and 0.25 ns, respectively. The trigger demands two photons within $|\eta| < 2.5$, with transverse energies E_T of 35 and 25 GeV. In addition, to guarantee the event comes from a proton–proton collision, a PV with five or more tracks with $p_T > 0.4$ GeV is required. The offline selection requires two photons with $E_T > 50$ GeV and $|\eta| < 2.3$, not in the barrel-endcap transition region ($1.37 < |\eta| < 1.52$), at least one of them in the barrel ($|\eta| < 1.37$) and with less than 4 GeV of energy deposited in the calorimeter in a cone of $\Delta R = 0.4$ around them (constituting the *isolation* criterion). In addition, the events are binned in \mathcal{E}_T : the $\mathcal{E}_T < 20$ GeV bin contains the prompt backgrounds, the $25 \text{ GeV} < \mathcal{E}_T < 75$ GeV bin is used as the control region, and finally the signal analysis is performed in the $\mathcal{E}_T > 75$ GeV bin. This study covers lifetimes from 0.25 to 100 ns in the GMSB framework, the lower limit being a hard cut-off imposed experimentally, as the similarity between background and signal samples in that region makes discrimination rather difficult. The excluded signal rates in this range of lifetimes vary between 0.8 and 150 fb, with the best-constrained value obtained for $\tau \sim 2$ ns.

3.4.2. CMS searches. The CMS study of delayed photons [303] follows a similar approach to ATLAS. The main difference is that it demands only one photon with $p_T > 80$ GeV, but in addition two jets are required. Furthermore, the vector sum of \mathcal{E}_T and E_T^γ , which is denoted $\mathcal{E}_{T\text{no}\gamma}$, is additionally used for background discrimination. Collisional backgrounds have small \mathcal{E}_T and large $\mathcal{E}_{T\text{no}\gamma}$, while the non-collisional backgrounds are characterised by large \mathcal{E}_T and small $\mathcal{E}_{T\text{no}\gamma}$. For the signal events the two variables are large, hence they are both requested to be larger than 60 GeV. The time resolution is 0.372 ns, slightly worse than in the optimal scenario of the ATLAS search. Their reach in lifetimes lies in the 2–30 ns range, excluding signal rates of 10–30 fb.

The CMS study of non-pointing photons [304] relies on the photon converting to e^+e^- pairs. It requires two photons, two additional jets, and $\mathcal{E}_T > 60$ GeV. The photon trajectory is obtained from the conversion vertex as the line segment along the momenta of the e^+e^- track pair, and the impact parameter, $|d_{XY}|$, is defined as the closest distance between the photon and the beam axis, which can be determined within approximately 1 mm. A comparison of the reach of these 8 TeV studies, as well as those using the 7 TeV dataset, can be found in figure 8.




Figure 8. Summary of the $\gamma + E_T$ searches from ATLAS and CMS, displayed assuming the same GMSB model. Taken from [303].

3.4.3. Summary. The gaps in these studies are straightforward to identify. The requirement of large E_T is due to the fact that all of these studies have an underlying theoretical picture of neutralinos decaying into gravitinos and photons, motivated from GMSB scenarios. Hence, these searches do not cover cases without the presence of missing energy, including LLPs that decay to $\gamma\gamma$, $l\gamma$ or $j\gamma$. It may be possible to extract a constraint on such LLP decay modes due to mis-measurement of jets or the photon decay geometry which could mimic large missing energy; however, this would be sub-optimal compared to a dedicated search. With the exception of the CMS study [303] which requires two additional hard jets, all of these analyses require two displaced photons. A single displaced photon signature can occur in motivated models: it can easily arise, for example, from a very slightly mixed electroweak triplet and singlet as in SUSY theories (see the UV models in section 2). Furthermore, as discussed in section 3.1, a single LLP in the detector can also arise for very large lifetimes of neutral LLPs, which limits the reach of current searches at longer lifetimes. In such scenarios, it is possible that the photons from LLP decay can be quite soft, and obtaining sensitivity to models with single photons from LLP decay and/or low momenta may require triggering on associated prompt objects, similar to the recommendations in section 3.1.4.

3.5. Other exotic long-lived signatures

In the preceding sections, we presented analyses sensitive to LLPs decaying into objects such as jets, leptons, and photons. In many cases, however, LLPs give rise to signatures that are completely distinct from more conventional prompt signatures. In this section, we present analyses that exploit properties of other exotic long-lived signatures, such as non-standard tracks. We summarize in detail the existing searches for HSCPs; DTs; SPs and monopoles, and describe existing ideas on how to look for quirks and SIMPs.

Note that the terminology employed in some of these searches can be confusing, with occasional conflation of signatures with the underlying model. We provide a detailed explanation of each search, and we attempt a classification here based strictly on signature. We distinguish between three classes of signals: *tracks with anomalous ionization*, *tracks with anomalous geometry*, and finally *out-of-time* decays.

3.5.1. Tracks with anomalous ionization. In this category, we collect all searches for charged-particle tracks with anomalous ionization, i.e. those that are inconsistent with a charge $|Q| = e$. Here we include (i) the so-called HSCPs, which apply to stable, electrically charged particles but also charged particles that decay in the calorimeters and/or MS; and (ii) magnetic monopoles.

3.5.1.1. Heavy stable charged particles. The searches for HSCPs at CMS [4, 306] and ATLAS [223, 307, 308] rely on two key properties. First, particles that are massive and/or electrically charged with $|Q| \neq e$ have a characteristic ionization loss (dE/dx) distinctively different than SM particles. This property can be measured in the tracker. Second, HSCPs are typically heavy and move with a speed less than the speed of light, $\beta = v/c < 1$. Thus, compared to a particle with $\beta \approx 1$, they require a longer TOF to reach the outermost components of the detector (calorimeters and muon chambers). As decays or interactions of the HSCP with the material in the detector can change the electric charge of the HSCP, both CMS and ATLAS perform separate *tracker-only* and *tracker + TOF* studies in the language of CMS¹⁷⁰. The event selection relies on standard single-muon or large-missing-energy triggers. The offline selection relies on identifying the signal events from quality requirements on the tracks using discriminator variables built from track observables.

The HSCP search conducted by LHCb [309] is slightly different. Instead of exploiting dE/dx and TOF, they use the lack of radiation in the ring imaging Cherenkov detector (RICH). Events are required to pass a high- p_T single muon trigger ($p_T(\mu) > 15$ GeV). Two opposite sign ‘muons’ are then required, each with p_T above 50 GeV and an invariant di-muon mass above 100 GeV, to suppress muons coming from DY production, the main background for this search. In addition, particles must have $\beta > 0.8$, set by the efficiency of the muon chambers to reconstruct slow particles. As electrons and hadrons interact more with the calorimeter than an HSCP, a deposit in the calorimeter of less than 1% of the momentum of the particle is required.

The theoretical interpretation of a signal or limit depends on whether the HSCP carries both color and electroweak charges. If it carries a color charge, the default benchmarks correspond to *R*-hadrons, namely HSCPs that hadronize with SM particles via the strong force, e.g. gluino-gluon or quark-squark states. In the absence of a color charge, the signal is exemplified by long-lived sleptons in the context of gauge-mediated SUSY. Both ATLAS [307] and CMS [4] studies employ these two scenarios, while LHCb [309] uses a stau benchmark model¹⁷¹. Finally, CMS also looks for HSCPs coupling only to hypercharge (and hence possessing only couplings to γ and Z), while ATLAS has a search inspired by electroweakinos in SUSY: it considers the associated production of a neutral and an electrically charged LLP (chargino-neutralino), and thus only one HSCP plus missing energy are required. These scenarios correspond, in our simplified model framework, to the direct

¹⁷⁰ ATLAS measures $\beta\gamma$ from dE/dX and β from TOF and extracts an independent mass, m_β and $m_{\beta\gamma}$, from each measurement.

¹⁷¹ The ATLAS *R*-hadron searches using the 13 TeV dataset have recently been presented in [223].

production of LLPs with electric or color charges and that do not decay, or decay at very large distances compared to the tracking volume.

To summarize, these searches present no obvious weak points. Standard triggers and tracking algorithms are used, and the analysis methods are well-understood and have been extensively validated against data. HSCP signatures do not suffer from the low-mass gap of many neutral LLPs due to constraints from LEP and other low-energy experiments. However, mCPs are not covered by these searches and require dedicated detectors (see section 5.3.3). While the HSCP search strategies are generally robust, we encourage the experimental collaborations to continue pursuing improvements for these searches. The small number of signal events that would be produced for HSCPs above current limits render the sensitivity highly dependent on the understanding of the background and control of the systematics.

3.5.1.2. Magnetic monopoles. ATLAS [215] has a dedicated search for highly ionizing particles (HIPs), which encompass a variety of new physics scenarios, such as magnetic monopoles, dyons (particles with both magnetic and electric charge), stable microscopic black holes, etc. For the sake of concreteness, we focus on magnetic monopoles but the interpretation in terms of other models is straightforward¹⁷².

The main phenomenological feature is that magnetic charge is quantized in units of $g_D = 2\pi/e \approx 68.5$. Hence, a magnetic monopole behaves as a particle with at least 68.5 electron charges, leading to an unusually large ionization power in matter, so that they would quickly stop in the detector because HIPs lose energy at spectacular rates. Because of the large QED coupling of magnetic monopoles, a perturbative calculation of the cross section is invalid and there is no accurate determination of the production rate, but a naïve Drell–Yan production cross section is provided for the purposes of comparison. The specifics of the detector restrict the sensitivity of this search to magnetic charges $g < 2g_D$ because a large fraction of the monopoles stop in the material upstream of the ECAL, as the latter is used for the L1 trigger [310]. We note that larger magnetic charges can be tested by the MoEDAL experiment [311], which is described in detail in section 5.3.2. Additionally, theories with monopoles can also be tested in heavy ion collisions [312].

ATLAS [215] has a dedicated trigger for HIPs based on identifying relevant RoIs in the ECAL and subsequently counting the number of hits in the TRT. As well, the fraction of TRT hits that are high-threshold (HT), meaning that they have an ionization larger than ~ 3 times that expected from a SM particle, is used as a discriminant. The analysis selects events based on the fraction of TRT-HT hits matched to an EM cluster deposit, and how the energy deposits are distributed in the different layers of the ECAL. It is important to note that due to the lack of a consistent theory, the signal simulation is performed by re-scaling Drell–Yan production at leading order and assuming no coupling to the Z boson. The HIPs are assumed to have either spin-0 or 1/2. The spin does not affect the interaction with the material, but the angular distributions are different according to angular momentum conservation (keeping in mind that there is, of course, no perturbative theoretical prediction for the angular distribution). Cross section limits for $0.5 < |g|/g_D < 2$ are set for masses in the 890–1050 (1180–1340) GeV range for the spin-0 (1/2) case.

The coverage in LHC Run 2 of magnetic monopoles in the $g_D - \sigma$ plane is displayed in figure 9.

¹⁷² At the time of writing there was no public result on a monopole search from CMS.




Figure 9. Comparison of the MoEDAL, ATLAS and CMS reaches for magnetic monopoles. The curves assume three signal events and a total integrated luminosity of 50 fb^{-1} for 13 TeV collisions. Updated version of existing figure in [310]. Note that the CMS curve relies on the expected performance of their detector.

3.5.2. Tracks with anomalous geometry. In this category we group searches where the tracks have an anomalous geometry, namely they disappear (track \rightarrow nothing), or follow non-standard trajectories (such as quirks). There are additional anomalous track signatures that we do not cover such as kinked tracks, where a charged LLP decays to another charged particle that travels at a non-zero angle with respect to the LLP direction [112, 313–315]; however, there are currently no dedicated searches for such signatures.

Within this category we also have the emerging jets signature that has been recently studied by CMS [316]. Since emerging jets arise from dark sector radiation, they are described in detail in chapter 7 in conjunction with the theoretical and phenomenological aspects of dark showers.

3.5.2.1. Disappearing tracks. Massive charged particles traveling in the detector can decay to a lighter, almost mass-degenerate neutral state, emitting a soft SM charged particle (typically a pion or a muon). While at first glance a small mass gap naïvely seems like a hallmark of tuning, near degeneracies often occur naturally as a result of a symmetry. In fact, electroweak symmetry generically leads to small mass splittings between components of a single electroweak multiplet. For example, $\mathcal{O}(100 \text{ MeV})$ splittings arise between the different components of an electroweak multiplet [64, 128] due to EW gauge boson loops¹⁷³. If the SM particle is sufficiently soft it cannot be reconstructed, and then a charged track seems to vanish: this is thus referred to as a DT¹⁷⁴. The actual lifetime of the charged particle is highly sensitive to the precise value of the mass splitting. For instance, the well studied cases of a fermionic doublet with $Y = 1/2$ and a fermionic triplet with $Y = 0$, reminiscent of a Higgsino and Wino in supersymmetry, respectively, have mass splittings of $\Delta = 355$ and 166 MeV , up to small corrections, but the corresponding $c\tau$ values differ by almost an order

¹⁷³ For a single fermion multiplet, the splitting can only be altered by higher-dimensional operators, and thus it is harder to vary Δ from the 1-loop EW value. For other cases, such as mixing with additional particles, the actual splitting can differ more substantially from this 1-loop EW value.

¹⁷⁴ If the charged particle could be reconstructed this case is often referred to as a *kinked track*. However, as the kinked portion has a very large impact parameter, without a serious attempt to capture the kink these tracks, too, simply disappear.

of magnitude: 6.6 mm versus 5.5 cm¹⁷⁵. This is because the lifetime, $c\tau$, depends on the third power of the mass splitting in these scenarios when the charged LLP decays to a charged pion [64, 128].

Before 2017, both ATLAS [317] and CMS [217] required a track to travel about 30 cm in order to be reconstructed, giving good coverage of the Wino scenario. This 30 cm value corresponds to four hits at ATLAS, three in the pixel layers plus one in the silicon tracker, and to seven hits in the pixel and trackers of CMS. The search employs a trigger requiring an ISR jet against which the charged particle recoils, along with the presence of large \cancel{E}_T . The DT is reconstructed offline and needs to fulfill quality criteria (isolation, p_T threshold, etc). A phenomenological study [99] has shown that reducing the distance from 30 to 10 cm would give coverage to the elusive Higgsino scenario, moving the expected reach up to 400 GeV, surpassing the expected mono-jet reach of 250 GeV [318–320]. Later, ATLAS presented a study [218, 321] using 13 TeV data and exploiting the presence of a new innermost pixel layer (IBL). This addition allows for all four hits to be in the pixel, with the outermost pixel layer now at 12.25 cm, enhancing sensitivity to lower values of $c\tau$. The summary for DTs at ATLAS for the Wino case can be seen in the left panel of figure 10, while in the right panel we show the constraints for Higgsinos from [99]. CMS also has a DTs search using 2015 and 2016 data at a center-of-mass energy of 13 TeV [322].

At LHCb the prospects for a DT analysis with the present detector are poor. Currently, the momentum of the track can only be measured if the particle passes through the tracking station (TT), which is about 3 m away from the IP. Particles decaying in the vertex locator (VELO) or RICH1 system will not leave a fully-measurable track and will be swamped in a background of SM processes such as kaon decays, which would give a similar signature in the detector components before the magnet. Detector improvements (additional magnets, better PID at low momentum, additional layers) might lead to some sensitivity for $\mathcal{O}(\text{cm})$ lived tracks, a golden opportunity for potential LLP discoveries at LHCb in the HL-LHC era.

To summarize, the search for DTs presents a few challenges. Using an ISR jet trigger means a price is paid in terms of signal efficiency. For example, [99] has shown that significantly lowering the p_T threshold of the jet or directly triggering on the momentum of the DT¹⁷⁶ would lead to a factor of two increase in the number of signal events. It is also clear that better access to lower lifetimes is needed; this may only be possible, for instance, by adding new tracking layers as close as possible to the beampipe (and/or having double layers instead of single ones).

3.5.2.2. Strongly interacting massive particles. SIMPs can be motivated by astrophysical observations of DM that do not fully agree with the WIMP paradigm (e.g. missing satellites, the core versus cusp problem; see, e.g. [327–330] for further discussion). These particles are assumed to interact strongly with baryons. Consequently, the experimental signature is little to no signal in the tracker and the ECAL, and large energy deposits in the HCAL. Such a final state with trackless jets also arises in the context of emerging jets [331], and ATLAS has a trigger for trackless jets in association with a soft muon (where the muon is required to fire L1 of the trigger) [256]. Additionally, the CalRatio trigger and associated search for displaced

¹⁷⁵ While these values set a concrete physics target, we stress again that the mass splitting can be arbitrary in other corners of the BSM parameter space (even within SUSY). For instance, $\tilde{\tau} \rightarrow \tau \tilde{\nu}$ (where the stau and sneutrino masses are free independent parameters) or for scalar particles (e.g. $H^+ \rightarrow \mu^+ H^0$), where the mass splitting and the overall mass scale are set by arbitrary quartic couplings.

¹⁷⁶ While currently there are some proposals to trigger on tracks [324], those predominantly apply to standard tracks. In particular, the new fast track reconstruction (FTK) system at ATLAS requires 10 hits in ID silicon, which corresponds to a decay radius of less than 9 cm. However, pattern banks for hits including high- $|d_0|$ tracks are currently being considered out to $d_0 \sim 2$ cm [325, 326].




Figure 10. Top: summary of ATLAS disappearing-track searches as applied to a Wino (electroweak triplet) benchmark scenario [323]. Bottom: HL-LHC projected constraints on the Higgsino scenario [99].

hadronic decays (addressed in section 3.1) is designed to be sensitive to a similar signature and could provide some coverage of this signature as well. Strictly speaking, SIMPs are not a track-based signature, but we include them here because the interactions of the SIMPs with the tracker are different from usual hadrons in jets, while the calorimeter signatures are similar.

An LHC phenomenological study of SIMPs was carried out in [332]. We summarize the main points of the study here. In their setup, SIMPs interact with the SM via an attractive potential (either scalar or vector mediator) coupling SIMP pairs to $q\bar{q}$ pairs. The proposed analysis selects events with high- p_T , back-to-back jets within the tracker, exploiting the charged energy fraction within a jet to discriminate signal from background. The astrophysical experimental constraints on this scenario are compared with the expected




Figure 11. Astrophysical and collider constraints on a simple SIMP setup. Note that the relevance of the astrophysical constraints depends on the contribution of the SIMPs to the relic density. Taken from [332].

reach of this search and that of mono-jets in figure 11. Currently there is an ongoing analysis in CMS pursuing this strategy.

3.5.2.3. Quarks. Quarks are particles charged under both the SM and a new confining gauge group [182], referred to here as ‘infracolor’ (IC). The defining property of quarks is that the tree-level quark masses m_Q are above the confinement scale Λ_{IC} (and thus similar to QCD but with no light-flavored quarks), so that there is never enough local energy density to create new quarks out of the vacuum. A pair consisting of a quark and an anti-quark can live in a quantum-mechanical bound state where the constituents are separated by macroscopic distances $\ell \sim \frac{m_Q}{\Lambda_{\text{IC}}^2}$, remaining connected by an IC flux tube. The IC flux tube exerts a force on each quark that causes its trajectory to differ from the expected helicoidal ones for SM particles.

The collider phenomenology depends greatly on the size of ℓ . If ℓ is much less than an Å, the rapid emission of IC glueballs results in the quarks annihilating before ever reaching the beampipe. For large enough confinement scales, the IC glueballs can decay back into SM particles on sufficiently long time scales that they can be distinguished from prompt signatures. While in this specific case, quarks produce hidden valley [66] or emerging jet [331] signatures due to these long-lived IC glueballs (see chapter 7), we stress that elsewhere in parameter space quarks exhibit their own distinct phenomenology and are not merely a subset of hidden valleys, contrary to popular lore.

If ℓ is larger than an Å but below the mm scale, the individual quarks are not distinguished from one another. However, the pair (which is overall neutral, and therefore does not bend in the magnetic field of the tracker) appears as a single, highly ionizing straight track with missing energy aligned with it, the latter arising from mis-measurement of the track

momentum. The *D0* collaboration has searched for precisely this signature [333], requiring an additional jet for trigger purposes. This search obtained lower bounds on quirk masses of 107, 119 and 133 GeV for SU(2), SU(3) and SU(5) gauge groups, respectively. However, no extensions of this search to higher mass have been performed at the LHC, and the existing HSCP searches require low uncertainties on the track momentum that a straight quirk track would not satisfy.

Conversely, if ℓ is very large then the existence of the confining force has no effect on the quirk motions, and HSCP searches apply directly to the quirk case, with quirks charged under QCD yielding R -hadrons and uncolored quirks leading to slepton-like signatures. We refer the reader to the section 3.5.1.1 for more information.

For intermediate values of ℓ , there are interesting phenomenological prospects at the LHC that have been recently studied theoretically [239, 240, 334], but for which there are no current public searches by the LHC collaborations. The first study [239] has recast mono-jet and HSCP searches, finding bounds up to 1.7 TeV for colored quirk masses. In addition, it has also proposed using the CMS dataset taken with zero magnetic field. In this dataset, all SM particles are expected to follow a straight trajectory, but the quirks would still bend due to the string tension. The second study [240] has proposed a new algorithm to search for quirks, exploiting the fact that the quirk and anti-quirk pair should lie in the same plane with the highest sensitivity in the $\ell \sim 1\text{--}10$ mm range. This avoids the necessity of fitting non-helicoidal trajectories, and has the potential to extend the sensitivity to quirks well beyond the current mono-jet and HSCP limits. The third study [334] considers quirks that lose energy through material interactions with the detector. A charged or colored quirk pair can come to a stop within the detector, and annihilate out-of-time with active pp collisions, allowing for sensitivity from SP searches across a wide range of characteristic length scales, $\ell \sim \text{Å km}$ range. We refer the reader to the section 3.5.3 for more information on SP searches.

Because of the non-standard nature of the tracks, quirk phenomenology poses substantial challenges in their experimental reconstruction, and the lack of constraints on quirks have already attracted the attention of the ATLAS, CMS and LHCb experiments. It would be desirable to test how the phenomenological proposals in [239, 240, 334], among others, can perform in a realistic detector simulation of one of the LHC experiments.

3.5.3. Out-of-time decays of SPs. This category is unique because LLP decays occur out-of-time with the collision. Indeed, decays can even occur when the LHC is not running! The only member of this class is the search for SPs, which we describe below.

If an HSCP is produced with very low kinetic energy, it can come to rest in the detector due to interactions with the detector materials. This most likely occurs in the calorimeters or the steel yoke in the MS as a result of their high material densities. The HSCP can then decay at a later time when no collision is taking place (known as an out-of-time decay). This experimental restriction reduces the types of background processes affecting the search, with the dominant backgrounds coming from cosmic rays (CRs), beam halo, and instrumental noise.

In the Run 1 analyses from ATLAS [236] and CMS [234], events are selected with a dedicated trigger selecting bunch crossings which are empty and have no bunches of protons nearby. The analyses require a jet with p_T (E) above 30 (50) GeV at ATLAS (CMS). ATLAS further supplements the hardware trigger by requiring $p_T(j) > 50$ GeV, $|\eta| < 1.3$ and $E_T > 50$ GeV, rejecting instrumental noise. In addition, CMS has updated the jet search in Run 2 and also provided a search that triggers on out-of-time muons, both of which use the 13 TeV dataset [235]. The latter also employs the displaced stand-alone (DSA) muon reconstruction algorithm [335].

An offline selection procedure is aimed at reducing the main backgrounds. Muons coming from CRs can be identified due to their distinctive topology. The ‘beam halo’ background is the result of protons interacting with residual gas in the beampipe, the beampipe itself, or collimators upstream from the IP. Most particles will not travel far before being absorbed by various structures, but muons will travel parallel to the beam and can leave calorimeter deposits out of time with a proton–proton collision. However, these deposits will often be accompanied by corresponding horizontal tracks in the MSs and can thus be efficiently vetoed. Instrumental noise is rejected in CMS by exploiting the anomalous response in the HCAL.

Stopped particle searches provide an alternative way of probing charged particles besides more conventional HSCP searches. HSCP searches are typically more sensitive to any signature with a charged particle, so SP searches are not often expected to be a discovery mode for most simple new physics scenarios with charged LLPs¹⁷⁷. The typical added value of SP searches is that they can help identify and characterise positive signals in HSCPs, for instance by providing a cleaner extraction of the lifetime and also to properly identify the decay products. However, in some cases (such as for quirks, where HSCP searches are insensitive in much of parameter space), it has been argued that an SP search could actually be a discovery mode if modifications are made to the search strategy to improve sensitivity to quirks [334]. These modifications that would increase quirk acceptance or lower backgrounds include expanding the η range, implementing higher energy thresholds, using the timing information, and considering shower evolution in the new CMS endcap calorimeter [336].

3.6. Discovery opportunities: overview of gaps in coverage

In the preceding sections 3.1–3.5.3, we have examined the so-called ‘coverage’ of existing searches for LLPs at the LHC with the explicit and express purpose of identifying uncovered realms and places where discoveries could be hiding. Here, we summarize these gaps and potential opportunities for LLP discovery in bullet form, as a to-do list for the experimental community.

1. All-hadronic LLP decays

- Associated-object triggers (especially motivated by Higgs-like VBF and VH production modes) need to be more comprehensively used to improve sensitivity to low- p_T objects.
- Improvements are needed in sensitivity at lower masses and lifetimes (e.g. for LLPs produced in Higgs decays).
- Single hadronic DVs need to be looked for in searches that currently use two (such as decays in ATLAS HCAL and MS).
- Possibilities need to be explored for ATLAS and CMS for online reconstruction of hadronic displaced objects, as the inclusive H_T triggers used by the two collaborations miss these objects unless they have a large p_T . (By contrast, LHCb can trigger on a displaced hadronic vertex [253, 254].)
- Low-mass hadronically decaying LLPs can look somewhat like tau leptons, so the question remains as to whether there is any possibility of using, for example, L1 tau triggers to seed displaced jet triggers at HLT and improve trigger efficiency; studies need to be performed by the experimental collaborations.

¹⁷⁷ The reason why the SP searches are less efficient than the HSCP ones is twofold. On one hand, only a fraction of LLPs stop in the detector, while the HSCP search only requires that the LLP crosses the detector. On the other hand, the SP is only looked for in a specific time-window that might fail to catch a large fraction of them.

- The prospects for dedicated searches for displaced hadronic taus need to be investigated, since no dedicated searches currently exist.
- The potential for flavor-tagging displaced jets (b-displaced jets, c-displaced jets, etc) needs to be explored.

2. Leptonic

- Coverage needs to be provided for the intermediate region between boosted, low-mass LLPs (lepton jets) and high-mass, resolved LLPs (resolved ATLAS/CMS searches).
- Improvements need to be made to extend coverage to lower masses and to lower p_T thresholds. Currently no prescription or plan for this exists, and so dedicated studies need to be done.
- Searches need to be done for different combinations of charge and flavor of displaced leptons (e.g. same-sign versus opposite-sign, opposite-flavor versus same-flavor).
- Searches need to be done for tau leptons in LLP decays, in particular if they come from the ID; an unanswered question remains as to whether displaced-jet triggers can be used for this purpose.

3. Semi-leptonic

- Searches do not exist and need to be done for LLP masses below about 30 GeV; this mass range is theoretically well motivated by Majorana neutrinos.
- Searches need to be performed for all flavor combinations (for example, one CMS search only covers $e^\pm \mu^\mp$), as well as same-sign versus opposite-sign leptons.
- Currently unknown improvements need to be made to relax or modify isolation criteria wherever possible to recover sensitivity to boosted semi-leptonic decays.
- Searches need to be done that better exploit triggering on associated objects for improved sensitivity to low-mass objects, or to employ high-multiplicity lepton triggers if there are multiple LLPs.

4. Photonic

- There is currently no coverage for LLPs decaying into $l\gamma$, $j\gamma$, or without E_T , and searches urgently need to be performed for this decay topology.
- There is currently poor coverage (i.e. there exists no dedicated search) for single- γ topologies. The only searches with sensitivity require two jets to be present in addition to E_T [303]. Studies are needed to assess the sensitivity of this search to signals with only one delayed photon and different jet multiplicities.
- There is currently no coverage for softer non-pointing or delayed photons, and searches need to be performed for these kinematic realms.
- Studies need to be performed to determine if triggers on associated objects may improve sensitivity to signals with a single photon, without E_T , or for lower- p_T photons

5. Other exotic long-lived signatures

- Disappearing tracks with $c\tau \sim \text{mm}$ are very hard to probe, and new ideas and detector components are needed to extend sensitivity to this potential discovery regime. It is unclear if the ATLAS insertable B-layer will be present in HL-LHC run and how sensitivity to the DT topology will improve with the replacement of the current ID with the new inner tracker (ITk), or whether new tracking layers very close to the beam line can be added. It is an open question as to what is the lowest distance at which new layers (or double layers) can be inserted. Another open question that needs to be answered is whether there are any prospects for DTs at LHCb with an upgraded detector.

- No dedicated searches for quirks exist at the LHC, a huge, open discovery possibility for ATLAS, CMS, and LHCb. Some LHC constraints exist by reinterpreting HSCP searches, but dedicated searches need to be performed. There are significant challenges in modeling the propagation and interaction of quirks with the detector, as well as in fitting tracks to their trajectories, but new ideas have been proposed that need to be explored by the experimental collaborations that might allow improved sensitivity to quirks with less ambitious analysis methods.

4. Common sources of backgrounds for LLP searches

Chapter editors: Juliette Alimena, Martino Borsato, Zhen Liu, Sascha Mehlhase

4.1. Introduction

For many searches for LLPs, the main backgrounds do not stem from irreducible SM processes, but arise instead from *external sources*. Indeed, there can even be backgrounds of instrumental and/or algorithmic nature. Often, LLP searches are designed to have a very small number of background events, sometimes even zero events, that pass the full selection criteria. This chapter gives an overview of common LLP search backgrounds and the means to estimate or control them.

4.2. Long-lived particles in the SM

Weak decays of SM particles can naturally give rise to DV at the boosts typically encountered at the LHC. Searches for LLP signatures at sufficiently low LLP mass and lifetime suffer from large backgrounds due to displaced SM decays. One simple example is found in the search for long-lived dark photons decaying to $\mu^+\mu^-$ at LHCb [269], which drastically loses sensitivity when the dark photon mass gets too close to the $K_S \rightarrow \pi^+\pi^-$ invariant mass, despite the very low $\pi \rightarrow \mu$ misidentification rate.

Moreover, b -hadrons can decay at displacements of a few mm and can be challenging to distinguish from LLPs with masses of a few tens of GeV that decay to a pair of jets. Requiring a large track multiplicity of the DV and performing a mass fit to the dijet invariant mass can help to significantly reduce the effect of this background (see, for example, [251, 253]). Backgrounds from heavy flavors are typically more abundant in the forward region, as arises, for example, if the signature under study is an LLP from the decay of a SM-like Higgs boson. However, the LHCb forward detector, which was designed to study these SM decays, is, in most cases, capable of rejecting heavy flavor backgrounds more effectively than can be done in ATLAS or CMS. Furthermore, displaced tracks from b -mesons, which usually have impact parameters (d_0) of less than 2 mm, can be rejected by using a larger criterion for the minimum track d_0 .

4.3. Real particles produced via interactions with the detector

Particles produced in the pp collision can interact with nuclei of the detector material, giving rise to DV, and can mimic LLP signals. Vertices from these interactions will be positioned in regions of the detector containing high densities of detector material and are therefore effectively vetoed by using detailed material maps.

The LHC detectors have developed tools internal to the collaborations to define a material volume to be vetoed. As the detector configurations changed slightly from Run 1 to Run




Figure 12. An example of a material map from ATLAS [237].




Figure 13. An example of a material map from CMS [339].

2, material maps have been determined separately for each data-taking period for both the ATLAS and CMS collaborations, using collision data. Maps can be found for ATLAS for Run 1 in [337], CMS for Run 1 in [338], ATLAS for Run 2 in [237], and CMS for Run 2 in [339]. Additionally, the Run 2 maps for both are shown here in figures 12 and 13 for ATLAS and CMS, respectively.

LHCb recently developed a precise material map of the VELO using beam-gas collisions [340], shown in figure 14. Beam-gas collisions can be distinguished from long-lived heavy flavor backgrounds and their utilization allows the map to cover precisely the whole VELO




Figure 14. Reconstructed secondary vertices in the LHCb VELO from beam-gas collisions in the zr plane integrated over ϕ . These vertices are used to build the material map [340] to veto backgrounds from material interactions.

geometry, not only the region close to the IP. This map was used to veto photon conversions to di-muons, which is the main background affecting displaced dark-photon searches at low mass [269]. In analyses, this material map, together with properties of a reconstructed secondary vertex and its constituent tracks, is used to construct a p -value that is assigned to the hypothesis that the secondary vertex originates from a material interaction. As a rule of thumb, LHCb material interaction background is dominant for vertices at a distance from the beam axis larger than 6 mm (where the VELO material begins). Below 6 mm the background is dominated by heavy flavor decays [341].

Because accurate material maps are essential to performing fully reliable sensitivity studies to signatures with DV, making them publicly available to the broader LLP community is of the highest priority. The availability of such tools in fast parametric simulations such as Delphes [342] would be very useful to reinterpret LLP search results. In section 6.4.6.2, an example of a reinterpretation of an LLP search is presented, where a rough material veto was performed because these material maps were not publicly available, highlighting the shortcomings of such approaches in the absence of accurate material maps and emphasizing the benefits of making them available.

High-energy collision muons originating mainly from W decays and creating secondary interactions in the tracker, calorimeters or MS can be an additional source for DV mimicking LLP signals. This mostly minor background arises as vetoing these DV is based on a not 100% efficient detection of the high- p_T track in the tracker.

Another important background, mainly for analyses targeting the reconstruction of decay vertices of LLPs reaching the MSs, is hadronic or electromagnetic showers not contained in the calorimeter volume, so-called punch-through jets [248]. These punch-through jets occur especially in regions of reduced total interaction length in the calorimeters (e.g. transition regions between the barrel and the end-caps) and can be suppressed by either rejecting these $|\eta|$ regions or requiring a minimal number of hits in the MS and isolating the DV from calorimeter jets as well as high-energy tracks and significant track activity in the inner tracking system. In order to not reject true vertices from displaced decays that occur near the end of the calorimeters, the calorimeter-jets veto should only consider jets with a minimum total energy deposit and, e.g. a minimum electromagnetic fraction of the total energy. The track isolation requirement aims at regions with a poor calorimeter measurement (again, transition regions in the calorimeters), where a single high-energy track or the sum of the track activity in a small cone around the DV could indicate a (punch-through) jet. On the other hand, punch-through jets, given their similarity to signal signatures, can also be used to evaluate systematic uncertainties due to imperfect modeling in the muon-system simulation.

4.4. Real particles originating from outside the detector

There are several types of real particles generated outside the detector that could be sources of background in an LLP search.

4.4.1. Cosmic muons. Cosmic rays from the atmosphere can enter the detector as cosmic-ray muons. These cosmic-ray muons can be reconstructed as displaced muons in the MS or as displaced jets in the calorimeters. If cosmic-ray muons are reconstructed in the MSs, they will typically appear as two back-to-back muons with ϕ values near $\pm\pi/2$. The rate of cosmic muons in the detector is about 500 Hz at L1, but depending on the HLT path and the offline selection used, the rate of cosmic-ray muons entering a given LLP analysis is generally much less.

Cosmic-ray muons are typically an important background source to consider for displaced signatures, especially those with large displacements [234, 236, 343–345]. Cosmic-ray muons are generally only an issue for LLP analyses in CMS and ATLAS since LHCb has coverage only in the forward direction.

For many analyses, cosmic-ray muons can be rejected with a simple veto on back-to-back dimuons. However, in some LLP analyses, this veto is not optimal for the signal acceptance or it is insufficient to suppress cosmic-ray backgrounds. Another often-used way to minimize the contribution from cosmic-ray muons is requiring high-momentum muons and/or high-energy jets, since cosmic-ray muons have a rapidly falling p_T spectrum. In addition, if a search primarily looks for inner-tracker or calorimeter objects, cosmic-ray-muon events can be rejected by requiring little MS activity [234, 343, 344].

If the cosmic-ray muon background is significant for an analysis, it can be estimated using data from dedicated cosmic data-taking runs or from empty bunches in pp collision runs [234, 343, 344]. Cosmic ray muon simulations can be made, but in many LLP analyses, a data-driven approach is favoured if the simulation modeling is found to be insufficient. Timing information in the calorimeters or the MSs can be used to discriminate the signal from cosmic-ray muons, sometimes in conjunction with impact parameter variables.

4.4.2. Beam halo. Another type of real particle generated outside the detector that could be a significant source of background for LLP searches is beam halo. Beam halo is produced when protons from the LHC beam scatter off the LHC collimators and produce debris, which can appear in the detector. Beam halo can create energy deposits in the calorimeters or hits in the MS, both of which would be largely in the beam direction. These energy deposits or MS hits would appear earlier than if they had been made from particles coming from the collision (see figure 15). Beam halo is usually not modelled in MC simulation, since it is highly dependent on the beam parameters.

Beam halo is most relevant for searches for displaced signatures without tracks in the inner tracker and for searches for decays in non-collision bunches (e.g. from SPs) [234, 343, 344], which are described in section 3.5.

The contribution from beam halo can often be reduced by requiring high-momentum or high-energy reconstructed objects. One can also decrease the number of beam halo events by requiring central objects or vetoing forward MS activity, since beam halo is usually in the very forward direction [234, 343, 344]. For inner tracker-based signatures, events from beam halo are rejected by requiring a minimum number of early hits; in this way beam halo is rejected due to its anomalous timing.




Figure 15. A sketch illustrating the timing differences due to the shorter, more direct path to the calorimeter cells between a beam-halo muon and a particle originating from the collision. The beam-halo muon is detected earlier than the particle from the collision.

Beam halo background can be estimated using data control regions near $\phi = 0$ and π . One could also identify cells with a low number of (or zero) tracks that are assigned an early time.

4.4.3. Cavern radiation. Diffuse backgrounds can also arise from proton–proton collisions filling the LHC caverns, consisting mostly of neutral, low-energy, and long-lived SM particles (i.e. neutrons and photons), leading to an overall increase in occupancy, especially in the MSs. This so-called ‘cavern background’ or ‘cavern radiation’ can constitute a significant background in a LLP search. As simulations are resource-intensive, it is usually not at all modelled in MC simulation samples.

Cavern radiation is most relevant for searches looking in non-collision data, that is, stopped-particle searches, and for searches using MS information to form tracks and vertices. It can be estimated from data by collecting events triggered by random triggers when there are no collisions, as was done in [236]¹⁷⁸. Cavern radiation can also be estimated by overlaying a cavern radiation simulation with minimum-bias events from data.

4.5. *Fake-particle signatures*

Another type of background for LLP searches is that from signatures that mimic real particles in the detector, but are in fact fake. Fake particles can originate from spurious detector noise. Noise appears differently for each detector, but in general, it is characterized by a single and concentrated energy deposit or hit that does not correspond in time or space to any other energy deposits or hits in the detector. Noise is usually difficult to model with MC simulation.

Calorimeter detector noise is most relevant for searches looking in non-collision bunches and low-energy collisions [234, 343, 344]. Muon system noise is most relevant for searches that are also highly affected by cosmic-ray muons.

Calorimeter noise can be rejected by vetoing single and concentrated energy deposits [234, 343, 344]. Muon system noise can be rejected by requiring high-quality muon tracks.

Noise in both the calorimeters and the MSs could be estimated by looking at dedicated cosmic data-taking runs and then applying some selection criteria to reject cosmic-ray muons. The remaining events would most likely be noise.

¹⁷⁸ Note that these triggers are unlike those used to collect the search data for stopped-particle searches, which instead select events with physics objects during empty bunch crossings.




Figure 16. Illustration of two close-by, low-mass vertices being reconstructed as one high-mass vertex.

4.6. Algorithmically induced fakes

For searches that aim to reconstruct the decay vertex of an LLP, and especially for LLPs decaying in the proximity of the interaction region, algorithmically induced fakes and/or instrumental backgrounds can be of importance. Algorithmic fakes can still be a significant background to LLP searches, even if a given detector is noise-free.

4.6.1. Random/Merged vertices. This type of background, illustrated in figure 16, is especially important in the environment close to the interaction region that experiences a high track density, and arises from two main sources. First, two or more individual tracks can cross each other and can be reconstructed as a DV. Second, two close-by, low-mass vertices can be reconstructed as one high-mass DV; such a final merging/cleaning step is often part of vertexing algorithms to reduce fakes in standard vertexing.

The former source is mostly suppressed by requirements originally targeting the removal of meta-stable SM particles: a minimum transverse impact parameter, $|d_0|$, for tracks and a minimum distance between the PV and a given DV.

The latter source is harder to suppress, though can be estimated by randomly merging vertices from distinct events. By studying the number of reconstructed ‘merged’ high-mass vertices as a function of distance between the two low-multiplicity low-mass vertices that were ‘merged’—both with vertices from the same event, as well as from different events and scaling them accordingly—an estimate for this background can be derived. This method has been successfully used in the ATLAS search for DV [237] and the ATLAS multitrack analysis [196].

4.6.2. Randomly crossing tracks. A background that is typically more relevant than merged vertices is the background stemming from low-mass DV crossed by unrelated tracks, resulting in the reconstruction of a high-mass vertex, as illustrated in figure 17. The mass of the reconstructed DV is especially increased when the random track crosses the vertex in a direction that is perpendicular to the distance vector pointing from the PV to the displaced one.

As demonstrated in detail in [196, 237], this background can be estimated by constructing vertices (n -track) from lower-multiplicity ones ($n - 1$ -track) by adding pseudo-tracks, drawn randomly from data-driven track templates derived for various radial detector regions. The normalization of the prediction is performed by comparing the $n - 1$ -track-based constructed vertices with the actual n -track vertices in all radial detector regions. One




Figure 17. Illustration of a low-mass vertex crossed by an unrelated track and being reconstructed as a high-mass vertex instead.

potential method for suppressing such backgrounds is to veto vertices where removing one track substantially decreases the mass of tracks associated with the vertex.

4.7. Summary

LLP searches often have very low backgrounds, as opposed to searches for prompt particles. This makes LLP searches highly sensitive to signals of new physics.

There are, however, a few common sources of background that arise in different LLP searches: other, known, LLPs such as b -hadrons; real particles produced in the detector, such as particles produced in collisions that interact with the nuclei of the detector material; real particles produced outside the detector, such as cosmic muons or beam halo; fake particles, such as detector noise; and algorithmically induced fakes, such as two tracks that cross and are reconstructed as a DV, as described above. These backgrounds are generally atypical, difficult to model in simulation, and challenging to estimate. Thus, the possible appearance of unexpected background sources should be taken into account in any new LLP search, and the development of novel techniques and methods to estimate them is encouraged.

5. Detector upgrades

Chapter editors: Juliette Alimena, Martino Borsato, Yangyang Cheng, Monica Verducci

Contributors: Cristiano Alpigiani, Xabier Cid Vidal, David Curtin, Elena Dall'Occo, Sven Dildick, Jonathan L Feng, Iftah Galon, Christopher Hill, Henning Keller, Felix Kling, Simon Knapen, Zhen Liu, Henry Lubatti, Philippe Mermod, Vasiliki A Mitsou, James L Pinfold, Jessica Prisciandaro, Dean Robinson, Livia Soffi, Sebastian Trojanowski, Carlos Vázquez Sierra, Si Xie, Charlie Young

The experimental searches for LLPs outlined in chapter 3 are limited by the abilities of the ATLAS, CMS, and LHCb experiments to trigger on and reconstruct the objects that are associated with each signature. In the Phase 2 high-luminosity upgrade of the LHC (HL-LHC), the extremely high PU conditions necessitate the upgrade of all three detectors to maintain triggers at thresholds needed for sensitivity to electroweak, Higgs, and BSM physics (see [346, 347] for an overview); to reject particles originating from PU vertices; and to maintain object reconstruction in the high-luminosity environment. These upgrades include the addition of tracking layers to the forward regions of ATLAS and CMS, improvements to

timing reconstruction in events, and the inclusion of tracking information at earlier stages of the trigger (or, in the case of LHCb, the full online reconstruction of every event).

While these upgrades are crucial to the success of the HL-LHC physics goals for conventional searches (such as for electroweak, Higgs, or SUSY signatures), they have the opportunity to be transformative for searches for LLPs. Signatures involving LLPs are often subject to small-to-negligible irreducible backgrounds, and improvements to the reconstruction, timing, and vertexing of displaced objects can typically suppress any instrumental or fake backgrounds. (See chapter 4 for a discussion of sources of backgrounds for LLP searches.) At the same time, the introduction of tracking information to earlier stages of the trigger could be used to trigger on events that may contain hadronically decaying LLPs that must otherwise pass jet triggers, leading to an improvement in sensitivity to low-mass LLPs. Indeed, many of the gaps in coverage from current searches identified in section 3.6 can be closed or reduced using the new technology from detector upgrades. Even more uniquely, LLP signatures may themselves motivate the introduction of new detector elements that are dedicated to exploring new lifetime frontiers in particle physics.

This chapter summarizes current and proposed plans for detector upgrades for the HL-LHC, paying special attention to those features of the detector upgrades that are most relevant for LLP searches. Where available, we show the results of projections for the sensitivity to various LLP scenarios of different improvements to the detector. We also highlight LLP studies for the Phase 2 upgrades that are not yet publicly available that should be done in order to assess (and, where possible, improve) the sensitivity of planned upgrades to LLP signatures. Finally, we include contributions from a number of existing and proposed experimental collaborations whose primary purposes are to search for LLPs produced at LHC IPs using additional, dedicated detectors. These detectors complement the capabilities of ATLAS, CMS, and LHCb, allowing sensitivity to signatures that are otherwise not possible to reconstruct at the main detectors.

In section 5.1, we combine the discussions of the planned ATLAS and CMS upgrades, facilitating for each detector component a direct comparison between the features of the upgraded detectors of both ATLAS and CMS. Since LHCb has a very different geometry and physics program from ATLAS and CMS, we have a separate discussion of planned LHCb upgrades in section 5.2. Finally, we present the contributions of the dedicated LLP experiments in section 5.3.

5.1. The ATLAS and CMS experiments

The planned upgrades to the ATLAS and CMS experiments for the HL-LHC will give the detectors increased coverage in the forward regions, better spatial and timing resolutions, and other new features including track triggers. The improved hardware capabilities, combined with software developments, give rise to exciting new prospects for future LLP searches. This section gives an overview of the upgrade scope (section 5.1.1), discusses their physics potential (sections 5.1.2 and 5.1.3), and presents new ideas for detector upgrade and LLP searches (section 5.1.4).

Unless specified otherwise, the subsequent CMS experimental results are from its Technical Design Reports for the different sub-detector upgrades at HL-LHC, namely tracker [348], barrel calorimeter [349], endcap calorimeter [336], muon detectors [350], timing detector (Technical Proposal, [351]), and Level-1 Trigger (Interim Technical Design Report, [352]). ATLAS results are from the Technical Design Reports for the inner tracker pixel detector [353], the TDAQ system [354], the tile calorimeter [355], LAr calorimeter [356], MS [357], and inner tracker strip detector [358].

5.1.1. Detector and trigger upgrades for high-luminosity LHC. The high luminosity LHC (HL-LHC) will begin with the third long shutdown (LS3) of the LHC in the coming decade (as of this writing estimated to begin at the end of 2023), where the machine and detectors will be upgraded to allow for pp running at a luminosity of $5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ in the nominal scenario, or potentially $7.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ in the ultimate performance scenario. This will allow the ATLAS and CMS experiments to collect integrated luminosities ten times that of the current operations, which amounts to around 300 fb^{-1} per year and 3000 fb^{-1} during the projected HL-LHC lifetime of ten years (up to 4000 fb^{-1} if the ultimate instantaneous luminosity can be achieved).

The HL-LHC conditions create unique challenges in terms of high PU levels and high radiation dosage. About 140 PU events per bunch crossing, on average, are expected in the nominal scenario, and up to 200 PU events in the ultimate luminosity scenario. The radiation levels will be unprecedented: for the design integrated luminosity of 3000 fb^{-1} , a 1 MeV neutron equivalent fluence of $2.3 \times 10^{16} n_{\text{eq}} \text{ cm}^2$ and a total ionizing dose of 12 MGy (1.2 Grad) is expected at the centre of the detectors, where the innermost silicon pixel tracking layers will be installed.

To meet the challenges of the HL-LHC operating conditions, and to fully profit from its physics capabilities, comprehensive upgrade programmes are planned for both the ATLAS and CMS experiments. This section summarizes the main detector and trigger upgrade plans for each sub-detector component of both experiments.

5.1.1.1. Tracker. By the start of the HL-LHC, the inner trackers of both experiments must be replaced due to the significant radiation damage and performance degradation they have suffered. To maintain tracking performance in the high-density environment, and to cope with the increase of approximately a factor of ten in the integrated radiation dose, both the ATLAS and CMS experiments will entirely replace their inner tracking detectors.

CMS upgrade. The CMS tracker is composed of the inner pixel detector and the outer tracker. At the HL-LHC, the CMS inner pixel detector will include four cylindrical barrel layers covering the region of $|z| < 200 \text{ mm}$, and forward extensions of up to twelve endcap disks on both sides (compared to the current configuration with three disks), which will extend its $|\eta|$ coverage from the current value of 2.4 to approximately 4. To maintain radiation hardness and reasonable occupancy, as well as to improve resolution, small, thin pixels will be used. For the studies in the CMS tracker TDR [348], pixels with a thickness of $150 \mu\text{m}$ and $25 \times 100 \mu\text{m}^2$ in size are used in the simulation¹⁷⁹. The first layer of the barrel inner pixel detector will be positioned at a radius of 28 mm.

The CMS outer tracker is composed of six cylindrical barrel layers in the central region, covering the region of $|z| < 1200 \text{ mm}$, complemented on each side by five endcap double-disks, in the region of $1200 < |z| < 2700 \text{ mm}$. Modules are installed between $r \sim 21 \text{ cm}$ and $r \sim 112 \text{ cm}$. Three sub-detectors are distinguished: the Tracker Barrel with pixel-strip modules (TBPS), the Tracker Barrel with strip-strip modules, and the tracker endcap double-disks (TEDD). The inner rings of the TEDD disks use pixel-strip (PS) modules up to $r \sim 60 \text{ cm}$, and the rest use strip-strip modules. The outer tracker modules, called p_T modules, are composed of two single-sided, closely-spaced (1–4 mm separation) small pitch sensors read out by a set of front-end ASICs that correlate the signals in the two sensors and select the hit pairs (referred to as ‘stubs’) compatible with particles above the chosen p_T

¹⁷⁹ An alternative option being considered is that of $50 \times 50 \mu\text{m}^2$. Larger pixel sizes of $50 \times 200 \mu\text{m}^2$ or $100 \times 100 \mu\text{m}^2$ are being considered in outer barrel layers and outer rings of the endcap as a potential option to reduce power consumption.

threshold. A p_T threshold of 2 GeV corresponds to a data volume reduction of roughly one order of magnitude, which is sufficient to enable transmission of the stubs at 40 MHz. The ‘stubs’ are used as input to the hardware trigger at Level-1 (L1), which enables track-finding at L1 for all tracks with a p_T of 2 GeV or above. To improve the ‘stub’-finding efficiency and also to reduce material, the inner three outer tracker barrel layers, the TBPS, are made with flat modules in the center and tilted modules in the regions with larger z .

ATLAS upgrade The ATLAS collaboration will replace its ID with a new, all-silicon tracker to maintain tracking performance in HL-LHC conditions.

The new ATLAS ITk will consist of a greatly enlarged pixel system extending to roughly twice the radius and four times the length of the current pixel array, coupled with a much more finely-segmented strip detector. In total, the coverage of the full radius of the inner solenoid requires over three times the silicon area of the current detector.

The new system will consist of silicon barrel layers and disks (strips) or rings (pixels) with the possibility of inclined pixel modules to better cover the transition from the barrel to the end-cap regions. In detail, the strip detector has four barrel layers and six end-cap petal-design disks, both having double modules each with a small stereo angle. The strip detector, covering $|\eta| < 2.7$, is complemented by a five-layer pixel detector extending the coverage to $|\eta| < 4$. The combined strip-plus-pixel detectors provide a total of 13 hits for $|\eta| < 2.6$, with the exception of the barrel/end-cap transition of the strip detector, where the hit count is 11.

5.1.1.2. Calorimetry. Both the ATLAS and CMS calorimetry consist of ECALs and HCALs. Different materials and designs are used for the two experiments.

CMS upgrade. For the CMS detector, the existing scintillating crystals in its ECAL (ECAL) barrel (EB) will be kept for the duration of LHC. On the other hand, both front-end and back-end electronics will be replaced [349], which allows for higher transfer rates and more precise timing. The target timing resolution for the upgraded ECAL electronics is ~ 30 ps for particles with $p_T \gtrsim 30$ GeV, which is the fundamental limit allowed by hardware and an order of magnitude smaller than the current limit. Current studies on the CMS HCAL barrel radiation damage suggest there is no need for replacement at HL-LHC.

The CMS endcap calorimeter, including both the electromagnetic (EE) and the hadronic sections, will be replaced with a high-granularity, silicon-based calorimeter (HGCAL). The HGCAL, with fine granularity in both the lateral and longitudinal directions, enables 3D imaging in reconstructing energy clusters. The intrinsic high-precision timing capabilities of the silicon sensors will add an extra dimension to event reconstruction. The HGCAL is expected to provide a timing resolution of ~ 10 s of ps for high-energy particles with p_T of tens of GeV.

ATLAS upgrade. The ATLAS liquid argon (LAr) calorimeter will be improved in Phase 2 with an electronics upgrade that will provide optimized super cells and full-granularity data to the trigger system by means of a new pre-processor. A similar upgrade of the ATLAS Tile calorimeter readout will use on-detector digitization and a new back end pre-processor. Both the LAr and Tile calorimeters expect to implement a 40 MHz readout system for Phase 2. The transmission of high-granularity calorimeter data (all cells with a transverse energy of two times the noise threshold) drives the bandwidth requirement for the upgraded trigger and data acquisition (TDAQ) system.

In addition, the outermost Tile calorimeter layer can be used to identify muons in the range $|\eta| < 1.3$ by better identifying particle energy depositions above the minimum ionizing particle (MIP) threshold.

5.1.1.3. Muon system. The MS will be upgraded at both experiments to meet HL-LHC conditions, extend geometric coverage, and improve detector performance and trigger capabilities.

CMS upgrade. For the CMS detector, its current MS consists of three different types of muon detectors. In the barrel region, drift tubes (DTs) are installed along with resistive plate chambers (RPCs). In the endcaps, there are cathode strip chambers (CSCs) together with RPCs. At the HL-LHC, the existing muon detectors will be improved with upgraded electronics to enable 40 MHz readout, as well as improve the timing resolution from the current 12.5–1 ns [350]. New muon detectors, namely gas electron multipliers and a new version of RPCs, will be added to the endcaps, covering the regions of $1.6 < |\eta| < 2.4$. Additional muon chambers, labeled ME0, will cover the very forward regions of $2.4 < |\eta| < 2.8$, a region also covered by the upgraded inner tracker. This will be used for the muon trigger at L1. The additional muon detectors are essential to achieve a high trigger efficiency with an acceptable rate, especially in the forward region. The additional hits in the new endcap muon stations, combined with improved algorithms, permit efficient triggering on displaced muon tracks even in the harsh environment of the HL-LHC.

ATLAS upgrade. Most of the front-end and detector readout of the ATLAS MS, including the trigger electronics for the RPC, thin gap chambers (TGCs), and monitored drift tube (MDT) chambers, will be replaced to face the higher trigger rates and longer latencies necessary for the new Level-0 (L0) trigger required by the HL-LHC conditions. The MDT chambers will be integrated into the L0 trigger in order to sharpen the momentum threshold. Some of the MDT chambers in the inner barrel layer will be replaced with new, small-diameter MDTs. Additional RPC chambers will be installed in the inner barrel layer to increase the acceptance and robustness of the trigger, and some chambers in high-rate regions will be refurbished. New TGC triplet chambers in the barrel–endcap transition region will replace the current TGC doublets to suppress the high trigger rate from random coincidences in this region. The electronics of all the sub-detectors will need to be replaced due to obsolescence, aging, and radiation damage. During the Phase 1 upgrade (to take place from 2019 to 2021) the new small wheel chambers will replace the CSC and the MDT chambers of the innermost endcap wheel by micromegas and small-strip TGCs. The replacement of the MDT front-end readout will address the trigger rate and latency requirements of the TDAQ system in Phase 2 and allow the use of MDT hit information to improve the muon p_T resolution in the L0 trigger. Additionally, in the upgraded detector all data from the barrel and endcap detectors will be transmitted to FPGAs at L0, which can be used to implement more advanced and flexible algorithms for muon reconstruction, including the use of neural networks and/or dedicated tracking for non-pointing muons [354].

Some LLP scenarios (e.g. HV models [66]) predict muon-jet final states, which result in collimated muons that are not identified with high efficiency by the current triggers (see section 3.2). In figure 18, the opening angle between muons in a HV model is shown, for a particular combination of particle masses and parameters. The di-muon separation is much smaller than the current system can resolve (approximately 0.2 in $\Delta\phi(\mu, \mu)$). In the no-upgrade scenario, these can only be recorded by the single muon trigger. In the upgraded scenario, a dedicated trigger is under development for a dimuon trigger with a p_T threshold of ≈ 10 GeV.

5.1.1.4. Timing detector. Precision timing can be provided by the aforementioned calorimetry upgrades. However, the tens of ps timing resolution in the upgraded calorimeters is only achievable for particles with energy above tens of GeV. Moreover, timing information for delayed objects from calorimetry alone will be affected by the




Figure 18. Opening angle between muons in a Hidden Valley model, where a sub-GeV-mass particle decays to $\mu^+\mu^-$. The opening angle is well below the resolution of the current system. From [359].

beamspot smearing, which corresponds to about 180 ps of uncertainty. Therefore, global event timing with the ability to reconstruct the vertex time and exploit time information in charged particle reconstruction requires a dedicated fast timing detector.

CMS upgrade. For the CMS experiment, the proposed MIP timing detector (MTD) will comprise a barrel and an endcap region made up of a single layer device between the tracker and calorimeters, and cover $|\eta|$ up to ~ 3 . In the barrel, the proposal is to adapt the present Tracker Support Tube (TST) design by instrumenting the current location of the thermal screen with a thin, actively-cooled, stand-alone (SA) detector, based on lutetium–yttrium orthosilicate crystals activated with cerium (LYSO:Ce) and read-out with silicon photomultipliers. The endcap region can be instrumented with a hermetic, single layer of MIP-sensitive silicon devices with high timing resolution, with a pseudorapidity acceptance from about $|\eta| = 1.6$ to $|\eta| = 2.9$. The MTD is designed to provide timing resolution of a few tens of ps for charged tracks throughout the detector lifetime. The performance projection in section 5.1.3 is evaluated with a 30 ps resolution for a p_T threshold of 0.7 GeV in the barrel and a p threshold of 0.7 GeV in the endcap, and covering the expected MTD fiducial region of $|\eta| < 3$.

ATLAS upgrade. The high-granularity timing detector (HGTD) is intended to distinguish between collisions occurring very close in space but well-separated in time. Currently there is not yet a TDR for this project. The current proposed detector design is based on low-gain avalanche detector technology that will cover the $|\eta|$ region between 2.4 and 4, with a timing resolution of 30 ps for MIPs. High-precision timing will improve the track-to-vertex association in the forward region, impacting jet and lepton reconstruction, as well as offering unique capabilities for online and offline luminosity determination.

5.1.1.5. Trigger. The ATLAS and CMS experiments adopt a two-level trigger system: the hardware-based Level-1 trigger (L1) and the software-based high-level trigger (HLT).

CMS upgrade. For the CMS experiment, the L1 trigger currently only uses calorimeter and muon information. At the HL-LHC, with the aforementioned outer tracker upgrade of p_T modules and stub-finding capabilities, tracking information will be included at L1 [352]. The

L1 track trigger uses parallel processing and pattern recognition on stub information to achieve track finding at an output rate of 750 kHz.

The L1 tracking capability will be further complemented by the calorimeter and muon upgrades, which provide more precise position and momentum resolution, calorimeter shower shape, and more muon hits in the forward region. In the L1 trigger, the electron and photon trigger algorithms for HL-LHC will use information from the ECAL as well as from the outer tracking detectors. The algorithm should preserve the ability to reconstruct electromagnetic clusters with p_T above a few GeV with high efficiency (95% or greater above 10 GeV) as well as achieve high spatial resolution which should be as close as possible to the offline reconstruction. Following the upgrade of both on-detector and off-detector electronics for the barrel calorimeters at the HL-LHC, the EB will provide energy measurements with a granularity of (0.0174, 0.0174) in (η, ϕ) , as opposed to the current input to the L1 trigger consisting of trigger towers with a granularity of (0.087, 0.087). The much finer granularity and resulting improvement in position resolution of the electromagnetic trigger algorithms is critical in improving electron/photon trigger efficiency and suppressing background at high PU.

The L1 global trigger (GT) will be upgraded with more sophisticated and effective global trigger calculations based on topology, plus an additional intermediate correlator trigger to fully exploit the increased information in the trigger objects, such as matching tracking info with finely-grained calorimeter information, or a combination of muon and track information. The upgraded detector readout and DAQ systems will allow 12.5 μ s latency and a L1 rate of 750 KHz; the latter may be substantially reduced by adding L1 tracking information matched to improved L1 trigger objects from the calorimeters and MS. At the HLT, the processing power is expected to scale up by PU and L1 rate, with an output rate of 7.5 kHz and up to 10 kHz.

ATLAS Upgrade. The ATLAS trigger and the data acquisition system are being planned with the intention of fully exploiting the physics potential of the HL-LHC. A baseline architecture has been proposed and documented in [354].

The hardware-based Level-0 Trigger system will be composed of separate calorimeter and muon triggers, as well as a GT and Central Trigger sub-systems. The result of the Level-0 trigger decision is transmitted to the data acquisition system at 1 MHz, and is followed by an upgraded Event Filter system to achieve a maximum stored event rate of 10 kHz.

The upgraded trigger system will take advantage of increased granularity provided by the calorimeters, will improve efficiency for muon-based triggers and perform hardware-based tracking profiting from the extended coverage of the planned silicon ITk. Options exist to further develop a hardware-based track trigger for quicker and less CPU-intensive rejection of the expected large increase in PU at the HL-LHC, to take full advantage of extended coverage provided by the ITk. Such a hardware track trigger (HTT) would be an evolution of the current ATLAS FTK reconstruction system.

5.1.2. Object performance: tracking and vertexing. The ability of the detectors to reconstruct tracks and find vertices with high precision and efficiency in a high-density environment underlies the experimental reach for displaced objects. This section reviews the ATLAS and CMS experiments' projected tracking performance at the HL-LHC, highlighting improvements and new features with the upgrades.

5.1.2.1. CMS performance. L1 tracking. With the aforementioned tracker and L1 track trigger upgrades, the CMS experiment will be able to do track finding at L1 as well as offline at HL-LHC. Both L1 and offline tracking performance are discussed here.




Figure 19. Left: L1 tracking efficiency versus generated particle p_T for $|\eta| < 2.4$. Right: L1 tracking efficiency versus η for $p_T > 3$ GeV. Results for muons (electrons) are shown as filled black (open red) circles, and are produced with $t\bar{t}$ events in a scenario with 200 pile-up events per bunch crossing, on average [348].




Figure 20. Relative p_T (left) and z_0 resolution versus pseudorapidity for muons in $t\bar{t}$ events with zero (black dots), 140 (red triangles), and 200 (blue squares) pile-up events per bunch crossing, on average. Results are shown for scenarios in which truncation effects are (markers) or are not (lines) considered in the emulation of L1 track processing. The resolutions correspond to intervals in the track parameter distributions that encompass 68% (filled markers and solid lines) or 90% (open markers and dashed lines) of all tracks with $p_T > 3$ GeV [348].

All L1 tracking studies have been performed assuming 3 GeV stub p_T thresholds. In figure 19, the L1 tracking efficiency for prompt muons and electrons for $t\bar{t}$ events in a scenario with 200 PU interactions per bunch crossing, on average, is presented. The tracking efficiency for muons exhibits a sharp turn-on at the 3 GeV stub p_T threshold, and saturates at approximately 98%. The tracking efficiency for electrons turns on more slowly and flattens out at 90%, mostly due to interaction with the detector material.

In figure 20, the L1 tracking resolutions of the p_T and z_0 parameters of muons with $p_T > 10$ GeV in $t\bar{t}$ events is shown for various average PU scenarios. The resolutions are defined in terms of an interval centered on the residual distribution that contains 68% or 90% of the tracks. Loss in tracking efficiency due to truncation effects (where there is insufficient time to transfer all the stub data) is determined from hardware and emulation to be at the level of 10^{-3} when considering $t\bar{t}$ samples with a PU rate of 200. As expected, resolutions degrade at forward pseudorapidity due to a corresponding increase in multiple scattering. In general,




Figure 21. Relative resolution of the transverse momentum (left) and transverse impact parameter (right) as a function of the pseudorapidity for the current (black dots) and the upgraded (red triangles) CMS tracker, using single isolated muons with a transverse momentum of 10 GeV [348].

L1 parameter resolutions are excellent, which will provide for robust trigger object matching and charged particle reconstruction in the L1 trigger.

Offline tracking. Preliminary results on the offline tracking performance over the full acceptance of the CMS tracker are excellent, with further improvements expected as the detector design and simulation algorithms are optimized. In figure 21, the resolution of the transverse momentum and the transverse impact parameter for single muons with $p_T = 10$ GeV as a function of the pseudorapidity, both with the current detector and after the implementation of the HL-LHC upgrades, is shown. The better hit resolution of the HL-LHC tracker and the reduction of the material budget result in a significantly improved p_T resolution. The transverse impact parameter resolution is also improved with respect to the current detector, ranging from below 10 μm in the central region to about 20 μm at the edge of the acceptance.

For $t\bar{t}$ events, the efficiency to identify the PV correctly is $\sim 95\%$ at an average PU level of 140, and $\sim 93\%$ at an average PU level of 200. The vertex algorithm used is the same as the one used in Run 2 for a PU of about 35, therefore it is not yet optimized for vertex reconstruction at very high PU. In figure 22 the resolution of the vertex position in the x , y , and z coordinates is shown as a function of the number of tracks associated to the vertex. The vertex position resolution is almost independent of the amount of PU in the event and the longitudinal resolution is only 50% worse than the transverse one, as expected given the pixel dimensions of the inner tracker modules.

Given that the CMS HLT tracking is based on the offline tracking code, a similar level of performance is expected. Because of HLT time constraints, a parallelization of the algorithms is already under development and will be applied also in the HLT track reconstruction at HL-LHC.

5.1.2.2. ATLAS performance. Excellent tracking performance is also expected with the ITk upgrade of the ATLAS experiment for the HL-LHC era. The left panel of figure 23 shows the track reconstruction efficiency for jets in $Z' \rightarrow t\bar{t}$ events with 200 PU for different η ranges. The right panel of figure 23 shows the fake rate for reconstructed tracks in $t\bar{t}$ events, and there




Figure 22. Vertex position resolution in x and y (left) and z (right) as a function of the number of tracks associated to the vertex, for $t\bar{t}$ events with an average of 140 (full circles) and 200 (open circles) pile-up interactions per bunch crossing [348].




Figure 23. Left: track reconstruction efficiency for tracks in jets from $Z' \rightarrow t\bar{t}$ with 200 average pile-up events. The efficiency is shown as a function of jet p_T for different η ranges, and $M_{Z'} = 5$ TeV. Right: fake rate for tracks in $t\bar{t}$ events with 200 average pile-up events using ITk; Run 2 detector results are shown for comparison. Both figures are from [353].

is clearly substantial improvement over the Run 2 detector performance along with the improved coverage in the forward region.

In figure 24, the resolution of the transverse momentum and the longitudinal impact parameter for single muons with various p_T values is shown as a function of the pseudorapidity both with the current detector and projections for after the HL-LHC upgrade using digital clustering to find the tracks. The improvement is even more marked with analogue clustering: the transverse and longitudinal impact parameter resolutions are shown for different pixel pitches in figure 25.

5.1.3. Upgrade projection: LLP searches. Searches for LLPs are well motivated by various classes of extensions of the SM, discussed at length in chapter 2.2. Often, the production cross section for such processes is expected to be very small. The HL-LHC will allow for the




Figure 24. Resolution of the transverse momentum (left) and longitudinal impact parameter (right) as a function of the pseudorapidity for the current (solid line) and the upgraded (points) ATLAS tracker [353].




Figure 25. Track parameter resolutions using analogue clustering for (left) transverse impact parameter; (right) longitudinal impact parameter. The resolutions are shown for single muons with $p_T = 100 \text{ GeV}$. The results of ITk are shown for $25 \times 100 \mu\text{m}^2$ and $50 \times 50 \mu\text{m}^2$ pixels, along with the current Run 2 detector performance [353].

collection of much larger data sets needed to reach better sensitivity to such BSM scenarios. The prospects are further strengthened with detector and trigger upgrades. This section discusses these potential improvements, and presents sensitivity projections on a number of benchmark LLP search channels with the aforementioned upgrades at the HL-LHC.

5.1.3.1. Heavy stable charged particles in CMS. A number of new physics scenarios give rise to HSCPs with long lifetimes that move with subrelativistic speed through the detector, heavily ionizing the sensor material as they pass through. In split SUSY models, the




Figure 26. Left: distribution in CMS of the dE/dx discriminator versus track momentum (p) for tracks with high momentum ($p_T > 55$ GeV) in background events (red) and for candidate signal particles. Pair produced $\tilde{\tau}_S$ with a mass of 871 GeV (blue), and a gluino with a mass of 1400 GeV (green), are shown. Right: the performance of the dE/dx discriminator for selecting gluinos in events at rates of 0 pile-up (PU) and 200 PU. The signal versus background efficiency performance curves for a discriminator making use of both the pixel information and the outer tracker HIP flag (red and magenta) demonstrate a better performance compared to a discriminator trained to exploit only the dE/dx information from the pixel modules (blue and green), for a background rejection of 10^{-6} [348].

supersymmetric particles known as the stau ($\tilde{\tau}$) and the gluino (\tilde{g}) can have such characteristic signatures. The relevant simplified models are described in sections 2.4.2–2.4.3, and current searches are described in section 3.5.1.1.

Sensitivity projection with tracker upgrade. Depending on the mass and charge of the new particles, HSCPs experience anomalously high energy losses through ionization (dE/dx) in the silicon sensors with respect to the typical energy losses of SM particles, as can be seen in figure 26 (left). At the CMS experiment, the current strip tracker features analog readout. Furthermore, the pixel detector featured analog readout during Phase 0 in 2016 and before, and currently features digital readout during Phase 1, which started at the beginning of the 2017 LHC run. Therefore, these detectors allow for excellent dE/dx measurements.

At the HL-LHC, the upgraded CMS inner pixel detector will continue providing dE/dx measurements, enabled by its time-over-threshold readout, while the outer tracker cannot provide such information, given that the readout is binary. To increase the sensitivity for signatures with anomalously high ionization loss, a second, programmable, threshold has been implemented in the short strip ASICs of the PS modules of the outer tracker, and a dedicated readout bit signals if a hit is above this second threshold.

Searches for HSCPs can thus be performed by measuring the energy loss in the inner pixel detector and by discriminating HSCPs from minimum ionizing particles based on the ‘HIP flag’ in the outer tracker. The threshold of the minimum ionization needed to set the HIP flag is an adjustable parameter in each PS module independently. A threshold corresponding to the charge per unit length of 1.4 MIPs, resulting from preliminary optimization studies, is used in the simulation, and the gain in sensitivity obtained by using the HIP flag is studied.

An estimator of the degree of compatibility of the track with the MIP hypothesis is defined to separate candidate HSCPs from tracks from SM background sources. The high resolution dE/dx measurements provided by the inner pixel modules are used for the computation of the dE/dx discriminator. The tracks in background events have a low number of HT clusters with HIP flag, compared to those observed for tracks in HSCP signal events and slow-moving protons and kaons in minimum bias events.

In figure 26 (right), the performance of the discriminator is shown by evaluating the signal versus background efficiency curves to identify tracks from signal events and reject those originating from backgrounds. The performance curves are evaluated for two different strategies for the discriminator: the dE/dx discriminator, which relies solely on the inner pixel modules (dE/dx -only, ignoring the HIP flags), and a recomputed discriminator which includes the HIP flags from the outer tracker PS modules ($dE/dx +$ HIP flags). The signal versus background efficiency performance curves demonstrate that for a background efficiency of 10^{-6} , analogous to the current analysis performance, the $dE/dx +$ HIP-based discriminator leads to an expected signal efficiency of 40%, around 4–8 times better than the dE/dx -only discriminator. In the dE/dx -only scenario, the efficiency for the HSCP signal is about 8 times smaller than that obtained in current data. The inclusion of the HIP flag for the PS modules of the Outer Tracker restores much of the efficiency, so that the same sensitivity as in Phase 1 will be realized with about four times the luminosity of Phase 1. The Phase 1 sensitivity will be surpassed with the full expected integrated luminosity of the HL-LHC. This study demonstrates the critical impact of the HIP flag in restoring the sensitivity of the CMS tracker for searches for HIPs in the HL-LHC era.

Additionally, the current CMS inner pixel detector provides measurements of charge deposits in each pixel up to 9600 electrons over a range of 4 bits in the digitizer. While it may be difficult to increase the number of bits used to store the charge information due to data rate constraints, it is possible to adopt a dual-slope mapping from charge deposit to ADC counts in the digitizer, which will preserve the granularity for lower charge deposits, while giving more information for HIPs such as HSCPs. This option is currently being studied to evaluate the potential improvement to dE/dx measurements. Furthermore, tuning of the HIP flag threshold may bring additional improvements.

HSCP trigger with muon detector upgrade. The upgrade of the RPC system will allow the trigger and identification of slowly moving particles by measuring their TOF to each RPC station with a resolution of $\mathcal{O}(1)$ ns. The new RPC detectors have a two-end strip readout, which provides precise measurements of the hit position in the local y or the global η coordinate. The speed of muon-like particles and the time (bunch crossing) of their origin will be computed with a fast algorithm to be implemented in the Level-1 trigger at the HL-LHC.

The RPC detectors are synchronized to register muons moving at the speed of light with a local time equal to zero with respect to the collision event that produced the trigger. Slow-moving particles, as HSCPs, will arrive with a delay depending on their speed as shown in figure 27. This time delay, measured by each RPC layer crossed by the HSCP, is exploited in order to trigger on and reconstruct such particles.

The principles of the proposed HSCP trigger algorithm are illustrated in figure 28. In this figure, the vertical axis is the time of signals measured in RPC chambers, as synchronized so that muons moving nearly at the speed of light from a particular collision are measured at the time of the collision. The horizontal axis is the distance from the collision point to the position of the RPC at which the time is measured. The diagram shows three successive bunch crossings, two of which contain muons represented at horizontal lines. The diagram also shows the RPC time measurements from two HSCPs having slopes different from zero due to their traveling significantly slower than the speed of light. The time delay Δt is related to the




Figure 27. RPC hit time measurement distribution in CMS for muons from $Z \rightarrow \mu\mu$ events and for semi-stable $\tilde{\tau}$ particles with $m \approx 1600$ GeV, produced in $pp \rightarrow \tilde{\tau}\tilde{\tau}$ processes [350].




Figure 28. Diagram showing times measured at different RPC stations for particles originating at different bunch crossings and with different velocities in CMS. The x -axis represents the distances from IP to RPC detectors, while the y -axis corresponds to time. The clock at each RPC station is tuned so that particles moving with the speed of light are registered with the exact same ‘local’ times. Hence, relativistic particles are represented by horizontal lines on this diagram [350].




Figure 29. Left: resolution of a particle-speed measurement at L1 trigger with Phase 1 and upgraded RPC Link Board System. Right: the efficiency as a function of β of the standard L1 muon trigger without any p_T threshold, and the RPC-HSCP Phase 2 trigger with 1.56 ns sampling time. For both plots, an HSCP signal is shown [350].

speed v of an HSCP via the following equation:

$$\Delta t = d \left(\frac{1}{v} - \frac{1}{c} \right). \quad (5.1)$$

Here d is the distance between the IP and the point where an HSCP crosses an RPC. For RE4/1 chambers and $\beta = v/c = 0.2$, the delay time is > 6 bunch crossings, comparable to 150 ns.

A penetrating charged particle leaves a trail of hits in RPC chambers along its trajectory. The TOF can be computed in each RPC station with respect to a number of bunch crossing hypotheses. Should there be a common velocity solution, derived from equation (5.1), with $\beta < 0.6$, a trigger is formed. For $\beta > 0.6$, the delays are small and can be handled by the Phase 1 trigger. The performance of this algorithm has been studied with the CMS full simulation. All the detector effects (e.g. electronics jitter, signal time propagation along strips) are considered. A particle-speed measurement resolution is shown in figure 29 (right) for the case of 25 ns signal sampling time (Phase 1) and 1.56 ns sampling time provided with the upgraded RPC Link Board System. For both plots, an HSCP signal is shown.

The efficiency of the RPC-HSCP algorithm as a function of β is studied and compared with the standard L1 muon trigger. The results are shown in figure 29 (right). The current CMS-HSCP Phase 1 trigger performs well down to $\beta \approx 0.75$. The upgraded RPC Link Board System will allow for the triggering, at the correct bunch crossing, on possible HSCPs with velocities as low as $\beta \sim 0.25$.

Possible improvements for this trigger proposal in the β measurement could be achieved by matching tracks in the track trigger to the HSCP muon trigger.

5.1.3.2. Displaced muons in CMS. Many BSM theories predict particle decays with displaced muon or muon pairs in its final state, such as dark SUSY and GMSB with smuons. In order to demonstrate the physics potential of displaced muons at the HL-LHC with the CMS detector, a particular SUSY model is selected where the displaced signature consists of a dimuon final state emerging from the decay of heavy sparticles (smuons). Searches for the direct production of heavy sparticles with long lifetimes are difficult in the present LHC runs, owing to small cross sections and limited integrated luminosity, and will only become possible at the HL-LHC.




Figure 30. Left: the muon transverse impact parameter, $|d_0|$, for several simulated smuon decay lengths, $c\tau$, at the generator level. Right: distribution of the minimum number of valid hits in the CMS Phase 2 muon system for a SUSY $\tilde{\mu}$ with $m_{\tilde{\mu}} = 500$ GeV and $\tau = 1000$ mm for the Run 2 (blue) and Phase 2 (red) detectors [350].

In GMSB models, smuons can be (co-)NLSPs (next to lightest supersymmetric particles) and decay to a muon and a gravitino [360]. When the slepton is long lived, the final state signature is two displaced oppositely charged muons and significant missing transverse energy. The smuon pair production has the advantage that it can be characterized by a very clean final state topology, and we will therefore focus on the process $q\bar{q} \rightarrow \tilde{\mu}\tilde{\mu}$, where the two smuons decay far from the primary interaction vertex. For this process, the muon $|d_0|$ can reach up to approximately one meter (or longer) for sufficiently large lifetimes, as shown in figure 30 (left). Figure 30 (right) compares the number of hits created by these displaced muons in the CMS MS in Phase 2 and the current CMS detector. The hits plotted here are those associated with the DSA muon tracks, which is a muon track reconstruction algorithm specifically designed for displaced muons that can only be reconstructed in the MS [335].

Standard triggers and reconstruction algorithms that use the position of the PV will not efficiently reconstruct tracks with large impact parameters. Consequently, triggering on and reconstructing muons produced far from the IP is challenging and requires dedicated triggers and reconstruction algorithms. The upgrades to the MS in CMS, as well as the L1 tracking capabilities, significantly improve the experiment's ability to search for displaced muons at the HL-LHC.

Triggering on displaced muons. The momentum resolution of the L1 muon trigger for muons coming from the PV will be greatly improved by adding information from the L1 track trigger, discussed previously. The L1 track trigger can also be directly combined with trigger primitives at the first stage of the muon track-finder electronics; this would mirror the offline reconstruction of ‘Tracker Muons’ which improve the efficiency for very low- p_T muons, especially in the barrel region.

To trigger on both prompt and non-prompt muons effectively at L1, a SA L1 muon generates two p_T measurements for each muon, prompt and non-prompt, which are matched with L1 tracks. If the track match is successful, the L1 track trigger p_T is used and a prompt




Figure 31. L1 Muon trigger rate (left) and efficiency (right) versus muon p_T threshold for the barrel displaced muon algorithm [350].

candidate is formed. If the match is unsuccessful and the muon is not vetoed by L1 tracks, the non-prompt L1 muon p_T is used to form a displaced muon candidate. Figure 31 shows good performance for displaced muons with this method, i.e. there is a reasonably high efficiency and a trigger rate for a single muon trigger of around 10 kHz under HL-LHC conditions. Further improvements to the algorithm are underway to accommodate high PU conditions. The upgrade of the RPC system will allow slowly-moving particles to pass the trigger and be identified by measuring their TOF to each RPC station with a resolution of $\mathcal{O}(1)$ ns. The speed of muon-like particles and the time (bunch crossing) of their origin will be computed with a fast algorithm to be implemented in the L1 trigger for the HL-LHC.

Reconstruction. A dedicated muon reconstruction algorithm was designed for non-prompt muons that leave hits only in the MS. This DSA algorithm is seeded by groups of track segments in the muon chambers. For each seed, a muon track is reconstructed with the same Kalman-filter technique as for the standard SA muon reconstruction algorithm, but without constraining the IP. Figure 30 (right) shows the distribution of the number of hits in the Run 2 and HL-LHC detectors for displaced muons. The impact of the new stations is clearly visible. The charge mis-identification probability is expected to further decrease with the additional hits.

Sensitivity projection with a GMSB model. To study the impact on physics sensitivity, a particular GMSB model is selected where the displaced signature consists of a dimuon final state plus gravitinos, emerging from the decay of heavy long-lived sparticles (smuons), where the gravitinos escape detection. This maps to the direct pair production simplified model with neutral LLP decays to muon + invisible in section 2.4. This signal can serve as a proxy for other models with two LLPs decaying into muons. The final-state signature is then given by two displaced, oppositely-charged muons and significant E_T . Example LLPs with $c\tau = 10, 100, 1000$ mm and several mass hypotheses (0.2, 0.5, 1 TeV) are simulated.

The main background for this search comes from multi-jet production (QCD), $t\bar{t}$ production, and $Z/DY \rightarrow ll$ events where large impact parameters are (mis)reconstructed. Cosmic-ray muons have been studied in Run 2 and these studies can be directly applied to Phase 2 running. In the barrel, they are efficiently rejected by the timing of the hits in the upper leg. Cosmic-ray muons do not originate at the vertex and therefore pass the upper-barrel sectors in reverse direction from outside in. The fraction of cosmic-ray muons in the endcaps is negligible.

Given the very low cross section of the signal process, it is essential to reduce the background efficiently. The best background discriminator is the impact parameter




Figure 32. The 95% C.L. projected upper limits at CMS for $q\bar{q} \rightarrow \tilde{\mu}\tilde{\mu}$, $\tilde{\mu} \rightarrow \mu\tilde{G}$ for various mass hypotheses for $c\tau = 1$ m (left), and as a function of the decay length for $m_{\tilde{\mu}} = 200$ GeV (right). In both panels, the theoretical cross section for the specific model is represented by the blue solid line. For different SUSY breaking scales, $\tan\beta$ or otherwise modified parameters, the cross sections may be 100 times larger, reflected by the blue dashed-dotted line. Green (yellow) shaded bands show the one (two) sigma range of variation of the expected 95% C.L. limits. Phase 2 results with an average of 200 pile-up collisions per bunch crossing and an integrated luminosity of 300 fb^{-1} are compared to results obtained with 300 fb^{-1} . The black line shows the sensitivity without the DSA algorithm, which reduces the reconstruction efficiency by a factor three [350].

significance $d_0/\sigma(d_0) \geq 5$. Given the signal kinematics, the muons from a signal process are expected to move in roughly opposite directions and \not{E}_T can be expected to be larger than 50 GeV. After this selection the signal efficiency is about 4%–5% for $c\tau \approx 1000$ mm, nearly independent of the smuon mass, and 10^{-5} – 10^{-4} for QCD, $t\bar{t}$, and DY backgrounds.

In figure 32, the expected exclusion limits are shown for the GMSB model in which the smuon is a (co)-next-to-lightest supersymmetric particle (NLSP, where ‘LSP’ indicates the lightest supersymmetric particle), for the predicted cross section as well as for a factor 100 larger cross section. The exclusion limits are shown as functions of smuon mass in figure 32 (left) and decay length in figure 32 (right).

The sensitivity depends on $c\tau$ because shorter decay lengths shift the signal closer to background. In figure 32 (right), the resulting physics sensitivities in terms of production cross section for the HL-LHC, normalized to 3000 fb^{-1} , are shown for the dedicated reconstruction of displaced muons and for the standard reconstruction. Also shown is the expected sensitivity at the end of Phase 1. Systematic uncertainties for the Phase 1 scenario are taken from current Run 2 analyses and adapted for expected HL-LHC conditions based on the assumptions of reduced systematics described in [361]. Clearly, only the HL-LHC will allow this process to be studied. The expected exclusion limit is around 200 GeV for $c\tau = 1000$ mm with 3000 fb^{-1} . This also illustrates the importance of keeping lepton trigger thresholds at several tens of GeV, even in the environment of 200 PU interactions per bunch crossing. Similarly, the discovery sensitivity is assessed assuming that a signal is present in data, and is shown as a function of smuon decay length in figure 33.

Sensitivity projection with a dark SUSY model. The analysis presented above was reinterpreted using a Dark SUSY model [30, 155], in which an additional dark $U_D(1)$




Figure 33. The projected discovery sensitivity at CMS for $q\bar{q} \rightarrow \tilde{\mu}\tilde{\mu}$, $\tilde{\mu} \rightarrow \mu\tilde{G}$ as a function of the decay length for $m_{\tilde{\mu}} = 200$ GeV. Together with the discovery sensitivity the corresponding p -value is shown. Phase 2 results with an average of 200 pile-up interactions per bunch crossing and an integrated luminosity of 3000 fb^{-1} are compared to results obtained with 300 fb^{-1} . The black line shows the sensitivity without the DSA algorithm, which reduces the reconstruction efficiency by a factor three [362].

symmetry is added as a supersymmetric SM extension. Breaking this symmetry gives rise to an additional massive boson, the so-called dark photon (γ_D), which couples to SM particles via a small kinetic mixing parameter ϵ . If ϵ is very weak, the lifetime of the dark photon can range from a few millimeters up to several meters. The lower ϵ , the longer is the dark photon lifetime which then decays displaced from the PV. A golden channel for such searches is the decay to displaced muons.

In the model studied here [362], dark photons are produced in cascade decays of the SM Higgs boson that would first decay to a pair of MSSM-like lightest neutralinos (n_1), each of which can decay further to a dark sector neutralino (n_D) and the dark photon.

For the branching fraction $\text{BR}(H \rightarrow 2\gamma_D + X)$, where X denotes the particles produced in the decay of the SM Higgs boson apart from the dark photons, 20% is used. Neutralino masses $m(n_1) = 50$ GeV and $m(n_D) = 1$ GeV are assumed. Final states with two and four muons are included in the analysis. In the former case, one dark photon decays to a pair of muons while the other dark photon decays to some other fermions (2-muon final state). In the latter case, both dark photons decay to muon pairs (4-muon final state).

The main background for this search comes from multi-jet production (QCD), $t\bar{t}$ production, and $Z/\text{DY} \rightarrow \ell\ell$ events where large impact parameters are (mis) reconstructed. Cosmic ray muons can travel through the detector far away from the PV and mimic the signature of displaced muons. However, thanks to their striking detector signature, muons from CRs can be suppressed by rejecting back-to-back kinematics.




Figure 34. Left: distance of the closest approach of the displaced muon track with maximum p_T to the primary interaction vertex, $R_{\text{Muon-1}}$, for signal and background after the final event selection. Right: parameter scan in the $\epsilon - m_{\gamma_D}$ plane. The grey lines indicate the regions of narrow hadronic resonances where the analysis does not claim any sensitivity [362].

For each event, at least two DSA muons are required. If more than two, the ones with the highest p_T are chosen. The two muons must have opposite charge ($q_{\mu,1} \cdot q_{\mu,2} = -1$) and must be separated by $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2} > 0.05$. The three-dimensional angle between the two displaced muons is required to be less than $\pi - 0.05$ (not back-to-back) in order to suppress CR backgrounds. Additionally, $\cancel{E}_T \geq 50 \text{ GeV}$ is imposed to account for the dark neutralinos escaping the detector without leaving any signal.

In order to discriminate between background and signal, the three-dimensional distance from the PV to the point of closest approach of the extrapolated displaced muon track, called R_{Muon} , is used. The event yield after full event selection of both selected muons as a function of $R_{\text{Muon-1}}$ and $R_{\text{Muon-2}}$ is used to search for the signal. The left panel of figure 34 shows $R_{\text{Muon-1}}$ of the first selected muon for signal and background samples.

The search is performed using a simple counting experiment approach. In presence of the expected signal, significance of the corresponding event excess over the expected background is assessed using the likelihood method. In order to evaluate the discovery sensitivity the same input is used as in the limit calculation, now with the assumption that one would have such a signal in the data. The discovery sensitivity is shown in the two-dimensional $m_{\gamma_D} - \epsilon$ plane in the right panel of figure 34. This search is sensitive to large decay length of the dark photon.

In absence of a signal, upper limits at 95% C.L. are obtained on a signal event yield with respect to the one expected for the considered model. A Bayesian method with a uniform prior for the signal event rate is used and the nuisance parameters associated with the systematic uncertainties are modeled with log-normal distributions. The resulting limits for the Dark SUSY models are depicted in figure 35. While the results shown in the left panel of figure 35 are for a dark photon with a decay length of 1 m as a function of the dark photon mass, the right panel of figure 35 shows the results for a dark photon mass of 20 GeV as a function of the decay length [362].




Figure 35. Upper limits at 95% C.L. on production cross section $\sigma/\sigma_{\text{theory}}$ for various dark photon mass hypotheses and a fixed decay length of $c\tau = 1000$ mm (left) and a fixed mass of $M_{\gamma_0} = 20$ GeV (right). Green and yellow shaded bands show the one and two sigma range of variation of the expected 95% C.L. limits, respectively. The grey lines indicate the regions of narrow hadronic resonances where the analysis does not claim any sensitivity [362].

5.1.3.3. Displaced photons at CMS. A number of new physics scenarios predict new particles that, upon decay, result in displaced photons in the final state (see sections 2.4 and 3.4). At the CMS experiment, with the scintillating crystal design of the ECAL that provides excellent resolution but lacks pointing capability, the photon arrival time in the ECAL is the main observable used to distinguish signal from background in displaced photon searches.

One benchmark model for displaced photon searches is the GMSB model where the lightest neutralino ($\tilde{\chi}_0^1$) is the NLSP, can be long-lived, and decays to a photon and a gravitino (\tilde{G}), which is the LSP, as illustrated in figure 36 (left). For a long-lived neutralino, the photon from the $\tilde{\chi}_0^1 \rightarrow \tilde{G} + \gamma$ decay is produced at the $\tilde{\chi}_0^1$ decay vertex, at some distance from the beam line, and reaches the detector at a time later than that of prompt, relativistic particles produced at the IP. The time of arrival of the photon at the detector can be used to discriminate the signal from the background. The aforementioned upgrade to the ECAL electronics in the barrel region, and the HGCAL upgrade in the endcaps, will improve photon timing resolution at HL-LHC by an order of magnitude to as little as ~ 30 ps for photons with p_T of tens of GeV or above, hence significantly improve the experimental reach of displaced photon searches.

Moreover, the proposed MTD will be able to provide another dimension of information to reconstruct LLP decays. The TOF of the photon inside the detector is the sum of the TOF of the neutralino before its decay and the TOF of the photon itself, until it reaches the detector. Since the neutralino is a massive particle the latter is clearly negligible with respect to the former. In order to be sensitive to short neutralino lifetimes of order 1 cm, the performance of the measurement of the photon TOF is a crucial ingredient of the analysis. Therefore, the excellent resolution of the MTD apparatus can be exploited to determine with high accuracy the TOF of the neutralino, and similarly the photon, also in case of a short lifetime.




Figure 36. Left: diagrams for a SUSY process that results in a diphoton final state through gluino production at the LHC. Right: sensitivity to GMSB $\tilde{\chi}_0^1 \rightarrow \tilde{G} + \gamma$ signals expressed in terms of neutralino lifetimes for 300, 180 and 30 ps resolution, corresponding to the current detector, the HL-LHC detector with photon timing without MTD and with MTD, respectively [351].

An analysis has been performed at generator level in order to evaluate the sensitivity power of a search for displaced photons at CMS in the scenario where a 30 ps timing resolution is available from the MTD [351]. The events were generated with Pythia 8 [219], exploring neutralino lifetimes ($c\tau$) in the range 0.1–300 cm. The values of the Λ scale parameter were considered in the range 100–500 TeV, which is relevant for this model to be consistent with the observation of a 125 GeV Higgs boson. After requiring the neutralino to decay within the CMS ECAL acceptance and the photon energy being above a ‘trigger-like’ threshold, the generator-level photon TOF was smeared according to the expected experimental resolutions. A cutoff at a photon time greater than 3σ of the timing resolution is applied and the ‘signal region’ is assumed to be background free to estimate the sensitivity. The signal efficiency of such a requirement is computed and translated, assuming the theoretical cross sections provided in [66], to an upper limit, at 95% C.L., on the production cross section of the $\tilde{\chi}_0^1 \rightarrow \tilde{G} + \gamma$ process.

In figure 36 (right), the analysis sensitivity in terms of the Λ scale (and therefore of the neutralino mass) and lifetime is shown for three different assumptions on the timing resolution. The 300 ps resolution is representative of the TOF resolution consistent with current CMS detector performance. The 180 ps resolution is representative of the TOF resolution of the upgraded CMS detector without the MTD, in which case the TOF measurement will be dominated by the time spread of the luminous region. The vertex timing provided by the MTD detector will bring the TOF resolution to about 30 ps. As visible in the figure, a full-scope upgrade of the CMS detector with photon and track timing will provide a dramatic increase in sensitivity at short lifetimes and high masses, even with the first 300 fb^{-1} of integrated luminosity.

5.1.3.4. Displaced jets in ATLAS. Neutral LLPs that can decay into jets displaced from the proton–proton IP arise in many BSM theories (see chapter 2 for an extensive discussion). For example, in hidden sector models a new set of particles and forces is proposed that is weakly coupled to the SM via a communicator particle. The hidden sector is otherwise invisible to the SM sector, but its particles (some of which can be long-lived) may decay to SM particles via




Figure 37. Simulation projections of the sensitivity of the ATLAS experiment upgrade plans to neutral LLPs. Top: the fraction of the jet energy deposited in the A-layer (left) and BC-layers (right) of the ATLAS tile calorimeter as a function of the transverse decay position of the LLP in events with a 125 GeV Higgs-like boson decaying to two 40 GeV LLPs. Bottom: the same for events with a 600 GeV Higgs-like boson decaying to two 50 GeV LLPs.

the communicator. The lifetimes of these LLPs are typically unconstrained and could be long enough for the LLPs to decay inside the ATLAS detector volume. Such particles produce unique signatures which may have been overlooked by more traditional searches for new physics. One ATLAS search for such displaced jet signatures focuses on LLPs which decay in the ATLAS HCal and consequently deposit most of their energy there and very little or none in the ECAL, and also have few or no charged tracks pointing at the hadronic energy deposits. A signature-driven trigger that optimizes the acceptance for this class of events is required for the online selection.

The existing ATLAS analysis described in section 3.1 can likely be improved by planned upgrades by using HCal information splitting the B-C layers in the calorimeter to identify the LLP decay position. The splitting between B-C layers will provide more information on the longitudinal shower profile, see figure 37.

This analysis uses a hidden sector benchmark model and considers the decay of a heavy boson to two long-lived neutral scalars; this maps to the simplified model with Higgs boson production of hadronically decaying LLPs in section 2.4. The heavy scalar bosons decaying into LLPs have masses ranging from 125 to 1000 GeV, and the LLP scalars have masses ranging between 5 and 400 GeV. Background processes, dominated by QCD dijet production,




Figure 38. Reconstruction efficiency of the disappearing chargino as a function of decay radius using two track reconstruction techniques: ‘tracklets’ refers to the short tracks mentioned in the text, while ‘tracks’ refers to standard track reconstruction. The corresponding reconstruction efficiencies for Run 2 are also shown [353].

are suppressed in this analysis by requiring both scalars to decay in the calorimeter. Final results of the sensitivity projections are pending.

5.1.3.5. Disappearing tracks. In the ATLAS experiment, the planned ITk upgrade of the tracking volume for the HL-LHC can be exploited to improve the existing searches for BSM particles with a DT [353].

Such particles are predicted by many well-motivated supersymmetric models such as anomaly-mediated SUSY-breaking scenarios, where the supersymmetric partners of the SM W bosons, the wino fermions, are the lightest SUSY state. In such models, the lightest neutralino and chargino can be nearly mass-degenerate, with a mass splitting around 160 MeV, and the chargino is consequently long lived. The combination of long lifetime and boost when produced in a high-energy collider allows the chargino to leave multiple hits in the traversed tracking layers before decaying. When performing searches for such signatures in ATLAS, selected events are typically required to contain at least one short track, hereafter called a tracklet. To study how such searches can be improved with the ITk upgrade, some assumptions have been made for simulating events and projected detector response. The detector response is parametrised by using response functions based on studies performed with Geant4 simulations of the upgraded detector in high-luminosity PU conditions.

The tracklet reconstruction efficiency for signal charginos as a function of the decay radius inside the detector is shown in figure 38. A 30% systematic uncertainty on the background yields is assumed, as observed for the Run 2 analysis. The expected background is largely dominated by fake tracklets due to random crossings.

Expected limits at 95% C.L. are shown in figure 39 as a function of the chargino mass and lifetime for a pure wino and a pure Higgsino LSP scenario. For comparison, the current Run 2 limit for 36.1 fb^{-1} in the wino LSP scenario is also shown. The HL-LHC dataset of 3000 fb^{-1} extends the sensitivity of neutrinos and charginos up to 250 GeV, assuming a pure Higgsino scenario.




Figure 39. Projected 95% C.L. limits on a degenerate chargino-neutralino scenario assuming the chargino is a (left) pure wino; (right) pure neutralino. The limits include both pair production $\tilde{\chi}_1^{\pm}\tilde{\chi}_1^{\mp}$ and associated production $\tilde{\chi}_1^{\pm}\tilde{\chi}_1^0$ (the Higgsino model also includes the $\tilde{\chi}_1^{\pm}\tilde{\chi}_2^0$ mode) [353].

5.1.3.6. Displaced vertices in ATLAS. Massive, LLPs with lifetimes of the order of $O(10)$ ps to $O(10)$ ns can decay inside the inner tracker into charged and stable particles. The products of these decays are reconstructed as tracks with measurably distant impact parameters with respect to the IP. The reconstruction of such displaced tracks is very challenging compared to the track reconstruction of prompt particles, due to fewer hits along the track and a larger parameter phase space for track finding. In order to identify a DV, one must first identify the tracks from the decaying LLP.

Several physics models predict the existence of long-lived, massive particles. For example, a standard SUSY scenario can contain a gluino with a mass of 2 TeV and a lifetime of 1 ns. The long-lived gluino hadronises into an R -hadron, which then decays into a 100 GeV neutralino and hadrons.

In ATLAS, this topology has been investigated in Run 2 by using a dedicated algorithm for reconstructing DV [255]. The performance expected to be achieved for the ITk upgrade for this signal has been tested with a simplified simulation which has a description of the ITk active sensors and a modeling of the magnetic field. The kinematics and location of the decay products of the R -hadron are injected into the simulation and their trajectories are extrapolated through the detector model. The probability of producing at least seven silicon hits in the ITk geometry is shown in figure 40 for a 2 TeV R -hadron with $\tau = 1$ ns. The increased volume and number of ITk layers leads to improved efficiency for the reconstruction of tracks from DV (and hence the vertices themselves) out to radial distances of up to 550 mm. The increased number of silicon layers also gives more room to veto tracks with missing hits, further suppressing backgrounds.

5.1.3.7. LLP searches with precision timing at CMS. The CMS MTD will provide new, powerful information in searches for LLPs. In addition to the aforementioned displaced photon search, the additional timing information can be used to provide full kinematic reconstruction of LLP decays, and can be a powerful tool in background suppression.

Possible improvements in the ability to reconstruct LLP mass. A precision MTD allows each reconstructed vertex to be assigned a time and therefore to measure the TOF of LLPs between primary and secondary vertices. Using the measured displacement between primary and secondary vertices in space and time, the velocity of an LLP in the laboratory frame, $\vec{\beta}_{\text{LAB}}^p$ (or, equivalently, the boost γ^p), can be measured. In such scenarios, the LLP can decay to




Figure 40. Probability that tracks with $p_T > 1$ GeV will traverse at least seven silicon ITk layers if they are produced in the decay of an R -hadron with mass 2 TeV and $\tau = 1$ ns. The probability is also shown for at least seven hits with the Run 2 layout [353].

fully visible or partially invisible systems. Using the measured energy and momentum of the visible portion of the decay, one can calculate its energy in the LLP rest frame and reconstruct the mass of the LLP, assuming that the mass of the invisible system is known.

This capability can be demonstrated in scenarios where the LLP decay produces a Z boson, which then decays to an electron–positron pair. For example, in the GMSB model where the $\tilde{\chi}_0^1$ couples to the gravitino \tilde{G} via higher-dimension operators sensitive to the SUSY breaking scale, the $\tilde{\chi}_0^1$ may have a long lifetime, and can be produced in top-squark pair production with $\tilde{t} \rightarrow t + \tilde{\chi}_0^1$, $\tilde{\chi}_0^1 \rightarrow Z + \tilde{G}$, $Z \rightarrow ee$.

Studies with simulated event samples have been performed to estimate the possible improved sensitivity of the search with the CMS MTD upgrade. The events are generated with Pythia 8, and the masses of the top-squark and neutralino are set to 1000 GeV and 700 GeV, respectively. Generator-level quantities are smeared according to the expected experimental resolutions. A position resolution of $12 \mu\text{m}$ in each of the three spatial directions is assumed for the PV. The secondary vertex position for the electron–positron pair is reconstructed assuming $30 \mu\text{m}$ position resolution in the transverse direction. The momentum resolution for electrons is assumed to be 2%. Finally, the time resolution of charged tracks at the DV is assumed to be 30 ps.

The mass of the LLP is reconstructed assuming that the gravitino is massless. The fraction of events with separation between primary and secondary vertices exceeding 3σ in both space and time as a function of the MTD resolution is shown in figure 41 (left). The mass resolution, defined as half of the shortest mass interval that contains 68% of events with 3σ displacement is shown in figure 41 (right), as a function of the MTD resolution.

A similar study has been performed with another SUSY scenario where the two lightest neutralinos and light chargino are Higgsino-like. The light charginos and neutralinos are nearly mass degenerate and may become long-lived as a consequence of the heavy higgsinos. Figure 42 shows the mass resolution as a function of the MTD resolution in this SUSY scenario. In both studies, the additional timing information from the MTD facilitates the reconstruction of the LLP mass, the resolution and efficiency of which are further improved with the excellent timing resolution of the MTD.




Figure 41. Efficiency (left) and mass resolution (right) as a function of the timing resolution of the MTD for reconstruction of the $\tilde{\chi}_1^0$ mass in the SUSY GMSB example of $\tilde{\chi}_0^1 \rightarrow \tilde{G} e^+ e^-$, with mass of $\tilde{\chi}_0^1 = 700$ GeV, considering events with a separation of primary and secondary vertices by more than 3σ in both space and time [351].




Figure 42. Resolution for the reconstruction of the LLP mass in the Higgsino scenario outlined in this section. The LLP mass resolution is shown as a function of lifetime and the MTD resolution [351].

Possible improvement to LLPs searches in general using timing information. Precision timing at CMS will provide a new tool to suppress the background and enhance the reach for LLPs in the HL-LHC era.

A schematic of a typical signal event containing an LLP is shown in figure 43. An LLP, denoted as X , travels a distance ℓ_X into a detector volume and decays into two light SM particles a and b , which then reach a timing layer at a transverse distance L_{T_2} away from the beam axis. In a typical hard collision, the SM particles generally travel very close to the speed of light. Hence, the decay products of X (using particle a as an example) arrives at the timing layer with a time delay of




Figure 43. An event topology with an LLP X decaying to two light SM particles a and b . A timing layer, at a transverse distance L_{T_2} away from the beam axis (horizontal gray dotted line), is placed at the end of the detector volume (shaded region). The trajectory of a potential SM background particle is also shown (blue dashed line). The gray polygon indicates the primary vertex. Taken from [363].

$$\Delta t = \frac{\ell_X}{\beta_X} + \frac{\ell_a}{\beta_a} - \frac{\ell_{\text{SM}}}{\beta_{\text{SM}}}, \quad (5.2)$$

with $\beta_a \simeq \beta_{\text{SM}} \simeq 1$. An ISR jet could easily be present for all processes, and can be used to ‘timestamp’, i.e. to derive the time of the hard collision at the PV.

For the CMS MTD located just outside the tracker volume, $\ell_{\text{SM}}/\beta_{\text{SM}}$ is about $O(1 \text{ ns})$. As a result, with tens of picosecond (ps) timing resolution, a sensitivity to percent-level time delay caused by slow LLP motion, e.g. $1 - \beta_X > 0.01$ with boost factor $\gamma < 7$, is expected to be achieved.

A theory study has been done in [363] examining potential gains in sensitivity for hadronic DV reconstruction using timing information. A new trigger strategy for a delayed jet is studied by comparing the prompt jet with $p_T > 30 \text{ GeV}$ that reconstructs the four-dimensional primary vertex (PV4d) with the arrival time of another jet at the timing layer. The delayed and displaced jet signal, after requiring a minimal decay transverse distance of 0.2 m (L_{T_1}), will typically not have good tracks associated with it. Consequently, the major SM background is from trackless jets. The origins of this background can be classified into two categories: hard collision from a same vertex (SV), and PU from different vertices. Other types of background such as CRs, beam halo, material interactions, etc, are out-of-time and will become important after most of the hard collision background is removed using selections based on reconstructing vertices and timing delay. Out-of-time backgrounds may have other distinctive features that would allow their effective suppression. While the study in [363] can serve as an optimistic inspiration for increasing LLP sensitivity with timing information, detailed experimental studies are needed to determine the actual sensitivity gain.

The jet faking a displaced signal, behaving as a trackless jet, has an intrinsic time delay $\Delta t = 0$. However, due to the limited timing resolution in reconstructing the PV4d, it can have a time spread. The background differential distribution with respect to apparent delay time (Δt) can be estimated as

$$\frac{\partial N_{\text{bkg}}(t)^{\text{SV}}}{\partial \Delta t} = N_{\text{bkg}}^{\text{SV}} \mathcal{P}(\Delta t; \delta_t^{\text{SV}}). \quad (5.3)$$




Figure 44. Theory projection from [363] of the 95% C.L. limit on $\text{BR}(h \rightarrow XX)$, where X is an LLP, for the production of $pp \rightarrow jh$ with the subsequent decay of $h \rightarrow XX$ and $X \rightarrow jj$ subject to assumptions in the text (one ISR jet with $p_{T,j} = 30$ GeV and $|\eta_j| < 2.5$, and at least one LLP inside the detector). Different colors indicate different masses of the particle X . The thick, long-dashed lines indicate searches with the CMS MTD plus the timing requirements. The thick solid and dotted lines indicate searches with a hypothetical timing layer outside the ATLAS muon spectrometer plus timing requirements. The numbers in parentheses are the assumed timing resolutions. This provides motivation to see whether these gains can be realized by studies from within the collaborations.

The time delay selection on Δt reduces such a background through a very small factor of $\mathcal{P}(\Delta t; \delta_t^{\text{SV}})$ for large $\Delta t/\delta_t^{\text{SV}}$, where δ_t^{SV} is the timing resolution for the SV background, dominant by the timing detector resolution. The LLP signal pays a much smaller penalty factor than the background due to its intrinsic delay.

The background from the PU requires the coincidence of a triggered hard event and objects from a PU (hard) collision whose PV4d fails to be reconstructed and that can mimic a signal. The differential background from PU can be estimated as

$$\frac{\partial N_{\text{bkg}}^{\text{PU}}(\Delta t)}{\partial \Delta t} \simeq N_{\text{bkg}}^{\text{PU}} \mathcal{P}(\Delta t; \delta_t^{\text{PU}}). \quad (5.4)$$

Similar to equation (5.3), δ_t^{PU} is the timing resolution for the PU background, dominated by the beam spread. The key difference between the background from the PU and the background from the same hard collision is that the typical time spread is determined by the beam property for the former, and by the timing resolution for the latter. They typically differ by a factor of a few, e.g. 190 versus 30 ps for CMS with the current upgrade plan.

Using this estimation, a sensitivity projection is presented with an example signal of a Higgs boson decaying to LLPs with the subsequent decay of the LLPs into $b\bar{b}$ pairs, with only minimal requirements of one low- p_T ISR jet, with $p_{T,j} > 30$ GeV and $|\eta_j| < 2.5$, and at least one LLP decay inside the detector. Timing information is used to suppress backgrounds. The 95% C.L. sensitivity is shown in figure 44. The decay branching ratio of the LLP $X \rightarrow jj$ is assumed to be 100%. The projection with 30 ps timing resolution of the CMS MTD is plotted with thick dashed lines. Compared to other 13 TeV HL-LHC projections (with 3 ab^{-1} of integrated luminosity) without the timing information (shown in thinner, dotted and dashed

lines), it is suggested that the addition of a selection on timing, under the set of assumptions described above, greatly reduces background and improves sensitivity. The possibility of such a significant improvement is enticing; this projection is determined from theoretical studies [363], however, and it remains to be seen whether it is possible to realize these gains in the experiment. For example, the low threshold of the single jet would require timing information at early levels of the trigger, and out-of-time backgrounds could also reduce the gains from timing information. Nevertheless, this high-level theory analysis provides an inspiration to the experimental collaborations to perform more detailed, internal studies that will ultimately determine how realistic the projections are.

In general, the prospect of improvements in LLP searches at the HL-LHC due to precision timing upgrades for CMS (and ATLAS) remains understudied, and deserves more comprehensive experimental and phenomenological studies, including understanding and reducing out-of-time backgrounds.

5.1.3.8. LLP searches with a level-1 track trigger in CMS. As discussed in section 5.1.1, a central feature of the CMS upgrade at the HL-LHC will be a new silicon outer tracker which allows track reconstruction for every LHC bunch crossing (at a rate of 40 MHz). The p_T selection for stubs (hit pairs in the p_T modules of the outer tracker) to be read out is determined by the bandwidth from the detector to the back end electronics, and is fixed at about 2 GeV. On the other hand, the choice of track finding algorithm and hardware is still being finalized, and there could be significant benefits to extending the L1 track trigger capabilities to trigger on off-pointing tracks.

To illustrate the case, a simple toy simulation to study rare Higgs boson decays into new particles with lifetimes of order of a few mm has been performed [364]. This study considers all-hadronic final states with low H_T , taking SM Higgs boson decays into four jets as an example. Theoretical motivation to look for such decays is very strong; see section 2.2 for a detailed discussion. In particular, there is currently a blind spot for lifetimes of order 1 cm in searches for new long-lived scalars in Higgs decays, i.e. $h \rightarrow \phi\phi$. The goal is to probe very small branching fractions of the Higgs boson, so for this study, $BR[h \rightarrow \phi\phi \rightarrow 4q] = 10^{-5}$ is assumed. For prompt decays, the background is overwhelming, but if the ϕ has $c\tau$ of a few mm, the offline analysis has very low backgrounds. The problem is getting such events on tape, in particular through L1. This toy study estimates how an off-pointing track reconstruction at L1 can help. To estimate the efficacy of the approach, the resulting projections are compared with the best alternatives in the absence of an off-pointing track trigger, by using associated Higgs boson production with a W boson that decays leptonically to pass a lepton trigger, or considering L1 calorimeter jets with no associated prompt tracks.

Once these positive results were obtained, a more detailed study using the full Phase 2 simulation of the CMS detector was performed [365]. The more mature exploration found that a plausible extension of the L1 track trigger to tracks with an impact parameter of a few centimeters results in dramatic gains in the trigger efficiency. The gains are even larger for additional heavy SM-like Higgs bosons with the same decay. These results are in agreement with the toy study described above. A few details of the mature study will be described below.

The study focuses on small or moderate decay lengths of the new particles, 1–50 mm, and assumes that the offline selection can remove all SM backgrounds with only a moderate loss of efficiency. While this study focuses on the specific Higgs boson decay to light scalars, the results and the proposed triggers are relevant for a broad spectrum of new physics searches, with or without macroscopic decay lengths.

The authors propose a simple jet clustering algorithm implementable in firmware, and compare it with anti- k_t jets [366] with a size parameter of $R = 0.3$, as produced by FASTJET

[367]. This simple algorithm produces a similar performance, in both L1 trigger efficiency and rate, to the full jet clustering using the anti- k_t algorithm.

Then, the performance of an algorithm for reconstruction of tracks with non-zero impact parameter is presented. This approach extends the baseline CMS L1 Track Trigger design to handle tracks with non-zero impact parameter and to include the impact parameter in the track fit. This enhanced design is feasible without greatly altering the track finding approach, but will require more FPGA computational power than the current proposal, which only considers only prompt tracks. Tracks passing the selection are clustered using the same algorithm as described above, and clusters containing tracks with high impact parameters are flagged as displaced jets. Though the baseline design of the L1 Track Trigger currently is optimized to find prompt tracks, these studies show that an enhanced L1 Track Trigger can extend the L1 trigger acceptance to include new BSM physics signals.

For now, the extended L1 track reconstruction is limited to the barrel region. In order to compare the results with prompt and extended track reconstruction, one needs to make a correction for the rapidity coverage: prompt tracks are found in $|\eta| < 2.4$, while the extended track algorithm currently only reconstructs tracks in $|\eta| < 1.0$. To scale the efficiency for finding track jets to the full $|\eta| < 2.4$ range, a scale factor based on efficiency in the full η range and the central η range was used. The scale factor was comparable to the increase in L1 rate.

Figure 45 shows the expected trigger rate as a function of efficiency for the SM and the heavy SM-like Higgs bosons.

The available bandwidth for the triggers described above, if implemented, will be decided as a part of the full trigger menu optimization. Here, two cases are considered: 5 and 25 kHz. The expected event yield for triggers using extended and prompt tracking are shown in figure 46, assuming branching fraction $\mathcal{B}[h \rightarrow \phi\phi] = 10^{-5}$ for the SM Higgs boson. For the heavy Higgs boson, the expected number of produced signal events is set to be the same as for the SM Higgs by requiring $\sigma_{pp \rightarrow h(250)} \mathcal{B}[\Phi \rightarrow \phi\phi] = 10^{-5} \sigma_{pp \rightarrow h(125)}$.

For the exotic Higgs decays considered, given the expected total dataset for the HL-LHC of 3 ab^{-1} , and assuming a branching fraction of 10^{-5} , CMS would collect $\mathcal{O}(10)$ events, which should be sufficient for discovery. A plausible extension of the L1 track finder was considered, to select tracks with impact parameters of a few cm. That approach improves the yield by more than an order of magnitude. The gains for the extended L1 track finding are even larger for the events with larger H_T , as demonstrated by the simulations of heavy Higgs boson decays.

5.1.4. Open questions and new ideas. The higher data rate and more powerful machinery in the high-luminosity era bring new prospects to LLP searches. Searches that are too challenging for the current machine may become feasible with the HL-LHC upgrades, and it is important to explore such possibilities to the full extent. Existing search methods, triggering, and reconstruction algorithms should also be updated to take full advantage of the new hardware capabilities. As the upgrade scope is being defined and finalized in the near future, it is of particular importance to evaluate the physics cases, which will also motivate the upgrade and inform its designs.

5.1.4.1. New studies for the HL-LHC. In section 5.1.3, various analyses with displaced signatures are presented in the context of how detector and trigger upgrades at the ATLAS and CMS experiments for the HL-LHC can improve their sensitivity. As was shown in section 5.1.1, while the upgrades for both experiments differ in detail, they are similar in concept and scope. It is therefore important to perform the same LLP search projections for




Figure 45. The rate of the track jet H_T trigger as a function of signal efficiency using extended track finding for the SM Higgs (left) and the heavy SM-like Higgs (right). The extended track finding performance is extrapolated to the full outer tracker acceptance as described in the text and in [365].

both experiments to compare performance and evaluate the complementarity of the two experiments.

For example, both experiments explore the idea of a fast timing detector at HL-LHC. The CMS MTD aims for full-coverage of a barrel layer plus endcap disks, while the ATLAS HGTD plans for multiple layers of endcaps. How the different η coverage of the timing layer might impact LLP searches will be an interesting question to answer. In the case of calorimetry, the ATLAS LAr ECAL is segmented in z , providing the additional pointing capability to find displaced photons. On the other hand, the CMS ECAL electronics upgrade in the barrel, and HGCAL upgrade in the endcaps, will considerably improve its timing resolution to identify photons that are out-of-time. A sensitivity comparison of both experiments will provide helpful information in understanding photon reconstruction and the combined reach of both experiments.

Additional opportunities exist for improving LLP sensitivity in Run 3 or at the HL-LHC, such as the ability to perform an analysis using physics objects at the trigger level, rather than




Figure 46. The number of triggered Higgs events (assuming $B[h \rightarrow \phi\phi] = 10^{-5}$, corresponding to 1700 events) as a function of $c\tau$ for two choices for the trigger rates: 25 kHz (left), 5 kHz (right). Two triggers are compared: one based on prompt track finding (dotted lines) and another that is based on extended track finding with a displaced jet tag (solid lines) [365].

their more complex, offline counterparts. This analysis method—sometimes referred to as *trigger-level analysis* or *data-scouting*—has been used in searches for di-jet resonances using data from Runs 1 and 2 by the CMS Collaboration [368, 369], data from Run 2 by the ATLAS Collaboration [370] and by the LHCb Collaboration for the searches for dark photons [269]. Such data-scouting methods can be used, for example, to reduce the p_T thresholds of jets and muons used in dark photon or hidden particle searches. Such possibilities could potentially assist searches for LLPs by, for example, allowing changes in the L1 muon patterns in order to trigger on non-pointing muons.

Moreover, the applications of machine learning techniques to LLP searches is a currently under-studied realm. Such techniques could be used in the context of LLPs in particle identification, reconstruction, or generation, in addition to analysis or reinterpretation. Further studies to supplement those that already exist are of high priority for the community.

Furthermore, as new models are being proposed and new channels open up in the realm of LLP searches, many with final states challenging for the current detector conditions, it is important to evaluate their sensitivities in the HL era. Here are a few such interesting searches that are on the agenda.

- *Inelastic DM (iDM).* While stringent limits are currently placed on WIMP-type DM from direct detection, indirect detection, and collider searches, DM particles can exist in a hidden sector with additional particles and forces. A representative example of a hidden sector consists of a DM particle which is charged under a hidden gauge or global symmetry. The DM can have both a symmetry-preserving mass and, if the symmetry is spontaneously broken, also a symmetry-violating mass, which splits the mass eigenstates. This iDM scenario [132, 133, 135] consists of two DM states that couple only off-diagonally to one another. Such iDM can be probed at colliders with the production of $DM + DM^*$ (where DM^* denotes the heavier DM state) in association with a hard SM object X , followed by the subsequent decay of $DM^* \rightarrow DM + Y$ for some potentially different SM states Y . The production is summarized as

$$\begin{aligned} pp &\rightarrow X + DM + DM^* \\ &\rightarrow X + DM + (DM^* \rightarrow DM + Y) \equiv X + \cancel{E}_T + Y, \end{aligned}$$

where X is any state that can be used to trigger on the event and reconstruct \cancel{E}_T , such as an ISR jet, and Y depends on the mode by which DM couples inelastically to the SM. One concrete, representative version of such models can be realized when the mediator is a dark photon and Y is a pair of leptons. The weak coupling between the SM and the hidden sector suggest the heavy eigenstate is meta-stable, creating a displaced signature. Because the mass splitting is small between the two eigenstates, the lepton pair is also typically softer compared with GMSB models, with p_T values of a few to tens of GeV for a $O(10\text{ GeV})$ DM. The displaced muon trigger and reconstruction strategy with the MS upgrade at CMS can likely improve searches for this scenario. The soft p_T spectrum in this case particularly motivates the lowering of the p_T threshold in the displaced muon trigger turn-on. Moreover, the additional timing information from the fast-timing detector opens up the possibility of reconstructing the mass splitting, while taking advantage of the good resolution of the timing detector for even low- p_T particles. Sensitivity studies for iDM with a dark photon mediator are planned for the CMS MTD upgrade. The projections will also be of value to searches for other types of dark sector models, such as self-interacting DM [371], that give rise to soft displaced lepton pairs.

- *Dark showers.* If the hidden sector has a QCD-like structure with dark quarks and hidden forces, a mediator between the SM and the hidden sector, such as a Z' or heavy Higgs boson, can decay into some number of dark quarks that subsequently shower and hadronize into dark mesons, some of which are meta-stable and decay back into SM particles after a macroscopic distance from the proton–proton IP. A showering dark sector can yield a particularly rich collider phenomenology that may give rise to a high multiplicity of displaced objects often low in p_T . Depending on the final state, searches for a hadronic shower with emerging [331] or semi-visible [372] jets can be improved with the increased acceptance and enhanced resolution due to the tracker upgrades at both ATLAS and CMS, as well as the finer granularity endcap calorimeter (HGCAL) at CMS. A dark shower with displaced photon pairs can benefit from the improved timing resolution with the ECAL electronics upgrade, as well as the new timing detectors. A dark shower with displaced lepton pairs, similar to the aforementioned case of iDM, can

be probed with better sensitivity with the MS upgrades, as well as the fast timing information. For more details on dark showers, see chapter 7.

5.1.4.2. New detectors at future collider. If we are only limited by our imagination, what new detectors may exist at a future collider that can open up new capabilities for LLP searches? Regardless of practical constraints, studies for such bold new proposals also help us understand our current experimental reach and optimize our search strategies.

- *Four-dimensional tracker with timing.* Extending from the ‘timing layer’ approach for HL-LHC upgrades, it is clear more time measurements on the track are favorable. In particular, the TOF between layers is a powerful discriminant for choosing hits on a track. A future tracker that provides four-dimensional information, including precise timing information for every layer, can significantly improve track reconstruction by reducing combinatorics, providing purer track seeds, and remove fake stubs. Assuming high- p_T particles, each layer’s time measurement is advanced by 30 ps and so preserves the differences in vertex times. Low p_T particles will have even more discrimination power since their paths leave longer times between hits in consecutive layers of the tracker. Major technical advances will need to happen to achieve such an implementation, including development of fine-pitch sensors with good timing resolution, improvements in scaling and power consumption, electronics upgrade, and alterations to present pattern recognition methods in track finding.
- *Timing detector outside ATLAS muon system.* As discussed in section 5.1.3, by using a prompt object, such as an ISR jet, to ‘timestamp’ an event, and requiring a timing delay from the LLP, the new timing layer (MTD) at CMS can help significantly reduce the background and improve LLP search sensitivity. The CMS MTD is located between the tracker and the ECAL with a distance to the beamline of about 1.2 m. If one imagines a fast timing layer outside the ATLAS MS with a distance of approximately 10 m [363], timing information at such distances could provide additional discriminating power for particles with longer lifetimes. In figure 44, the projection with a 30 ps timing resolution of such a hypothetical detector outside the ATLAS MS is plotted in a thick solid line. A less-precise timing resolution (150 ps) has been also considered with a selection $\Delta t > 1$ ns to suppress background. While this study is optimistic in assuming that selections on timing can eliminate all backgrounds to the LLP search, it nevertheless serves as important inspiration for more detailed experimental studies to understand the actual extent of the sensitivity gain. The LLP efficiency is largely unaffected by this change, while low-mass LLPs lose sensitivity by a factor of a few. The CMS MTD timing upgrade for the HL-LHC already provides significant improvement [363]. The timing detector outside the ATLAS MS has the notable benefits of lower background, a larger volume for the LLP to decay and more substantial time delay for the LLP signal due to longer travel distance. Moreover, due to the extended time delay of the LLPs in the volume of the MS, less-precise timing can still achieve similar physics goals. As a result, with the above caveats it can serve as an estimate of the best achievable sensitivity using timing information in LLP searches.
- *New double-sided tracking layer very close to the beamline.* Inspired by the L1 track trigger design at CMS for the HL-LHC, a scenario may be imagined where an additional tracking layer close to the beamline is added, mechanics and radiation hardness permitting. This would allow a track veto close to the IP to be implemented, and extend the LLP sensitivity to even smaller displacements. Moreover, if such a layer can be designed with double-sided modules, similar to the CMS outer tracker approach for the

HL-LHC, hit pairs (stubs) from this layer can be used in a L1 track trigger for displaced tracks or DTs. The feasibility of such an approach would depend on the ability to store and process data at a rate not allowed by current technology. On the other hand, it is helpful to consider more innovative data-scouting methods at trigger level to store and select information, or introduce additional discriminating factors to reduce the rate for particular data streams, for the potential of exploring such signatures.

5.2. *LHCb upgrade*

The LHCb experiment is designed to detect decays of LLPs in the SM, namely bottom and charmed hadrons. As such, it is naturally suited for the search of BSM LLPs in a mass and lifetime range comparable to these hadrons. It is the only LHC experiment to be fully instrumented in the forward region $2 < \eta < 5$, where b - and c -hadrons are abundantly produced and their decay length is enhanced due to the large longitudinal boost. In this region, detector occupancy is extremely high and thus the LHCb experiment has been run at reduced luminosity compared to ATLAS and CMS during Runs 1 and 2. However, an upgrade of the detector is planned to allow running at a luminosity of $2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ in LHC Run 3 (starting in 2020) while maintaining or improving the current physics performance [373]. This is five times larger than the luminosity during Runs 1 and 2. This first upgrade phase (Phase 1a) will entail a novel trigger paradigm where all sub-detectors are fast enough to be read out in real time and the first trigger decisions are done in software. This trigger scheme is flexible and offers a great opportunity for searches of striking signatures like those of BSM LLPs. This upgrade comes earlier than the planned ATLAS and CMS upgrades for the HL-LHC phase which are planned to be installed during LHC long shutdown 3 (by 2025). In this shutdown, LHCb plans to consolidate and modestly enhance the Phase 1a upgrade detector (Phase 1b), while a Phase 2 upgrade to run at an even higher luminosity up to $\sim 2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ is planned to be installed later, during long shutdown 4 (by 2030) [374].

Section 5.2.1 gives a brief overview of the Phase 1 upgraded-LHCb detector design and the expected performance of the LHCb sub-detectors. An overview of the upgraded LHCb capabilities in the context of LLP searches is given in section 5.2.2 with a few example signatures. Finally, an overview of the plans for the Phase 2 upgrade and some thoughts on the opportunities given by putative additional detector features are reported in section 5.2.3.

5.2.1. *LHCb detector and trigger upgrade for run 3 (Phase 1a).* In LHC Run 3, LHCb plans to take data at an instantaneous luminosity of $2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$, a factor of five larger than the current luminosity. The LHCb detector needs to be upgraded to cope with the higher radiation dose and, most importantly, to avoid the saturation of the trigger rate and exploit the higher luminosity. The main bottleneck in the current trigger is in the first stage, which reduces the accepted rate from 30 to 1 MHz at hardware level. For the upgrade, the hardware trigger will be removed and the full event will be read out at the bunch crossing rate of the LHC (40 MHz), with a flexible software-based trigger.

In table 4 the current and upgraded conditions of the detector are summarised. To cope with the larger occupancy and higher rate of the upgraded detector, the electronics of all of the sub-detectors must be upgraded, and some sub-detectors must be fully replaced. For example, the tracking system, which plays a crucial role in LLP searches, must be replaced.

The upgraded tracking system consists of the VELO, surrounding the IP, the Upstream Tracker, a TT placed before the magnet, and the scintillating fibers tracker (SciFi), three

Table 4. LHCb current and upgraded operating conditions [374]. Instantaneous luminosity \mathcal{L} , integrated luminosity $\int \mathcal{L}$ (including previous runs), pp collision energy \sqrt{s} and average number of visible proton interactions μ are listed.

	Current	Upgrade 1a	Upgrade 1b	Upgrade 2
$\mathcal{L}/(\text{cm}^{-2} \text{s}^{-1})$	4×10^{32}	2×10^{33}	2×10^{33}	2×10^{34}
$\int \mathcal{L}$	8 fb^{-1}	23 fb^{-1}	50 fb^{-1}	300 fb^{-1}
\sqrt{s}	7, 8, 13 TeV	14 TeV	14 TeV	14 TeV
μ	~ 1	~ 5	~ 5	~ 50

stations after the magnet. In the VELO [375], the current strips will be replaced by pixel detectors, with a custom developed ASIC (VeloPix) able to cope with a maximum hit rate of 900 Mhits/s/ASIC. The Upstream Tracker [376] is composed of four silicon micro-strip planes, with finer granularity and larger acceptance compared to the current tracker. Each station of the SciFi has four planes of 2.5 m long scintillating fibres read out by silicon photomultipliers.

The upgrade components most important for LLP searches are the VELO, the tracking and the software trigger. In the following subsections a brief description of the design and capabilities is given for each.

5.2.1.1. VELO upgrade. The VELO plays a fundamental role in LLP searches at LHCb: due to the large boost particles typically experience in the forward direction, a precise measurement of the LLP vertex position allows LHCb to efficiently reject tracks that belong to a different b or c vertex (see section 4.2).

In upgrade conditions, the number of tracks and primary vertices will increase by about a factor of five, making it much more difficult to identify DV close to the beam-line; it is an additional challenge to accomplish this in real time. The VELO was thus completely redesigned [375] to cope with the new expected high-luminosity conditions, maintaining high physics performance and allowing real-time readout for the software trigger. The new VELO has a pixel rather than strip geometry and its distance from the LHC beams is reduced from 8 to 5 mm. This leads to an improvement in the vertex resolution (see figure 47) and reduces the rate of unphysical (ghost) tracks.

The pattern recognition efficiency for track reconstruction is superior to the one of the current VELO when evaluated in high-luminosity conditions. Particularly important for LLP searches is the efficiency of track reconstruction as a function of the displacement from the origin along the beam axis. As shown in figure 48, for the upgraded VELO the efficiency approaches 100% and is uniform in a window of 20 cm around the IP, thanks to the new configuration of the modules in the z direction and the shorter distance from the beam. The VELO acceptance degrades quickly after 20 cm in z , giving an upper limit for LLP decay lengths with vertices reconstructible in the VELO. A display of the upgraded VELO geometry and its acceptance in both the forward and backward directions is shown in figure 49.

Another important metric of detector performance is impact parameter resolution, especially in LLP searches where it can be exploited to reduce the background due to fake tracks. With the upgrade, the impact parameter resolution significantly improves for low- p_T tracks. For example, the impact parameter resolution along x for tracks with p_T of 0.5 GeV is 40 μm in the upgrade versus 70 μm in the current VELO. The replacement of strips with pixel sensors also makes the pattern recognition faster for the same multiplicity. This can be used in




Figure 47. Left: the impact parameter resolution (in the x direction) as a function of the inverse p_T . Right: the resolution on a vertex as a function of number of tracks. The current LHCb VELO performance is shown in black points while the one of the upgraded VELO in red points [375].




Figure 48. Left: a comparison of the current and upgrade VELO z -layouts for LHCb. The top layout (black) is the current VELO while the bottom layout (red) is the upgrade VELO. Right: the track reconstruction efficiency as a function of the origin in z for the current (black) and upgrade (black) VELO in upgrade conditions [375].

the trigger to find tracks and identify DV in the trigger, making it possible to soften or remove inefficient p_T requirements.

Real displaced tracks created by interactions in the VELO material can be a significant background to LLP searches (further details are discussed in chapter 4). The total material budget of the upgraded VELO is similar to that of the current detector (with a radiation length of about 20%) and is dominated by interactions in the radio-frequency foil separating the beam vacuum from the vacuum of the sensors. However, the average percentage of radiation length before the first measured point is significantly reduced in the upgrade VELO, passing




Figure 49. Left: a display of the upgrade VELO geometry and a comparison with the LHCb spectrometer acceptance which is shaded in yellow. Right: the η acceptance of the upgrade VELO geometry both for forward and backward tracks. The fraction of tracks crossing three and four modules is given in red and black, respectively [375].

from $4.6\%X_0$ in the current design to $1.7\%X_0$ in the upgraded VELO (where X_0 is the radiation length).

5.2.1.2. Upgraded trigger. The online event selection in the LHCb experiment during the 2010–2018 running period has been performed by a trigger composed of a hardware level (L0), and two software levels: High Level Trigger 1 (HLT1) and High Level Trigger 2 (HLT2). The L0 reduces the rate from 40 MHz (the LHC bunch-crossing rate) to 1 MHz using information from the calorimeter and MSs. Typical requirements in the L0 are $p_T > 1.4$ GeV for muons and $E > 2.5$ GeV for electrons. The software trigger performs a partial event reconstruction at HLT1, reconstructing tracks and primary vertices for any particle down to $p_T = 500$ MeV, followed by a complete event reconstruction at HLT2, reducing further the rate to 12.5 kHz (in LHC Run 2).

In the Phase 1 upgrade, LHCb foresees to run at a luminosity of $2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$, about a factor of five larger than that experienced during LHC Run 2. For this reason, a new trigger system, able to fully exploit the LHC potential, has been designed [377, 378]. The upgraded LHCb trigger is based on two paradigms: a triggerless readout and a full software trigger. In addition, as already tested in Run 2, a real-time alignment and calibration will achieve offline-quality reconstruction already in the trigger, allowing a higher signal purity of interesting decay channels. Figure 50 shows the current and the upgraded trigger schemes.

Triggerless readout and full-software trigger. With the LHCb trigger upgrade, the 1 MHz readout limitation will be removed, allowing the full event rate to be processed in software. This will increase the efficiency for several channels which otherwise would not benefit from the higher luminosity because they would saturate the L0 trigger rate. Figure 51 shows how the rate for non-muonic B decays saturates the L0 trigger with increasing luminosity. Most of these channels saturate the L0 trigger already at the Run 2 luminosity ($4 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$). Moreover, a purely software trigger will not be subject to the p_T requirements currently applied at the hardware trigger level. Several physics programs involving low- p_T particles, currently prohibited by of the low L0 efficiency, will therefore become possible.




Figure 50. Scheme of the current LHCb trigger (left), and of the upgraded trigger (right) [379].




Figure 51. Trigger yields for $B \rightarrow \pi^+\pi^-$, $B_s \rightarrow \phi(K^+K^-)\gamma$, $B_s \rightarrow J/\psi(\mu^+\mu^-)\phi(K^+K^-)$ and $B_s \rightarrow D_s^-(K^+K^-\pi^-)K^+$ as a function of the luminosity with the current LHCb trigger scheme. For non-muonic channels, the saturation effect due to the L0 trigger can be observed [380].

Turbo stream. Starting from Run 2, offline-quality alignment and calibration has been applied between the HLT1 and the HLT2 levels. This was made possible by the design of a new dedicated trigger output called Turbo Stream. The event record is written directly from

the trigger and processed so that it can be used for physics analysis without the need for offline reconstruction [277]. Several variations of Turbo have been introduced in the last several years. For 2015 data taking, the first version of Turbo allowed only the triggered signal candidate objects to be saved, without keeping the rest of the event and thereby discarding all sub-detector information. While the event size was an order of magnitude smaller than for full stream data, any analysis relying on additional information from the surrounding event could not use Turbo Stream data.

For this reason, full event reconstruction was implemented as Turbo++ in 2016. Finally, a new intermediate solution between Turbo and Turbo++ called Turbo SP (Selective Persistence) was implemented in 2017. With Turbo SP, both the trigger candidate objects and a subset of the other objects in the event are saved. This flexible solution allows the analyzer to choose which objects to save, minimizing the size of the stored event. Once the trigger is upgraded to run purely at the software level, the large majority of LHCb analyses will be moved to Turbo. The reduced event size allows the storage of particle candidates at a high rate, fully exploiting the reduction in p_T thresholds due to the removal of the L0 hardware trigger.

Triggers on downstream tracks. Most of the LLP searches at the LHCb experiment use so-called ‘long tracks’ to reconstruct the candidates, which are tracks where inputs from both the VELO and the TTs are considered. These tracks have an excellent spatial and momentum resolution, and result from an LLP decaying within the VELO region. This is the reason why most of the LLP searches done by LHCb correspond to long-lived candidates with displacements typically up to $\mathcal{O}(10)$ cm, with the VELO tracking algorithm optimized for displacements of up to 20 cm. However, for long-lived candidates with displacements larger than 2 cm, a different type of track which considers information only from the TTs, ‘downstream tracks’, has to be used instead of the ‘long track’ type. Unfortunately, these tracks have worse vertex and momentum resolution, limiting the capabilities of LHCb for this displacement range. In order to improve the detector sensitivity in the displacement ranges between 20 and 200 cm, a proposal to develop new trigger lines to select ‘downstream tracks’ is being studied [378]. An interesting idea which can help to achieve this task and to significantly reduce the CPU computing time is the implementation of a system of specialized processors used to rapidly find the downstream tracks through look-up tables, and present these tracks to the software trigger in parallel with all the raw detector information in the event. This system, named ‘retina’, is being studied and has its own R&D programme [381].

The much higher luminosity and improved capabilities of the upgraded LHCb detector are expected to significantly improve the capabilities for LLP searches in LHC Run 3. In the following sections, the projected sensitivities to several benchmark LLP signatures are shown to illustrate the potential of the upgraded detector. However, the potential of the upgraded triggerless readout has not been completely explored yet and its great flexibility could be exploited in several ways beyond those shown here.

5.2.2. LHCb upgrade phase 1a projections for LLP signatures

5.2.2.1. Displaced di-leptons. The upgraded LHCb experiment is expected to have exceptional sensitivity to low-mass, displaced dilepton signatures thanks to the mass resolution, excellent vertexing, and the online selection allowed by the triggerless readout.

The upgrade LHCb sensitivity to dileptons has been explored in the literature in the context of dark photon searches. Two complementary signatures have been considered: an inclusive search for BSM gauge bosons (dark photons) in $A' \rightarrow \mu^+ \mu^-$, and a search for A' using radiative charm decays $D^{*0} \rightarrow D^0 A', A' \rightarrow e^+ e^-$. The inclusive search [341] scans a




Figure 52. Current and expected limits in the dark photon parameter space mixing ϵ^2 versus A' mass [374]. The black line represents current LHCb limits from [269], while grey-shaded regions are existing limits from other experiments. Expected limits from the proposed inclusive search with 50 and 300 fb^{-1} are shown in shades of blue. The expected sensitivity of $D^{*0} \rightarrow D^0 A'(ee)$ at low mass is shown in shades of green. The arrows indicate the available mass range from light meson decays into $e^+ e^- \gamma$.

large region from the dimuon threshold $2m_\mu$ all the way up to the Z pole. The second proposed signature [382] exploits a tag of the radiative decay of the D^{*0} using its reconstructed invariant mass and a di-electron dark photon final state to probe a much lower mass range allowed by the decay kinematics, [2 m_e , 142 MeV]. For both signatures, searches for both a prompt and a displaced dark photon vertex are carried out. The prompt search is expected to probe mixing parameters ϵ^2 below 10^{-7} despite the large irreducible background from Drell–Yan and QCD¹⁸⁰. Since the dark photons with masses above the η mass decay promptly for couplings that are accessible within LHCb, no attempt is made to probe displaced dark photons in that region.

An inclusive search for dark photons has already been performed with LHCb data collected in 2016 [269]. This was possible due to the high reconstruction and identification efficiency of soft di-muons at LHCb. These results demonstrated the unique sensitivity that can be reached at LHCb. The planned increase in luminosity and removal of the hardware-trigger stage in Run 3 should increase the number of expected $A' \rightarrow \mu\mu$ decays in the low-mass region by a factor of $\mathcal{O}(100\text{--}1000)$ compared to the 2016 data sample. The limits placed by the current data and the sensitivity expected with future LHC runs is shown in figure 52.

The exclusive search for $D^{*0} \rightarrow D^0 A', A' \rightarrow e^+ e^-$ is much more challenging and not feasible prior to the upgrade, since the hardware trigger and the higher material budget degrade the sensitivity. This search highly relies on the online identification of $e^+ e^-$ pairs, since over 5 trillion of these D^{*0} decays are expected in Run 3. The expected sensitivity probes unexplored regions of phase space at low A' mass and mixing ϵ^2 that is usually in the realm of beam-dump experiments.

A displaced di-muon signature also appears in some HV scenarios [66], in which a hidden sector with strong dynamics showers and hadronizes into dark mesons that can have an appreciable decay rate to leptons. (For more information on dark showers, see chapter 7.) The upgraded LHCb prospects for this type of signature have been explored in [383] and are very promising. In the scenario explored therein, dark mesons are produced with large multiplicities of between 10 and 30. Selection criteria inspired by the proposed dark photon search [341] in the region after the first VELO module are applied. The expected reach for the

¹⁸⁰ In the limit $m_{A'} \ll m_Z$, the coupling of the dark photon to SM particles with charge Q is approximately $Q\epsilon\epsilon$.




Figure 53. Projected bounds from various ATLAS/CMS searches and the LHCb search for hidden valley-like models with a Z_p decaying to a $q_{\text{HV}}q_{\text{HV}}$ pair which then undergo hadronization in the dark sector, eventually leading to dark meson decays to di-muon final states [383].

proposed searches using Run 3 data from LHCb (15 fb^{-1}) and from ATLAS/CMS (300 fb^{-1}) are shown in figure 53. The model studied involves a $200 \text{ GeV } U(1)'$ gauge boson Z_p decaying to a HV quark pair; showering and hadronization in the dark sector leads to a large multiplicity of hidden hadrons ω_V (with $m_{\omega_V} = 0.3 \text{ GeV}/c^2$) that can decay to di-muons. In this context, the upgraded LHCb detector could have better sensitivity than other proposed searches at ATLAS and CMS.

5.2.2.2. Displaced jets. Signatures with displaced jets are common in the context of LLP searches. As summarised in section 3.1.3, LHCb has started to explore its sensitivity to displaced jets by using the data collected during Run 1. In LHCb Upgrade Ia, the background rejection for displaced jet searches is expected to improve thanks to the improved VELO resolution. The selection efficiency should also be significantly higher due to the prospects for online DV identification. The main focus of LHCb will be to probe the region at small lifetime where it has already proven to be the most competitive (see discussion in section 3). The background from QCD and material interactions is the main limiting factor, and the improved vertex resolution of the upgraded VELO together with the lower material budget and the use of a detailed material map are expected to bring large improvements in the upgraded detector.

In some dark sector models, the number of hadronic DV can be large, even in the limited LHCb acceptance (see, for example, [331] and the discussion in chapter 7). An inclusive DV search would likely include a requirement on the isolation of the DV from other tracks in the event (such as the one used for the dark photon analysis [269]). This requirement could be very inefficient in the context of a dark shower and so a dedicated search strategy would be needed. Furthermore, a dedicated software trigger looking for a large number of DV in the VELO and soft p_T requirements could in principle improve the sensitivity to this kind of models, but studies are needed to fully understand its potential and compare LHCb to other experiments.




Figure 54. Projected bounds from various proposed searches for confining hidden valley models exploiting $c\bar{c}$ decays of hidden valley mesons ω_v (taken from [383]). The sensitivities with various signatures at LHCb are shown: two displaced vertices (green), one reconstructed D meson and one DV (blue) and two reconstructed D mesons (red). Searches for this same particular model at ATLAS/CMS are shown not to be competitive (more details can be found in [383]).

5.2.2.3. Displaced mesons. Quark–antiquark pairs from the decay of a low mass particle in the SM often hadronize into SM mesons that subsequently decay with known branching ratios. For example, a dark sector meson with a displaced decay to $c\bar{c}$ often produces two D mesons which in turn have non-negligible lifetimes. In this scenario, the authors of [383] have investigated the prospects of using more or less inclusive reconstruction of the two D meson decays at the upgraded LHCb. A similar approach could be used to target decays to $b\bar{b}$ that hadronize to B mesons since the latter is likely to produce D mesons in its decay. Since LHCb is designed to reconstruct heavy flavor decays, it can be competitive in this kind of search [383] (see figure 54). Furthermore, searches for displaced mesons will greatly profit from the software trigger, which is mainly designed to improve the efficiency of similar-looking hadronic decays of heavy flavor mesons.

5.2.3. After phase 1a upgrade: phase 1b and phase 2. After a consolidation phase (Phase 1b) during long shutdown (LS 3) aiming to run in the same conditions as in Phase 1a, the LHCb detector plans a Phase 2 upgrade during LS 4 (by 2030). The Phase 2 upgrade is needed to face even more challenging conditions than during previous runs. For example, particle multiplicity, PU, and radiation damage are expected to be ten times higher than those experienced in Phase 1. The LHCb experiment expects to collect at least 300 fb^{-1} by the end of Upgrade 2 [374].

Major improvements to the LHCb detector during Phases 1b and 2 that are relevant for LLP searches are described in the following, as well as some naïve projections of Run 1 results to an integrated luminosity of 300 fb^{-1} .

Magnet stations. Along with the ‘long’ and ‘downstream’ track types, ‘upstream’ tracks are also considered useful for LLP searches. These tracks correspond to soft charged particles bending out of the detector acceptance, produced from LLP candidates which decay within the VELO region. Aside from the installation of a new tracker during LS2, the Upstream Tracker, a proposal to add magnet stations inside the LHCb magnet to improve low




Figure 55. Impact parameter resolution estimations from simulation using the Phase 1 VELO model for the LHCb experiment: (left) a comparison if considering Phase 1 or Phase 2 conditions; (right) the effect of the RF-foil removal under Phase 2 conditions [384].

momentum resolution is considered [384]. These magnet stations have been proposed to be installed for Phase 1b, and are foreseen to highly improve the tracking of low momentum particles produced from certain kind of LLP candidates, such as for example soft pions from ‘disappearing’ chargino tracks.

Material interactions. The presence of a VELO envelope at approximately 5 mm from the beam line in order to separate the VELO and the LHC vacuums, named ‘RF-foil’, strongly affects the background composition of LLP searches in the LHCb experiment. Namely, for LLP candidates decaying below 5 mm from the beam line, the main source of background is due to heavy flavour decays, while material interactions with the RF-foil compose the main background contribution for LLPs decaying above 5 mm from the beam line. While the former is purely due to QCD processes and hence not reducible, the latter is kept under control by the use of a detailed veto map (see [269]). However, the ideal case would be to completely remove the RF-foil during the Phase 2 upgrade [384], which would result in a large reduction of the background component due to material interactions. The improvements foreseen to the Phase 2 VELO (probably based on an updated version of the Phase 1 VELO), are expected to increase the sensitivity to shorter lifetimes and better PV and impact parameter resolution (see figure 55). Unfortunately, the removal of the RF-foil requires the development of new techniques to isolate the sensors from the beam radio frequency and is not necessarily seen as a viable option. Therefore, improving the material veto maps as much as possible by accurately modeling the material interactions would be desirable as a more realistic option for the HL-LHC era.

Naïve Projections of Run 1 Results to the HL-LHC Era. By taking the published Run 1 results from the single displaced dijet search [253] at LHCb, a naïve extrapolation to the integrated luminosity foreseen to be recorded by LHCb during Phase 2, 300 fb⁻¹, is presented in figure 56. These numbers have been obtained by simply scaling signal and background with the expected increase in cross section and luminosity, neglecting PU effects and expected detector improvements. The removal of neutral objects from jet reconstruction, and the use of machine-learning techniques are expected to assist in the required suppression of




Figure 56. Naive sensitivity projections for searches of displaced dijet vertices at LHCb for the expected luminosity to be collected with the Phase 2 upgraded detector (300 fb^{-1}) [385]. The projected limit is on Higgs boson decays to dark pions π_V with branching fraction \mathcal{B} , and the dark pions each in turn decay to a pair of jets.

PU. Furthermore, jet substructure techniques are foreseen to improve the quality of di-jet object with lower mass.

5.3. Dedicated detectors for LLPs

5.3.1. Introduction. Despite a wide and seemingly comprehensive research program—both existing and, in this document, proposed to be expanded—for LLPs at the LHC, in cases of ultra-low-mass particles, ultra-long lifetimes, or unusual LLP charges, it is hard or impossible to trigger on and/or reconstruct such events in the main ATLAS, CMS, and LHCb detectors. This has led to new proposals for dedicated experiments to look for LLPs in new regimes that are otherwise inaccessible at the LHC. These experiments provide the best sensitivity to new millicharged LLPs, magnetic monopoles, and other LLPs arising from models such as those containing Higgs-portal hidden sectors, dark photons, and Majorana neutrinos.

As discussed in section 2.2, LLPs BSM are theoretically well motivated and come in wide range of masses and lifetimes. ATLAS and CMS have excellent sensitivity for fairly high mass LLPs, regardless of their lifetime (see, e.g. [4, 246, 386, 249]). Low mass and/or softer final states are more challenging due to both background and triggering limitations. In the short-lifetime regime, for $c\tau$ of the scale of the VELO, LHCb has sensitivity to somewhat lower masses and can trigger on softer muonic final states, generating complementary reach provided the LLP has a significant branching ratio to muons [253, 254, 279, 309, 387]. Finally, the low-mass/soft final states with rather long lifetimes are challenging for all three experiments. These signatures can be covered partially by NA62 [388] operating in beam dump mode, or by SHiP [389], or by dedicated LHC experiments like CODEX-b [104] (see section 5.3.5), FASER [105] (see section 5.3.6), or MATHUSLA [103] (see section 5.3.4). Each of these dedicated experiments is sensitive to different LLP lifetimes, masses, and production modes based on their position and orientation and thus each can be considered a necessary component of a comprehensive, coordinated search program for very long-lived particles at the LHC.




Figure 57. A three-dimensional schematic view of the MoEDAL detector (in the yellow circle) around the LHCb VELO region at Point 8 of the LHC.

5.3.2. MoEDAL experiment and future developments. Monopole and exotics detector at the LHCbEDAL (MoEDAL) [101]¹⁸⁰ is designed to search for manifestations of new physics through highly-ionising (HI) particles in a manner complementary to ATLAS and CMS [310]. The main motivation for the MoEDAL experiment is to pursue the quest for magnetic monopoles at LHC energies. Nonetheless the detector is also designed to search for any massive, long-lived, slow-moving particle [390, 391] with single or multiple electric charges arising in many scenarios of physics BSM [392].

The MoEDAL detector [311] is deployed around the intersection region at LHC Point 8 (IP8) in the LHCb VELO cavern. A schematic view of the MoEDAL experiment is shown in figure 57. It is a unique and largely passive detector comprising different detector technologies.

The main sub-detector system is made of a large array of CR-39, Makrofol® and LexanTM nuclear track detector (NTD) stacks surrounding the intersection area. The passage of a HI particle through the plastic detector is marked by an invisible damage zone along the trajectory. The damage zone is revealed as a cone-shaped etch-pit when the plastic detector is chemically etched. Then the sheets of plastics are scanned looking for aligned etch pits in multiple sheets. The MoEDAL NTDs have a threshold of $z/\beta \sim 5$, where z is the charge and $\beta = v/c$ the velocity of the incident particle.

Another type of NTD installed is the very high charge catcher ($z/\beta \sim 50$). It consists of two flexible low-mass stacks of Makrofol®, deployed in the LHCb acceptance between RICH1 and the Trigger Tracker. It is the only NTD (partly) covering the forward region, adding only $\sim 0.5\%$ to the LHCb material budget while enhancing considerably the overall geometrical coverage of MoEDAL.

A unique feature of the MoEDAL detector is the use of paramagnetic magnetic-monopole trappers (MMTs) to capture magnetically-charged HI particles. The high magnetic charge of a monopole—being at least one Dirac charge $g_D = 68.5e$ —implies a strong magnetic dipole moment, which may result in strong binding of the monopole with the nuclei of the aluminium MMTs. In such a case, the presence of a trapped monopole would be detected through the induction technique by measuring the *persistent current*, defined as the

¹⁸⁰ For general information on the MoEDAL experiment, see: <http://moedal.web.cern.ch/>.

difference between the superconducting magnetometer currents before and after the passage of the MMT bar through the sensing coil [393, 394].

The only non-passive MoEDAL sub-detector is an array of TimePix pixel devices distributed throughout the MoEDAL cavern, forming a real-time radiation monitoring system of HI beam-related backgrounds. The operation in time-over-threshold mode allows a 3D mapping of the charge spreading in the volume of the silicon sensor, thus differentiating between various particles species from mixed radiation fields and measuring their energy deposition.

The MoEDAL detector is designed to fully exploit the energy-loss mechanisms of magnetically charged particles [395–398] in order to optimise its potential to discover these messengers of new physics. Multiple theoretical scenarios have been proposed over the years in which magnetic charge would be produced at the LHC [392], resulting in such possible new particles as light 't Hooft–Polyakov monopoles [397–399], electroweak monopoles [400–404], global monopoles [405–410], and monopolium [396, 411–413]. Magnetic monopoles that carry a non-zero magnetic charge and dyons possessing both magnetic and electric charge are predicted by many theories including grand-unified and superstring theories [414–416].

A possible explanation for the non-observation of monopoles so far is Dirac's proposal [395, 396, 411] that monopoles are not seen freely because they form a bound state called *monopolium* [412, 413, 417, 418] being confined by strong magnetic forces. Monopolium is a neutral state, difficult to detect directly at a collider detector, although its decay into two photons would give a rather clear signal for ATLAS and CMS [419]. Nevertheless the LHC radiation detector systems can be used to detect final-state protons $pp \rightarrow pXp$ exiting the LHC beam vacuum chamber at locations determined by their fractional momentum losses [420]. Such a technique would be appealing for detecting monopolia.

The MoEDAL detector is also designed to search for any massive, long-lived, slow-moving particles [390, 391] with single or multiple electric charges arising in many scenarios of physics BSM. Supersymmetric LLPs [421], quirks, strangelets, Q-balls, and many others fall into this category [392]. A generic search for high-electric-charge objects is currently underway.

For the 2016 run at 13 TeV, the MMT included 672 aluminium rods (for a total mass of 222 kg) that were placed 1.62 m from the IP8 LHC IP under the beam pipe on the side opposite to the LHCb detector. The MMT bars were analysed and no magnetic charge $> 0.5g_D$ was detected in any of the exposed samples when passed through the ETH Zurich SQUID, which is a DC SQUID long-core magnetometer [422]. Hence cross section limits are obtained for Drell–Yan pair production of spin-1, spin-1/2 and spin-0 monopoles for $1g_D \leq |g| \leq 5g_D$ at 13 TeV [422] improving previous bounds set by MoEDAL at 8 TeV [311] and 13 TeV [423]. Monopole production via photon fusion is also now considered in MoEDAL monopole search analyses [424] following recent studies [425]. However, the large monopole-photon coupling invalidates any perturbative treatment of the cross section calculation and hence any result based on the latter is only indicative. This situation may be resolved if thermal production in heavy-ion collisions—that does not rely on perturbation theory—is considered [426], or by including a magnetic-moment term in monopoles with spin [425].

To recapitulate, under the assumption of Drell–Yan cross sections, MoEDAL has derived mass limits for $1g_D \leq |g| \leq 5g_D$, complementing ATLAS results [215, 427], which placed limits for monopoles with magnetic charge $|g| \leq 1.5g_D$, as shown in figure 58. The ATLAS bounds are better than the MoEDAL ones for $|g| = 1g_D$ due to the higher luminosity delivered in ATLAS and the loss of acceptance in MoEDAL for small magnetic charges. On the other




Figure 58. Magnetic monopole mass limits from CDF [429], ATLAS [215] and MoEDAL searches [311, 422] as a function of magnetic charge for various spins, assuming a Drell–Yan pair-production mechanism.

hand, higher charges are difficult to be probed in ATLAS due to the limitations of the level-1 trigger deployed for such searches. Limits on production cross sections of singly-charged magnetic monopoles set by various colliders are presented in [414, 415], while general limits including searches in cosmic radiation are reviewed in [428].

MoEDAL is proposing to deploy the MoEDAL apparatus for detecting penetrating particles (MAPP) in a tunnel shielded by some 30–50 m of rock and concrete from the IP8 [430, 431]. A prototype of the MAPP detector was installed in 2017. It is envisaged that the full detector will be installed in LS2 to take data in LHC Run 3. The purpose of the detector is to search for particles with fractional charge as small as one-thousandth the charge of an electron. This detector would also be sensitive to neutral particles from new physics scenarios via their interaction or decay in flight in the ~ 10 m decay zone in front of the detector or in the detector itself. The isolation of the detector means that the huge background from SM processes in the main detectors is largely absent. Also, the detector can be placed at various angles to the beam axis (from 5° to 25°). The ability to vary depth and angle enhances MoEDAL to be able to distinguish between theoretical scenarios in the event a signal is observed.

The first apparatus specifically designed to detect mini-charged particles was the Stanford Linear Accelerator Centre (SLAC) ‘beam dump’-type detector, comprising scintillator bars read out by photomultiplier tubes [432]. MoEDAL’s new detector, shown in figure 59, and milliQan (discussed above in section 5.3.3) proposed for deployment near to the CMS detector [102] also designed to search for mCPs, both have a design that harks back to the original SLAC detector. In order to reduce backgrounds from natural radiation the photomultiplier tubes and scintillator detectors of the MoEDAL apparatus will be constructed from materials with low natural backgrounds currently utilised in the astroparticle-physics arena. Its calibration system utilises neutral density filters to reduce the received light of high incident muons that manage to penetrate to the sheltered detector from the IP, in order to mimic the much lower light levels expected from particles with fractional charges.




Figure 59. A depiction of the MoEDAL's MAPP sub-detector [431].




Figure 60. The MALL sub-detector designed to monitor MoEDAL trapping volumes for the decays of trapped electrically charged particles with lifetimes as long as 10 years [431].

MoEDAL is also planning another new sub-detector called MALL (MoEDAL apparatus for detecting extremely long lived particles) [431]. In this case MoEDAL trapping volumes, after they have been scanned through the ETH Zurich SQUID facility to identify any trapped monopole will be shipped to a remote underground facility to be monitored for the decay of pseudo-stable massive charged particles that may also have become trapped. MALL is the detector that monitors MoEDAL trapping volumes for decays of captured particles. It is envisaged that MALL will be installed deep underground at SNOLAB in Canada where cosmic backgrounds are minimised to one muon per $0.27 \text{ m}^2 \text{ d}^{-1}$. Background is further reduced by the ability to determine if a detected track originated within the monitored volume and also by energy cuts on deposited signals. A schematic view of the detector is shown in figure 60. Initial estimates indicate that lifetimes up to around 10 years can be probed. The MALL detector is designed to be sensitive to charged particles and to photons, with energy as small as 1 GeV. It is envisaged that construction of the detector will begin after the MAPP detector is full deployed.

The possibility of analysing decommissioned parts of the LHC beam-pipe system at the ATLAS, CMS and LHCb/MoEDAL sites with a SQUID to search for trapped magnetic monopoles has been proposed [433]. In this context the MoEDAL experiment may serve as a formal platform for coordinating machining, scanning and analysis work, in collaboration with interested ATLAS, CMS and LHCb members.

The induction technique has been successfully employed at the LHC with the dedicated MoEDAL trapping detector. Additional searches for trapped monopoles in beam pipe material would access wide windows of magnetic charges and production cross sections to which other LHC experiments are insensitive. The decommissioned central beryllium beam pipe sections of ATLAS and CMS, with a 4π coverage and exposure to the highest rates of 7 and 8 TeV pp collisions, are by far the most attractive samples to be analysed. The analysis on the CMS beam pipe is expected to be carried out in 2019.

5.3.3. The milliQan experiment. MilliQan is a dedicated experiment at the LHC to search for mCP [102, 434]. The milliQan experiment is part of a general program to search for hidden sectors [435] and other BSM scenarios [436]. As an illustrative example, we show the sensitivity of milliQan to an extra Abelian gauge field coupled to a massive Dirac fermion ('dark QED') that mixes with hypercharge through the kinetic term [122]. The result is that the new matter field is charged under hypercharge with a fractional electric charge of ϵ , where $\epsilon \ll 1$. The milliQan experiment targets an unexplored part of the parameter space, namely mCP masses $0.1 \lesssim M_{\text{mCP}} \lesssim 100$ GeV, for charges Q at the $10^{-3} e - 10^{-1} e$ level.

The experimental apparatus consists of three scintillator detector layers of roughly 1 m^3 each, positioned near one of the high-luminosity IPs of the LHC. The experimental signature consists of a few photo-electrons (PE) arising from the small ionization produced by the mCPs that travel unimpeded through material after escaping the LHC detectors.

The milliQan experiment is planned to be sited in the PX56 Observation and Drainage gallery above the CMS underground experimental cavern. The proposed gallery is limited in space. The detector will be located in this tunnel at an optimized location that is 33 m from the CMS IP, behind 17 m of rock, and at an angle of 43.1 degrees from the horizontal plane. The selected location in a 3D model is shown in figure 61.

The milliQan detector is a $1 \text{ m} \times 1 \text{ m} \times 3 \text{ m}$ plastic scintillator array. The array will be oriented such that the long axis points at the nominal CMS IP. The array is subdivided into 3 sections each containing 400 $5 \text{ cm} \times 5 \text{ cm} \times 80 \text{ cm}$ scintillator bars optically coupled to high-gain photomultiplier tubes (PMTs). The detector will be shielded from other background sources such as activity in the scintillator and environmental radiation. With an estimated detection efficiency of about 10%, milliQan expects an average of $\mathcal{O}(1)$ PE from each attached phototube for each mCP with $Q = \mathcal{O}(10^{-3}) e$ that traverses 80 cm of plastic scintillator. The signal is a longitudinal triple-incidence of hits with one or more PEs; a triple-incidence within a 15 ns time window along longitudinally contiguous bars in each of the three sections is required to reduce the background from dark-current noise. Requiring triple-incidence is expected to control background to $\mathcal{O}(10)$ events per year with $N_{\text{PE}} \geq 1$. The milliQan detector will self-trigger to a dedicated readout with no dead time from readout, up to rates of ~ 1 kHz. Energy calibration will be done *in situ* using an ^{241}Am source.

The dominant background is expected to come from dark-current pulses in the PMTs. Pulses from background radiation, including cosmic muons, will consist of 1000s of PEs that can be easily vetoed offline. Assuming a total background rate per PMT of $\nu_B = 500$ Hz, with a time window of $(\Delta t)_{\text{online}} = 100$ ns, milliQan expects a double coincident trigger rate per board of 1.5 Hz. The entire detector will be read out if one board triggers and there will be 50 such boards in total. Therefore, the full background trigger rate is expected to be 75 Hz.




Figure 61. A 3D model showing the optimal position of milliQan within the PX56 Drainage and Observation gallery located above CMS UXC.

Offline, the time window will be tightened to $(\Delta t)_{\text{offline}} = 15$ ns, yielding an offline background rate for a triple coincidence of 2.8×10^{-8} Hz. Since there are 400 such sets, the total offline background rate is estimated to be 1.1×10^{-5} Hz. With these background rates, milliQan estimates a total of 165 (330) background events in 300 (3000) fb^{-1} of integrated luminosity.

The milliQan collaboration has performed a full simulation of the experiment to evaluate the projected sensitivity to various mCP electric charges and masses. The simulation is performed in two stages. In the first, the production of mCP particles via Drell–Yan, J/Ψ , Υ (1S), Υ (2S), and Υ (3S) channels at a center-of-mass energy of 14 TeV is performed. Particles produced at the IP are propagated to the proposed experimental site described above using a map of the CMS magnetic field. The effects of multiple scattering and energy loss are included using a simplified model of the CMS detector material budget and a region of rock spanning 17 m between the CMS experimental cavern and the proposed experimental site. The number of expected mCP particles per fb^{-1} of integrated luminosity incident at the detector is computed as a function of the mass of the mCP. The signal efficiency is then estimated after processing the calculated particles as they would emerge from the IP through a full GEANT4 simulation of the detector; this is necessary because the small charge regime is sensitive to details such as the reflectivity, the light attenuation length, and the shape of the scintillator. These details, as well as the quantum efficiency, light-emission spectrum and the fast-time constants were modeled in GEANT4 using the specifications provided by the manufacturers for the scintillator and PMTs. Combining the estimated background rates discussed above with the cross sections, acceptances and efficiencies calculated for all masses and electric charges, the sensitivity projections of the milliQan experiment for LHC and HL-LHC are shown in figure 62.

A 1/100 scale ‘demonstrator’ of milliQan to validate the detector concept was installed in the PX56 location at CERN during Technical Stop 2 of 2017 and was upgraded during the 2017–2018 year-end technical stop. This demonstrator, shown in figure 63, has been recording data since its installation, and expects to have first results later this year. If funding




Figure 62. The expected sensitivity of the milliQan experiment for different LHC luminosity scenarios. The black line shows the expected 95% C.L. exclusion (solid) and 3σ sensitivity (dashed), assuming 300 fb^{-1} of integrated luminosity. In blue is shown the corresponding expectations for 3000 fb^{-1} .



Figure 63. The 1/100 scale ‘demonstrator’ of milliQan installed in the PX56 drainage gallery 33 m from the CMS IP.

is secured, construction of the full milliQan apparatus is planned for 2019, with installation in the tunnel in 2020.

5.3.4. The MATHUSLA experiment. The basic motivation for the MAssive Timing Hodoscope for Ultra-Stable neutrAL pArticles (MATHUSLA) detector [103] is the search




Figure 64. Simplified MATHUSLA detector layout showing the position of the $200\text{ m} \times 200\text{ m} \times 20\text{ m}$ LLP decay volume.

for LLPs with lifetimes much greater than the size of the LHC main detectors, $c\tau \gg 100\text{ m}$. 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 at the LHC, suppression of backgrounds is therefore crucial for discovery.

For LLP searches with high-energy or lepton-containing final states, the spectacular nature of displaced-vertex (DV) signals leads to very low backgrounds in searches at LHC detectors such as ATLAS or CMS. Any other class of LLP signature suffers from backgrounds and triggering limitations that can be significant. This greatly curtails the main detectors' ability to discover LLPs with very long lifetimes.

To address this broad blind spot of existing detectors, MATHUSLA is proposed to be a large, relatively simple surface detector that can robustly reconstruct DVs with good timing resolution. This gives MATHUSLA a similar geometric acceptance to LLP decays in the long-lifetime limit as the main detectors, while providing shielding from QCD backgrounds and sufficient tracking to reject ubiquitous CRs. As a result, MATHUSLA is able to detect LLPs produced with $\sim 1\text{ pb}$ cross sections at the HL-LHC with lifetimes near lifetimes of $\sim 0.1\text{ s}$, which is generally the limit imposed by big-bang nucleosynthesis (BBN).

The simplified detector design for MATHUSLA is shown in figure 64. The main component of the detector is a 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 should have on the order of five planes to provide robust tracking with $\sim \text{ns}$ timing and $\sim \text{cm}$ spatial resolution. The current MATHUSLA design employs proven and relatively cheap technologies to allow for MATHUSLA's construction in time for the HL-LHC upgrade. Therefore, the trackers are envisioned to be implemented with resistive plate chambers (RPCs), which have been used for very large area experiments in the past [437, 438], while plastic scintillators provide the surrounding veto.

This minimal design has been shown to be capable of measuring the LLP boost on an event-by-event basis [439] using the geometry of the LLP visible decay products. Final-state multiplicity provides a straightforward discriminant between hadronic and electromagnetic decays. Additional particle identification capability, as well as detection of final-state photons, might be possible by inserting an additional material layer between tracking layers to induce an electromagnetic shower that can be used to distinguish electrons, photons and muons.

As argued in [103], MATHUSLA could search for LLPs decaying into charged particles with little or no backgrounds. In figure 64 is shown, schematically, the two main MATHUSLA signals, LLPs decaying into at least two charged leptons, or into jets that




Figure 65. Sensitivity of MATHUSLA to hadronically decaying LLPs produced in exotic Higgs decays, where the solid lines correspond to four decays in the detector [103]. The dotted lines are the most optimistic ATLAS projection, using a very inclusive search for a single DV in the muon chamber [262].

contain $\mathcal{O}(10)$ charged hadrons [439]. In 50%–90% of leptonic decays and practically 100% of hadronic decays [439], two or more charged particles hit the ceiling due to the LLP boost and are recorded by the tracker. The charged particle trajectories can be fitted to reconstruct a DV. Unlike most 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 originate at the DV: there can be no hits along the line between the vertex and the LHC main IP, nor along the lines obtained by extrapolating the individual charged particle trajectories backwards. These exhaustive geometric and timing requirements make it extremely difficult for backgrounds to fake the LLP signal. Cosmic rays can be rejected since they travel in the wrong direction (as well as occurring mostly in CR showers with extremely highly correlated particle multiplicity in the detector). Muons from the LHC collision do not satisfy the DV signal requirement. The rare occasions in which muons scatter inside the decay volume can be vetoed with the scintillators. Finally, neutrinos from CRs and LHC collisions can scatter in the decay volume and produce DV, but those can also be rejected with geometric and timing requirements. We refer the reader to [103] for more details. More comprehensive studies of these backgrounds and their rejection strategies, including full simulations, are currently in progress.

An important general class of signals are LLPs with masses $\lesssim 100$ GeV that decay hadronically and are produced without other highly visible signals like high MET or high-energy leptons. MATHUSLA can improve the sensitivity to these LLP production cross sections by a factor of $\sim 10^3$ compared to searches with the main detectors alone. This is illustrated in figure 65, which shows MATHUSLA’s sensitivity to hadronically decaying LLPs produced in exotic Higgs decays [103] compared to an optimistic projection for searches for a single DV in the ATLAS MS [262]. For branching ratios of $\sim 10\%$, the BBN lifetime limit [440] can be probed.

Apart from reaching the BBN upper limit or ceiling of LLP parameter space, MATHUSLA extends the power of other LHC measurements. If a $\sim 10\%$ invisible branching ratio for the Higgs was detected via coupling measurements at the HL-LHC, the absence of a MATHUSLA signal would lend strong support to the interpretation that the Higgs decayed to a stable component of DM.

Beyond LLPs produced in rare Higgs decays, MATHUSLA is a general-purpose LLP detector that is sensitive to a wide range of other models. Its reach for other BSM scenarios has been explored in [172, 174, 198, 441–447]. The general physics case for MATHUSLA’s construction is made systematically in the recently released whitepaper [2].

It is important to point out that MATHUSLA is a very flexible detector concept that is completely scalable. The sensitivity of MATHUSLA is roughly proportional to the decay volume (which scales with surface area) and roughly independent of the precise geometry or location of the detector, as long as it is $\mathcal{O}(100\text{ m})$ horizontally displaced from the IP. Therefore, a detector with smaller or larger volume than the benchmark in figure 64 may be implemented, depending on available space and budget. Furthermore, the detector volume can be divided into smaller, independent modules (which cooperate for triggering purposes). These can be mass-produced economically and arranged according to the requirements of the experimental site. Such a modular construction also allows for a natural way to improve or extend the physics program of the experiment in an efficient way, by eventually upgrading, for example, one or a few modules with additional capabilities, such as higher tracking resolution for reconstruction of very low-mass LLPs below 10 MeV, or an additional material layer between the trackers for particle ID.

The MATHUSLA collaboration is currently studying such a modular design with the aim of producing a Letter-of-Intent in 2018. Crucial to this endeavor is the data from the MATHUSLA test stand, a $\sim 3 \times 3 \times 5\text{ m}$ MATHUSLA-type detector that took a few days of data in the ATLAS instrument hall in 2017 and will take data for a few months in 2018. This allows local CR backgrounds to be measured, background rejection and signal reconstruction strategies to be tested, and simulation frameworks to be calibrated.

In conclusion, the MATHUSLA detector concept calls for a dedicated LLP surface detector above ATLAS or CMS. A detector volume of $\sim 10^6\text{ m}^3$ gives sensitivity to LLPs near the BBN lifetime limit if they are produced with $\sim\text{pb}$ cross section. This improves LLP sensitivity, compared to the main detectors alone, by several orders of magnitude for many LLP scenarios. The detector is simple and relies on proven technology, making its construction in time for the HL-LHC upgrade feasible. Once constructed, MATHUSLA could function without modification as a detector for the HE-LHC (with increased sensitivity). A small-scale test stand detector is already taking data at CERN, and studies are underway to finalize a detailed design for the full-scale detector.

5.3.5. A COntract for EXotics at LHCb (CODEX-b). The CODEX-b proposal involves housing a small detector in the LHCb cavern—external to the LHCb detector itself—in a space approximately 25 m from the interaction point (IP8), behind the 3 m thick concrete UXA shield wall. This space is presently occupied by the LHCb data acquisition (DAQ) system, but will become available before Run 3 once the DAQ system is relocated to the surface. The layout of the cavern is shown in figure 66, with the location of CODEX-b overlaid. The nominal CODEX-b configuration features a $10 \times 10 \times 10\text{ m}$ volume instrumented with RPC tracking layers or other off-the-shelf tracking technology, as well as roughly 25 interaction lengths of shielding near IP8—e.g. 4.5 m of Pb—to suppress primary and secondary K_L , neutron, and other hadronic backgrounds. This shield requires an active muon veto with an efficiency of $\mathcal{O}(10^{-5})$, in order to reject backgrounds induced by muons or other charged particles in the downstream parts of the shield. The veto is located several metres within the shield such that backgrounds induced by neutral particles remain suppressed. See [104] for a study of a proof-of-concept example detector layout and corresponding tracking efficiency, as well as a detailed study of the backgrounds. More




Figure 66. Layout of the LHCb experimental cavern UX85 at point 8 of the LHC [448], overlaid with the proposed CODEX-b location. From [104].

ambitious technologies, including calorimetry, precision TOF measurements, or integration into the LHCb readout may also be feasible.

For the current discussion, the reach of CODEX-b for two benchmark models is quantified: first, a light scalar field, φ , that mixes with the SM Higgs boson is considered. If $m_\varphi \lesssim 5$ GeV, the production mode is primarily through inclusive $b \rightarrow s\varphi$ decays [449–451]. The LLP φ subsequently decays back to SM fermions through the same Higgs portal. The reach in terms of m_φ and the mixing angle s_θ is shown in the left-hand panel of figure 67. CODEX-b significantly extends the projected reach of LHCb using only VELO-based DV reconstruction, and covers part of the parameter space to which SHiP [452] and MATHUSLA [443] are projected to be sensitive. Studies of the potential LHCb reach to longer lifetimes using downstream tracking are ongoing [378, 453]. The right-hand panel of figure 67 indicates the reach for more general models, where the lifetime and production rate of φ are unrelated.

For the second benchmark, a dark boson, γ_d , produced through the exotic Higgs decay $h \rightarrow \gamma_d \gamma_d$ is considered. For concreteness the γ_d is taken to be a spin-1 field which can decay through mixing with the SM photon [70, 230, 456, 457]. In this benchmark, the production and decay are therefore controlled by different portals. The projected reach is shown in figure 68, overlaid with the reach of ATLAS [262, 273] and MATHUSLA [103]. In particular, at low γ_d masses, CODEX-b complements and significantly extends the reach of ATLAS and CMS.

Finally, it might be possible to install a larger version of the CODEX-b detector at IP2, after the ALICE collaboration concludes its heavy ion program. This option is named ‘A Laboratory for Long-Lived eXotics’ (AL3X) [459], and can make use of the existing ALICE TPC, supplemented with a thick hadron absorber. The B-field from the L3 magnet would provide a good momentum measurement of the tracks, something which is absent in MATHUSLA and CODEX-b. The feasibility of this proposal is however contingent upon a luminosity upgrade of IP2, as well as the possible continuation of the ALICE physics during Run 5.

5.3.6. The forward search experiment (FASER). If new LLPs are light compared to the weak scale and very weakly coupled, the focus at the LHC on searches for new particles at high p_T




Figure 67. Top: CODEX-b reach for $B \rightarrow X_s \varphi$ in the $s_\theta^2 - m_\varphi$ plane. Solid (dotted-dashed) line assumes $\mathcal{L} = 300 \text{ fb}^{-1}$ ($\mathcal{L} = 1 \text{ ab}^{-1}$). Bottom: inclusive CODEX-b $B \rightarrow X_s \varphi$ reach (solid lines). The shaded regions (dashed lines) indicate current LHCb limits (300 fb^{-1} projection) from $B \rightarrow K(\varphi \rightarrow \mu\mu)$, rescaled to the inclusive process and assuming $\text{Br}[\varphi \rightarrow \mu\mu] \simeq 30\%$ and 10% for $m_\varphi = 0.5 \text{ GeV}$ and 1 GeV , respectively. The gray shaded area and the dashed line indicate the approximate current [454] and projected [455] limits, respectively, for Belle II, from $B \rightarrow K^{(*)}\nu\bar{\nu}$ precision measurements. From [104].

may be completely misguided. In contrast to TeV-scale particles, which are produced more or less isotropically, light particles with masses in the MeV–GeV range are dominantly produced at low p_T . In addition, because the new particles are extremely weakly coupled, very large SM event rates are required to discover the rare new physics events. These rates are not available at high p_T , but they are available at low p_T : at the 13 TeV LHC, the total inelastic pp scattering cross section is $\sigma_{\text{inel}}(13 \text{ TeV}) \approx 75 \text{ mb}$ [460, 461], with most of it in the very forward direction. This implies




Figure 68. Comparison of experimental sensitivity to the BSM decay of an SM-like Higgs to dark photons with and without the muon shadow (' μ Sh') for a few different existing and proposed detectors. The $\gamma_d \rightarrow \mu\mu$ branching ratio is taken from e^+e^- data [458]. The CODEX-b sensitivity here is the so-called 'optimistic' reach, with $\mathcal{L} = 1 \text{ ab}^{-1}$ and a larger volume, assuming DELPHI is removed. From [104].

$$N_{\text{inel}} \approx 2.3 \times 10^{16} (2.3 \times 10^{17}) \quad (5.5)$$

inelastic pp scattering events for an integrated luminosity of 300 fb^{-1} at the LHC (3 ab^{-1} at the HL-LHC). Even extremely weakly-coupled new particles may therefore be produced in sufficient numbers in the very forward region. Due to their weak coupling to the SM, such particles are typically long lived and travel a macroscopic distance before decaying back into SM particles. Moreover, such particles may be highly collimated. For example, new particles that are produced in pion (B -meson) decays are typically produced within angles of $\theta \sim \Lambda_{\text{QCD}}/E$ (m_B/E) off the beam-collision axis, where E is the energy of the particle. For $E \sim \text{TeV}$, this implies that even $\sim 500 \text{ m}$ downstream, such particles will have only spread out $\sim 10 \text{ cm}$ – 1 m in the transverse plane. A small and inexpensive detector placed in the very forward region may therefore be capable of extremely sensitive searches to LLPs, provided a suitable location can be found and the signal can be differentiated from the SM background.

Forward search experiment (FASER) [105, 462–465], the is an experiment designed to take advantage of this opportunity. It is a small detector, with volume $\sim 1 \text{ m}^3$, that is proposed to be placed along the beam-collision axis, several hundreds of meters downstream from the ATLAS or CMS IP. In the following, we present a promising location of FASER, discuss the properties of the signal and the required detector design, and present the new physics reach for representative models.

As shown in figure 69, FASER will be placed along the beam collision axis, several hundreds of meters downstream from the ATLAS or CMS IP after the LHC tunnel starts to curve. A particularly promising location is a few meters outside the main LHC tunnel, 480 m downstream from the ATLAS IP, in service tunnel TI12, as shown in the bottom panels of figure 69. (A symmetric location on the other side of ATLAS in tunnel TI18 is also possible.) This tunnel was formerly used to connect the SPS to the LEP tunnel, but is currently empty and unused. As shown on the tunnel map in the lower left panel of figure 69, the beam collision axis passes through TI12 close to where it merges with the main LHC tunnel. A




Figure 69. Proposed location for FASER in TI12. Top panel: a schematic drawing of the LHC and the very forward infrastructure downstream from the ATLAS and CMS interaction points; FASER is to be located 480 m from the IP, after the LHC ring starts to curve. Bottom panels: a map of the tunnel TI12 including the beam collision axis (left), a photo of this location (center), and a model of the experiment integrated in the TI12 tunnel (right).

more detailed study of the intersection between the beam collision axis and TI12 verifies that there exists space for FASER in the tunnel, as shown in the lower-right panel of figure 69.

In this location, FASER harnesses the enormous, previously completely unused and unexamined cross section for very forward physics ($\sigma \sim 100$ mb). This cross section implies that even very weakly coupled new particles can be produced in large numbers at the LHC. In addition, the production of LLPs at high center-of-mass energy results in long average propagation distances ($\bar{d} \sim \mathcal{O}(100)$ m) and decays that are far beyond the main LHC infrastructure in regions where the backgrounds are expected to be negligible.

FASER will search for LLPs that are produced at or close to the ATLAS IP in the very-forward direction, travel approximately 480 m, and then decay via $\text{LLP} \rightarrow \text{charged tracks} + X$. When LLPs are produced in the very-forward region of the beam collision axis, they typically have very high energies $E \sim \text{TeV}$. Although the identity of the LLP decay products depends on the mass of the LLP and the concrete new-physics model, some of the standard, characteristic LLP decay signatures are generically expected, such as two or more stable charged particles, such as electrons, muons or pions. This leads to a striking signature at FASER: two oppositely charged tracks with very high energy that emanate from a vertex inside the detector and which have a combined momentum that points back to the IP. A measurement of individual tracks with sufficient resolution and an identification of their charges is therefore imperative if the apparatus is to make use of kinematic features to distinguish signal from background. A tracking-based technology, supplemented by a magnet and a calorimeter to allow for energy measurements, will be the key components of FASER. Details of the detector design can be found in the Letter of Intent [462] and Technical Proposal [464].

The FASER signals consist of two extremely energetic ($\sim \text{TeV}$) coincident tracks or photons starting at a common vertex and pointing back to the ATLAS IP. Muons and neutrinos are the only known particles that can transport such energies through 100 m of rock and concrete between the IP and FASER. The CERN Sources, Targets, and Interactions group has computed the fluxes of muons and neutrinos at the FASER location using a FLUKA simulation [466–468]. These muon fluxes then allow one to estimate the rate and energy spectrum of muon-associated radiative processes near the detector. Preliminary estimates

show that muon-associated radiative processes and neutrino-induced backgrounds may be reduced to negligible levels [464, 464].

Emulsion detectors and battery-operated radiation monitors were installed in both TI12 and TI18 during Technical Stops in 2018. The results from these *in situ* measurements have validated the estimates of the FLUKA simulation, confirming that the high-energy particle background is highly suppressed and radiation levels are also very low and not expected to be problematic for detector electronics. Additional work is ongoing to refine background estimates, evaluate signal efficiencies, and optimize the detector.

In its first stage, FASER is an extremely compact detector, sensitive to decays in a cylindrical region of radius $R = 10$ cm and length $L = 1.5$ m. FASER is planned to be constructed and installed in LS 2 and will collect data during Run 3 of the 14 TeV LHC (2021–23). After FASER’s successful operation, FASER 2, a much larger successor with roughly $R \sim 1$ m and $L \sim 5$ m, could be constructed in LS 3 and collect data during the HL-LHC era (2026–35). More details on the FASER timeline can be found in the Letter of Intent [462] and Technical Proposal [464].

The sensitivity of FASER to detect LLPs has been studied in a plethora of new physics models, including dark photons [105], dark Higgs bosons [469], HNLs [297], axion-like particles (ALPs) [470], iDM [471], flavor-specific scalar mediators [472], $U(1)_{B-L}$ -gauge bosons [473], R -parity violating supersymmetry [60, 174] and SIMPs [474]. A summary on FASER’s physics reach for LLP can be found in [463]. Combined, these studies establish FASER’s significant impact on the LHC’s discovery reach for LLPs. The physics reach at FASER and FASER 2 for these models is shown in figure 70. Here we assume that backgrounds can be reduced to negligible levels. The gray-shaded regions of parameter space have already been excluded by previous experiments. For comparison we also show the projected reaches of other proposed experiments that search for LLPs.

Dark photons and ALPs (upper panels) are mainly produced in the decay of light mesons, via dark bremsstrahlung, or through the Primakoff process, and they are therefore very collimated around the beam collision axis. Already a very small detector is able to probe large and unconstrained regions of parameter space, making dark photons an ideal short-term goal for FASER. In contrast, dark Higgs bosons and HNLs (lower panels) define good long-term physics goals. They are both mainly produced in heavy meson decays, leading to a larger spread around the beam collision axis. A larger, but still relatively small, detector with $R \sim 1$ m is then required to exploit the full potential of FASER.

6. Reinterpretation and recommendations for the presentation of search results

Chapter editors: Giovanna Cottin, Nishita Desai, Sabine Kraml, Andre Lessa, Zhen Liu

Contributors: Juliette Alimena, Will Buttinger, Eric Conte, Yanou Cui, Jared Evans, Benjamin Fuks, Lukas Heinrich, Jan Heisig, Gavin Hesketh, David Michael Morse, Michael Ramsey-Musolf, Ennio Salvioni, Michele Selvaggi, Brian Shuve, Yuhsin Tsai

6.1. Introduction

Models and scenarios with LLPs have seen an enormous rise in interest in recent years. They include supersymmetric scenarios with almost mass-degenerate lightest states [63, 65], highly split spectra [113, 475], very weakly interacting LSPs like gravitinos or axinos [476, 477], or R -parity violation [55, 110, 111], as well as equivalent scenarios in other SM extensions (e.g. extra-dimensional models) with new SM gauge-charged particles. More recent ideas include models with feebly interacting DM [146] (supersymmetric or not), asymmetric DM [478],




Figure 70. Projected FASER exclusion reach for benchmark new-physics scenarios containing dark photons (top left), ALPs with dominantly photon couplings (top right), dark Higgs bosons (bottom left), and HNLs with dominantly τ -mixing (bottom right) in the corresponding coupling versus mass planes. The gray shaded regions are excluded by current experimental bounds, and the colored contours represent projected future sensitivities of other proposed experiments that search for LLPs. See [463] for details.

HV models [66], and other dark-sector models; for a comprehensive discussion, see the classification of existing well-motivated theories with LLPs in chapter 2.

All of these models can feature a large variety of possible LLP signatures. In HV models [66], for instance, new particles can either decay into invisible dark particles or back to the SM, thus possibly leading to a mix of long-lived and prompt signatures, with or without missing transverse energy (\cancel{E}_T). Furthermore, new theoretical frameworks are constantly emerging, often motivated by new approaches to the hierarchy problem or DM. It is therefore

of great interest to our community to be able to reinterpret the LLP experimental results for new models which may be developed in the future¹⁸².

The reinterpretation of experimental results can generically be done in two ways: by applying appropriate simplified-model results to more complete models, or by reproducing the experimental analysis in a MC simulation. Clearly, the former is easier and faster, while the latter is more generally applicable but also more difficult and much more time consuming¹⁸³.

In the context of searches for prompt signatures with Z_T , the use of simplified models has been shown to be a fruitful approach for both the experimental and theoretical communities [74, 76, 78, 87, 88, 480, 481]. Dedicated tools, notably SMODELS [482–484] and FASTSLIM [485], are publicly available and allow the user to reinterpret SUSY simplified-model results within the context of a full model. The coverage of a full model can, however, be severely limited by the kind of simplified-model results available, as discussed recently in [106] for the case of the phenomenological MSSM. Indeed, for some models the large number of relevant simplified-model topologies and their complexity can make the simplified-model approach inexpedient. In this case, a more complete and robust recasting procedure is necessary.

Again, for prompt signatures, a general recasting approach is available through several public tools, notably CHECKMATE [486, 487], MADANALYSIS5 [488, 489], RIVET [490] (v2.5 onwards) and Gambit’s COLLIDERBIT [491]. These tools allow the user to reproduce experimental analyses by means of MC event simulation coupled to an approximate emulation of detector effects¹⁸⁴. For the latter, CHECKMATE and MADANALYSIS5 rely on DELPHES [342], in some cases supplemented with appropriate tuning, while RIVET and COLLIDERBIT employ object smearing and analysis-specific reconstruction efficiencies. The ATLAS and CMS collaborations are helping these recasting efforts by providing more and more detailed information about their analyses and results. As an example, covariance matrices for the background correlations in the framework of simplified likelihoods were provided for the analysis of [492].

The situation is, so far, quite different for LLP searches. First, the presentation of results in terms of simplified models is still limited to a few topologies and does not always include all of the parameters required for a general purpose reinterpretation. It is worth noting that, compared with simplified models for prompt searches, simplified models for LLP searches always have at least one additional free model parameter: the lifetime of the LLP. Second, recasting LLP searches outside of the experimental collaborations is a difficult task since the searches are very sensitive to the detector response, which in most cases cannot be easily emulated by a fast detector simulation (and many backgrounds are also challenging to simulate due to their non-standard nature; see chapter 4 for a full discussion of backgrounds to LLP searches). As a result, while first steps towards the treatment of LLPs were achieved in SMODELS [493] and MADANALYSIS5 [494], the implementation of LLP searches in public recasting tools is still in its infancy, thus limiting the applicability of the experimental results.

¹⁸² In this context we refer the reader also to the activities of the ‘Forum on the Interpretation of the LHC Results for BSM studies’ [479].

¹⁸³ When no new backgrounds need to be considered and the hypothesized signal does not affect control regions, one can simply determine the event counts in the signal regions and compare them to the 95% C.L. observed limits, or take the numbers of observed events and expected backgrounds to compute a likelihood.

¹⁸⁴ For completeness it should be noted that, while all of these tools include a more-or-less extensive set of SUSY searches, many of the searches for other, ‘exotic’ types of new physics cannot yet be reproduced outside the experimental collaborations. This concerns in particular searches relying on multivariate techniques such as boosted decision trees.




Figure 71. A qualitative overview of the possibilities for the presentation of results discussed in this chapter. The axes represent the complexity of information required by each format and the corresponding level of model independence.

In order to allow for a more extensive reinterpretation or recasting of the experimental analyses, detailed information concerning the detector performance and object reconstruction is needed. These can in principle be provided in the format of efficiencies¹⁸⁵ for selection and reconstruction of relevant objects, as demonstrated already by some pioneering experimental publications [237, 495]. One difficulty in this respect is that the information needed for recasting LLP searches is clearly analysis dependent, which means an additional workload for the analysis groups to provide this information on a case-by-case basis.

The objective of this chapter is to discuss the presentation of LLP search results with the aim that they can be re-used for interpretations beyond the models considered in the experimental publications. To this end, we first discuss in section 6.2 the various options for presenting the LLP results (and for a summary, see figure 71), and compare their advantages and shortcomings. In section 6.3, we discuss in more details how the simplified models defined in section 2 can be used to reinterpret LLP searches. In section 6.4, we present several attempts at recasting LLP searches, according to the LLP signature, focusing on HSCPs and various displaced objects. For each case, we elaborate on the lessons learned from the reinterpretation effort. Section 6.5 presents a first attempt to extend the public detector simulator DELPHES to deal with LLP searches, while section 6.6 focuses on reinterpretations performed within the experiments themselves, including the RECAST framework [496]. In section 6.7, we discuss complementary constraints on LLPs from reinterpreting prompt searches. We conclude in section 6.8 with our recommendations for the presentation of LLP results.

¹⁸⁵ We employ the term ‘efficiency’ in a broad sense. It can refer to reconstruction efficiencies, selection efficiencies, overall signal efficiencies, etc, which will be further specified below.




Figure 72. Possibilities for the presentation of results: simplified-model efficiencies assuming a specific topology of LLP production and decay, signature efficiencies assuming only a specific LLP decay, and final-state object efficiencies which are independent of the specific decay mode.

6.2. Options for presenting experimental results

A qualitative view of the various possibilities for presentation of search results is illustrated in figure 71. We broadly classify these possibilities according to the type of information provided for reinterpretation (in other words, the type of efficiency). Each type refers to distinct signal objects, as illustrated in figure 72. As we can see, each possibility relies on distinct assumptions about the signal, resulting in different levels of model dependence. Below, we provide a brief discussion of the advantages and limitations of the various possibilities of presentation of LLP search results. We gradually progress from the simplest case towards more complexity but also better re-usability.

Simplified models. In most cases, the simplified-model topology corresponds to single or pair production of the LLPs, though in principle simplified models for production through cascade decay of heavier states can also be envisaged. In any case, the simplified model incorporates important assumptions on the LLP production mode, decay mode and quantum numbers. Simplified-model results can be presented at different levels of sophistication and re-usability:

- exclusion curves in, e.g. a mass-versus-mass or mass-versus-lifetime plane, are highly model dependent and can rarely be used for reinterpretation without significant effort;
- cross section upper limits¹⁸⁶ can be applied to a larger variety of models in which the same LLP production and decay mode occurs through a rescaling of the cross section times branching ratio factor;

¹⁸⁶ For the sake of re-usability, cross section upper limits in absolute terms are much preferred over limits on the signal strength.

- simplified-model efficiencies go one step further: they make it possible to combine different topology contributions to the same signal region and compute an approximate likelihood using the number of expected and observed events.

The main advantages of simplified models are a parametrization of the signal sensitivity in terms of few physical parameters and a unified language and format applicable to a wide range of searches. Also, when re-using simplified-model results one avoids detector simulation uncertainties. The disadvantages are that the simplified-model result cannot be applied to other LLP production or decay modes, resulting in overly conservative limits if the LLP signal is composed from multiple topologies. This can be important if the LLP is a color singlet, but there are several heavier color-charged states which can be produced and decay to the LLP. In principle this can be overcome if efficiencies are provided for a sufficiently large number of simplified models (including cascade decays), as a function of the simplified-model parameters, which should include the LLP lifetime. These efficiencies can then be combined in order to compute the corresponding constraints to complex models, where multiple topologies are present¹⁸⁷. We stress, however, that in order for this combination to be possible, signal efficiencies and not cross section upper limits must be provided. The major drawback of this approach is that in order for the results to be applicable to a broad class of models, the number of required simplified models and their complexity can easily become very large. For achieving a high level of model independence, it is therefore desirable that the experimental analysis can be recast with MC event simulation. Two ways of presenting results are useful to this end: *signature efficiencies* and *object efficiencies*.

Signature efficiencies. Signature efficiencies are efficiencies for the reconstruction of the main LLP signature treated as a single ‘object’ (single charged track, DV, DT, etc) as a function of the LLP kinematic parameters and the lifetime. Signal efficiencies require the assumption of a specific LLP decay mode, but are otherwise highly model independent since they make no assumption on the LLP production mode. In addition, they are fully model independent for stable particles (within the detector volume), since in this case no assumptions about the LLP decay mode are required at all. In many cases, however, the reconstruction efficiencies depend on multiple kinematic variables, such as the LLP p_T , its transverse decay position, impact parameter, etc. As illustrated in section 6.4.6, these efficiencies can be very useful for recasting LLP searches, but they are not often provided by the experimental collaborations. Many recasting efforts consist of extracting these efficiencies from the provided information, but this procedure can result in large uncertainties.

Object efficiencies. Object efficiencies are efficiencies for the reconstruction of the physics objects relevant for building the LLP signature; for example, if an analysis reconstructs DV out of tracks, the associated object efficiencies would be the efficiency of reconstructing each displaced track that is later combined to form the vertex. They can clearly be applied to a wide range of LLP decay and production modes, since no specific assumptions about either is required. For instance, a displaced lepton reconstruction efficiency can be provided as a function of the lepton p_T and its transverse impact parameter d_0 . As discussed in section 6.4.2, these efficiencies can be used to recast LLP searches to an acceptable accuracy ($\sim 20\%$). Furthermore, as illustrated in section 6.4.6, knowledge of object efficiencies are essential for a general purpose recasting of the search. Within this approach the model dependence is minimal and can be restricted to a few general assumptions about the nature of the LLP. Also, object efficiencies could be included in fast-detector simulators, thus

¹⁸⁷ These points are illustrated by the reinterpretation of the CMS search for HSCPs [495] discussed in section 6.3: for the specific model considered in [97], constraints obtained using only a single simplified model (direct production of $\tilde{\tau}$ s in this case) can underestimate the bounds on the LLP mass by almost a factor of two.

providing a way of recasting LLP searches on a similar footing as prompt searches. The main difficulty with providing such object efficiencies is the potentially large number of parameters required for their parametrization.

6.3. Reinterpretation using simplified models

One of the possibilities for extending the experimental results from LLP searches to a large variety of scenarios is through the use of simplified-model topologies. Simplified models (or simplified-model spectra, SMS) have been widely used for the interpretation of prompt and LLP searches. As discussed in chapter 2, a large number of SMS topologies are possible for the distinct LLP signatures, which can be grouped by the LLP production mode, decay and lifetime. These SMS topologies aim to capture the main physical properties of the LLP signal and can then be used to constrain other scenarios containing similar topologies. The use of simplified-model results to constrain full models has been shown to be possible [93, 106, 482, 485, 497, 498], even though it has its shortcomings [106]. Also within the context of LLP searches, the use of simplified model results for reinterpretation can be a good alternative, e.g. when a recasting based on a MC simulation is difficult or is too computationally expensive. In this section, we briefly review how SMS results can be used to reinterpret searches for full models as well as the particular challenges presented by LLP searches. A concrete example of reinterpretation using simplified models is given in section 6.3.2, based on the results of [97].

6.3.1. From simplified to full models. The interpretation of experimental results using simplified models typically corresponds to producing upper limits on the production cross section or signal efficiencies for a specific SMS topology (production and decay channel). These results are provided as a function of the simplified model parameters, which have been largely taken to be the masses of the BSM particles appearing in the topology. For LLP topologies, however, a new parameter must be considered: the LLP lifetime (see chapter 2). With the exception of searches for stable particles, the lifetime is one of the main parameters affecting the topology efficiency and upper limit.

Once signal efficiencies¹⁸⁸ (ϵ) are provided for one or more SMS topologies, these can be used, under some approximations, to quickly compute the number of expected signal events (S) for a full model:

$$S = \mathcal{L} \times \left(\sum_{\text{SMS}} \sigma_{\text{SMS}} \times \text{BR}_{\text{SMS}} \times \epsilon_{\text{SMS}} \right), \quad (6.1)$$

where \mathcal{L} is the luminosity for the respective search and the sum runs over simplified model topologies. Since the production cross section (σ_{SMS}) and branching ratios (BR_{SMS}) for each topology can be quickly computed for any full model, the simplified model signal efficiencies (ϵ_{SMS}) can be directly used to obtain the signal yield. This procedure does not rely on any MC simulation or recasting of LLP searches and can be easily applied to a wide variety of models, provided ϵ_{SMS} is known. The main limitation of this approach comes from the limited (although growing) number of SMS results available. Since ϵ_{SMS} is typically known only for very few simplified models, the sum in equation (6.1) is limited to the number of available topologies, resulting in an under-estimation of S .

¹⁸⁸ For simplicity we will refer to the signal acceptance times efficiency as ‘signal efficiency’. This efficiency is a function of the simplified model parameters, including the LLP lifetime.




Figure 73. Region of the CMSSM parameter space with long-lived NLSP staus. The light gray regions are excluded by the requirements $m_{\tilde{\tau}} \gtrsim m_{\tilde{\chi}_1^0}$ and $120 \text{ GeV} \leq m_h \leq 130 \text{ GeV}$. The top red region is excluded by the upper limit on the neutralino relic density, while the lower blue region is excluded by the CMS constraints on direct production of long-lived staus. For more details see [97].

For prompt SUSY searches, a systematic approach for reinterpreting simplified model results based on the procedure outlined above has been developed in [482, 485]. Furthermore, using the large number of available SUSY SMS results, public tools are available for constraining full models using these results [483, 485]. The same procedure can also be applied to LLP SMS results, as shown in [97] for the case of HSCPs and implemented recently in SMODELS [493]. In the next section we review some of the results found in [97]. Although these have been obtained within the context of HSCPs, the main results can be generalized to other LLP signatures, and demonstrate some of the advantages and shortcomings of reinterpretations using LLP simplified models.

6.3.2. Reinterpretation using HSCP simplified models. The CMS search for HSCPs in [495] provided signal efficiencies for the simplified model topology $pp \rightarrow \tilde{\tau}\tilde{\tau}$ as a function of the stau mass. In the language of the simplified models of section 2.4.2, this is the direct pair production mode of a charged LLP. The stau is assumed to be stable (at detector scales), thus producing a highly ionizing track, which can be used to search for this scenario (see section 3.5). Since the stau lifetime (τ) is assumed to be $\gg 10 \text{ ns}$, the signal efficiencies do not depend on τ , thus simplifying the SMS parameter space, which reduces to the stau mass¹⁸⁹. The relevant selection efficiencies required for a general purpose MC recasting of the HSCP search have also been provided by the CMS analysis; see section 6.4.1 for details.

The efficiencies for the stau simplified model can be used to constrain a full BSM scenario which contains HSCPs. In [97], the region of the CMSSM parameter space with $m_{\tilde{\tau}} - m_{\tilde{\chi}_1^0} < m_{\tau}$ has been considered, since it provides a possible solution to the Lithium problem [499, 500]. Due to the small mass difference, the stau is long-lived and decays outside the detector, thus generating a HSCP signal. In figure 73, we show the constraint on

¹⁸⁹ We point out that it is still possible to apply these simplified model results to models with smaller LLP lifetimes if we include the suppression factor from the LLP decay length distribution, as discussed in section 6.4.1.2.

Table 5. Definitions of the HSCP simplified models considered in this section. The symbol X represents the HSCP, Y_i represents intermediate BSM particles, SM represents any Standard Model particle and inv represents an invisible final state, such as the neutralino. The correspondance with the simplified models language of chapter 2 (section 2.3.1) is also given.

SMS topology	Notation in chapter 2
$pp \rightarrow X X$	DPP
$pp \rightarrow Y_1 Y_1, Y_1 \rightarrow \text{SM } X$	HP
$pp \rightarrow Y_1 Y_1, Y_1 \rightarrow \text{SM } Y_2, Y_2 \rightarrow \text{SM } X$	—
$pp \rightarrow Y_1 Y_2, Y_1 \rightarrow \text{SM } Y_2, Y_2 \rightarrow \text{SM } X$	—
$pp \rightarrow Y Y, Y \rightarrow \text{SM SM } X$	HP
$pp \rightarrow \text{inv } X$	CC
$pp \rightarrow Y_1 Y_2, Y_1 \rightarrow \text{SM inv}, Y_2 \rightarrow \text{SM } X$	—
$pp \rightarrow Y_1 Y_2, Y_1 \rightarrow \text{SM inv}, Y_2 \rightarrow Y_3 \text{ SM}, Y_3 \rightarrow \text{SM } X$	—

the CMSSM parameter space obtained using only the simplified model provided by CMS (direct stau production). Since the simplified model only contains one parameter, it translates to a limit on the stau mass ($m_{\tilde{\tau}} < 260$ GeV), as shown by the blue region in figure 73. In this CMSSM scenario, however, direct production of staus only contributes to a small fraction of the total HSCP signal, since staus are typically produced from cascade decays of heavier SUSY states, such as charginos, squarks and gluinos. (Note that these correspond to the HP modes of section 2.4.2.) Furthermore, there are several possible topologies which contain a stau and the LSP ($\tilde{\chi}_1^0$) in the final state, thus resulting in a mixed missing energy-HSCP signature. Therefore using only the CMS constraints for the direct stau production simplified model largely underestimates the sensitivity of the CMS search.

In order to improve the constraints shown in figure 73, one must have efficiencies for several SMS topologies. Fortunately, a MC recasting of the 8 TeV CMS search is possible (see section 6.4.1 for details) and can be used to compute simplified model efficiencies. In [97], seven additional simplified models containing cascade decays were considered and their efficiencies computed as a function of the masses appearing in the topology. A summary of the topologies considered are shown in table 5. It is important to point out that it is not necessary to specify the SM final states appearing in the simplified models, since the HSCP search is inclusive and the efficiencies do not depend on the additional event activity. Using this extended database of simplified model efficiencies and equation (6.1), we can compute a more inclusive signal yield for each point of the CMSSM parameter space and improve the constraints on the model. The results are shown in figure 74, where we see a drastic improvement in the region excluded by the constraints on HSCPs, as expected. For this specific scenario (with $\tan\beta = 10$), all the parameter space is excluded either by the CMS or DM constraints [97].

Figure 74 illustrates the feasibility of using simplified model efficiencies to constrain full models. This approach has the advantage of being computationally inexpensive (once the efficiencies are known) and can be used to quickly test a large number of model points. However the approach relies on a few approximations and is never fully inclusive, since the number of SMS topologies with published efficiencies is always limited. Hence, it is important to verify how close the simplified model reinterpretation comes to the full recasting using a MC simulation. In [97] it was shown that, within the CMSSM scenario discussed above and using eight simplified model topologies, the SMS results reproduce the full




Figure 74. Region of the CMSSM parameter space with long-lived NLSP staus, as in figure 73, but here additionally using all the simplified models listed in table 5.

simulation within 20% or better. Since this error is of the order of the uncertainties in recasting, the use of simplified models becomes a viable alternative to full recasting. The SMS reinterpretation is even more relevant for the cases where a straightforward MC recasting is not possible.

Lessons learned. The example discussed in section 6.3.2 illustrates how simplified models for LLPs can be used as a reinterpretation tool. One important point which becomes clear once we compare figures 73 and 74 is the importance of a sufficiently inclusive database of simplified model efficiencies. In particular, depending on the full BSM scenario considered, the minimal set of simplified models proposed in chapter 2 and the appendix may not be sufficient to allow for a reinterpretation based on simplified model results alone. The reason is that this set was derived as the minimal set to generate a collection of relatively inclusive signatures, but may not correctly model the efficiency of every UV theory leading to that signature. In these cases, results for additional simplified model topologies are necessary and can be provided by the experimental collaborations or generated by theory groups if a full recasting is available. Furthermore, since LLP searches can be very inclusive, the simplified models considered can also be defined inclusively, as discussed above. In this way a limited number of simplified models can cover a large number of event topologies, thus increasing the SMS coverage of full models.

6.4. Recasting examples for specific searches

Here we provide examples of recasting specific experimental searches for several LLP signatures: searches for HSCPs, displaced leptons, displaced jets, displaced LJs, non-pointing photons, and DV. These recasting attempts have been made outside the experimental collaborations, making use of the public information provided by the experimental note or publication. The aim here is to highlight the challenges faced when recasting LLP searches and also to highlight the cases where the experimental information provided is straightforward and useful for recasting.

6.4.1. Heavy stable charged particles. Searches for HSCPs are based on the signature of highly ionizing tracks and/or an anomalous TOF between the particle's production at the IP and its arrival in the muon detector [390] (see section 3.5.1.1 for more details). Both signatures are sensitive to the particle's velocity and exploit the production of HSCPs outside of the ultra-relativistic regime, allowing for a powerful discrimination against the highly boosted SM backgrounds. HSCP searches assume particles are sufficiently long-lived to traverse the entire detector. They have been interpreted for HSCPs that are purely electrically charged or carry color charge, the latter of which hadronize to form R -hadrons as they propagate through the detector [10]. Typically, the HSCP signature yields high sensitivities providing a very strong background rejection while still allowing for large signal efficiencies. As a consequence, search strategies for new physics models with HSCPs typically do not benefit from more model-dependent selection criteria, like requiring additional particles in the event [95]. The corresponding searches can, hence, be performed in a mostly inclusive manner concentrating on the HSCP candidate itself. This fact opens up the possibility of providing a widely applicable recasting based on signature efficiencies. This approach has been followed by the CMS Collaboration [495], which has provided probabilities for HSCP candidates to pass the on- and off-line selection criteria for Run 1 as a function of the relevant kinematical parameters.

In this section we describe the recasting of the 8 TeV CMS search for HSCPs and discuss its validation and applicability. Furthermore, we comment on the attempt to extrapolate the 8 TeV signature efficiencies to the corresponding 13 TeV analysis, for which the corresponding efficiencies have not been provided by CMS.

6.4.1.1. Recasting using signature efficiencies. Reference [495] provides efficiencies for the reconstruction and selection of HSCP candidates with $|Q| = 1$ in the form of on- and off-line probabilities, $P_{\text{on}}(\mathbf{k})$ and $P_{\text{off}}(\mathbf{k})$. These are given as a function of the truth-level kinematic variables velocity (β), pseudo-rapidity (η) and transverse momentum (p_T) of isolated HSCP candidates, so the vector \mathbf{k} is defined as: $\mathbf{k} = (\beta, \eta, p_T)$. The on- and off-line probabilities must be applied to isolated HSCP candidates, which are required to fulfill

$$\left(\sum_i^{\Delta R < 0.3} p_T^i \right) < 50 \text{ GeV}, \quad \left(\sum_i^{\Delta R < 0.3} \frac{E^i}{|\mathbf{p}|} \right) < 0.3, \quad (6.2)$$

where the sums include all charged and visible particles, respectively, within a radius of $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} < 0.3$ around the HSCP candidate track, p_T^i denotes their transverse momenta and E^i their energy. Muons are not counted as visible particles and the HSCP itself is not included in either sum.

If an event contains one or more HSCPs satisfying the above isolation criteria, the efficiency for the event to pass the analysis selection is given by:

$$\epsilon = \epsilon_{\text{on}} \times \epsilon_{\text{off}}. \quad (6.3)$$

For an event with one HSCP candidate $\epsilon_{\text{on/off}}$ is directly given by the signature efficiencies $P_{\text{on/off}}(\mathbf{k})$. For an event with two candidates, the efficiency reads [495]

$$\epsilon_{\text{on/off}} = P_{\text{on/off}}(\mathbf{k}^1) + P_{\text{on/off}}(\mathbf{k}^2) - P_{\text{on/off}}(\mathbf{k}^1)P_{\text{on/off}}(\mathbf{k}^2), \quad (6.4)$$

where $\mathbf{k}^{1,2}$ are the kinematical vectors of the two HSCPs in the given event. Therefore the on- and off-line probabilities combined with the isolation criteria allow for the complete recasting of the HSCP search using only truth-level events generated in MC.




Figure 75. Left: signal efficiency ϵ of the HSCP search. Right: 95% C.L. cross section upper limit for the GMSB model as the function of the stau mass. We compare the CMS analysis [495] from the full detector simulation (red solid lines) with the recasting using signature efficiencies (blue dashed lines). In the lower frames we show the respective ratios $\epsilon^{\text{Full}}/\epsilon^{\text{Recast}}$, $\sigma_{\text{limit}}^{\text{Full}}/\sigma_{\text{limit}}^{\text{Recast}}$. Taken from [97].

The recasting of the 8 TeV search was performed in [97] using the procedure described above. Events were simulated using PYTHIA 6 [501] and the total signal efficiency for a given model was then computed using:

$$\epsilon = \frac{1}{N} \sum_{i=1}^N \epsilon_i,$$

where the sum runs over all the (N) generated events and ϵ_i is the efficiency for each event computed using equation (6.3). Since the probabilities $P_{\text{on/off}}(\mathbf{k})$ are given for four distinct cuts on the reconstructed HSCP mass (m_{rec}), these were considered as four different signal regions. The number of observed events and the expected background for which of these cuts are reported in [495].

6.4.1.2. Validation and applicability. A validation of the method described above was done in [97] using the same GMSB model considered by CMS [495]. This supersymmetric model features a gravitino and a long-lived stau as the lightest and NLSP, respectively. Since the stau only decays outside the detector volume, all cascade decays of the produced sparticles terminate in the lightest stau, which provides the HSCP signature. The left pane of figure 75 compares the resulting signal efficiency obtained by the recasting and the full CMS detector simulation. The signal efficiencies agree within 3%, demonstrating that the recast is an excellent approximation to the full CMS simulation. The differences are of the order of the statistical uncertainties from the MC simulation of the signal. In the right pane of figure 75, we show the 95% C.L. upper limits on the inclusive production cross sections, which, again, agree (within $\sim 3\%$) with the ones obtained by the full simulation in [495]. Note that both limits are based on the discrete mass cuts on m_{rec} mentioned above. In the full CMS analysis [306], an event-based mass cut is used, resulting in slightly stronger constraints for some HSCP masses.

Due to the inclusive nature of the search, the above recasting provides a widely applicable and highly reliable way to reinterpret the HSCP search for arbitrary models

containing detector-stable HSCPs. Accordingly it has been used in a variety of phenomenological studies. For instance, it has been used for reinterpretations of supersymmetric models [96, 205, 502, 503] and non-supersymmetric models of very weakly interacting DM [142, 148]. In [96, 142], the recasting has been used to reinterpret the HSCP search for finite lifetimes by convolving the signature efficiency with the fraction of HSCPs that decay after traversing the relevant parts of the detector. The recasting has also been used for a reinterpretation in terms of simplified models, as discussed in section 6.3.2.

6.4.1.3. Extrapolation to 13 TeV. While the CMS search for HSCPs at 8 TeV has provided the signature efficiencies discussed above, the same is not true for the 13 TeV analysis [4]. Therefore a straightforward recasting of the Run 2 search is not possible. Nonetheless, since the 8 TeV CMS search has proven to be extremely useful in constraining models with long-lived charged particles, it would be desirable to recast the 13 TeV analysis as well. In the following we discuss an attempt [504] to obtain a similar recasting for the corresponding HSCP search at 13 TeV. Our aim is to extrapolate the public 8 TeV efficiencies to Run 2 by introducing a correction function F that accounts for the differences between both runs:

$$P_{\text{off}}^{13 \text{ TeV}}(\mathbf{k}) = F(\beta) \times P_{\text{off}}^{8 \text{ TeV}}(\mathbf{k}), \quad (6.5)$$

where we have assumed that the correction function is mainly dependent on the HSCP velocity. If $F(\beta)$ is sufficient to account for the difference between both runs and can be computed, we can directly obtain $P_{\text{off}}^{13 \text{ TeV}}$ and, using the procedure described in section 6.4.1.1, recast the 13 TeV analysis.

In order to compute the correction function $F(\beta)$, we use the total signal efficiencies reported by the 13 TeV CMS analysis [4] for direct production of long-lived staus. Since the signal efficiencies have been provided for six distinct values of the stau mass, we perform a fit of F to the efficiencies reported. We chose to parametrize the correction function $F(\beta)$ by eight parameters (C_i). Using MADGRAPH5_AMC@NLO [226] and PYTHIA 6 [501], we obtain generator-level events for each of the stau mass points at 13 TeV. Then, comparing the total signal efficiencies obtained for a given set C_i to the efficiencies reported in [4], we can determine the best-fit values for the C_i parameters and consequently the best-fit for the correction function ($F_{\text{best-fit}}$). The result of the best-fit function and its 1σ uncertainty is shown in figure 76. The deviation of F from 1 implies a decrease or increase of the respective detector and signal efficiency between the 8 and 13 TeV analyses. Figure 76 also shows that the function is loosely constrained for low values of β .

In order to verify the validity of the extrapolation to 13 TeV, we use $F_{\text{best-fit}}$ and equation (6.5) to compute the total signal efficiencies for the same six benchmark points used in the fit. A comparison between the results obtained through recasting and the efficiencies reported by CMS is shown by the second and third columns in table 6. The results reproduce the CMS values well within the expected uncertainties, thus validating the fitting procedure. Furthermore, the inclusion of the correction function significantly improves the agreement with respect to the direct extrapolation of the 8 TeV efficiencies ($F = 1$), as shown by the fourth column in table 6.

The high level of agreement obtained with $F_{\text{best-fit}}$ for the direct stau benchmark points is expected, since these points were used in order to fit the correction function. Therefore an independent test of the above fit must be performed in order to properly validate the recasting of the 13 TeV analysis. Fortunately CMS has also reported efficiencies for a second scenario, the GMSB model with long lived staus. This scenario not only contains direct stau production, but also includes production through cascade decays of heavier sparticles. Since




Figure 76. Best-fit correction function $F(\beta)$, used for recasting the CMS HSCP search (described in the text), and its $\pm 1\sigma$ band.

Table 6. Efficiencies for the 13 TeV LHC for the six benchmark masses in the direct stau production scenario reported by CMS (second column) and obtained through recasting using $F_{\text{best-fit}}$ and without the inclusion of the correction function ($F = 1$).

m_{HSCP} (GeV)	Direct production		
	$\epsilon(\text{CMS})$	$\epsilon(F_{\text{best-fit}})$	$\epsilon(F = 1)$
200	0.235	0.232	0.259
308	0.294	0.298	0.346
494	0.387	0.384	0.452
651	0.450	0.448	0.503
1029	0.497	0.501	0.466
1599	0.428	0.429	0.225

the GMSB model produces distinct event topologies, it provides a good test for the validity of the recasting procedure.

The results for the six GMSB benchmark points considered in [4] are shown in table 7. As we can see, they deviate from the CMS values by up to 20% for large stau masses, where our estimate undershoots the CMS efficiencies. Although the overall agreement is improved by the correction function, the result is not entirely satisfactory, given that the uncertainties for the 8 TeV recasting were under 5% (see figure 75). The observed difference might arise from several shortcomings in our description.

In particular, we assume F to only depend on β whereas the full probability maps are parametrized in the three kinematic variables β , η and p_T . However, assuming a dependence of all three kinematic variables is clearly not feasible given the very limited amount of information provided by the 13 TeV CMS analysis. Therefore we conclude that it is not possible to extrapolate the 8 TeV efficiencies in a straightforward way without additional information from the experimental collaboration.

Table 7. Efficiencies for the 13 TeV LHC for the six GMSB model benchmark points reported by CMS (second column) and obtained through recasting using $F_{\text{best-fit}}$ and without the inclusion of the correction function ($F = 1$).

m_{HSCP} (GeV)	GMSB		
	$\epsilon(\text{CMS})$	$\epsilon(F_{\text{best-fit}})$	$\epsilon(F = 1)$
200	0.276	0.297	0.279
308	0.429	0.401	0.423
494	0.569	0.494	0.556
651	0.628	0.524	0.580
1029	0.665	0.538	0.493
1599	0.481	0.442	0.228

Lessons learned. The prominent signature of HSCPs allows for a mostly inclusive search strategy concentrating on the HSCP track itself. Hence, searches for HSCPs can be reinterpreted using signature efficiencies in a widely applicable and highly reliable way. This possibility has been followed by the CMS Collaboration providing signature efficiency maps for the 8 TeV LHC. The validation reveals an excellent performance. The recast has been successfully used in the literature.

The signature efficiencies for 8 TeV can also be used to estimate the ones for Run 2 by applying a multiplicative correction function. While such an extrapolation introduces some level of approximation, a better knowledge of the underlying changes between both runs might reduce the uncertainties.

6.4.2. Displaced leptons. Searching for displaced leptons by requiring the leptons to have large impact parameters with respect to the PV is a very clean strategy due to the low backgrounds, and such searches are usually very straightforward to recast. The CMS displaced $e\mu$ search [275] demands two oppositely charged, different flavour (e, μ) leptons with large impact parameters (see also section 3.2). The recast is fairly straightforward to do, and the biggest difficulty in doing so is locating all of the relevant information as it is not all provided within the main document. The ‘standard’ isolation requirements used in the search can be found in an earlier version of the search [505]. The necessary cuts on the displaced decay position (v_T, v_Z) as well as the selection (as a function of p_T), reconstruction (as a function of impact parameter, d_0) and trigger efficiencies can be found on an additional website [506] containing auxiliary information for recasting. Although all of this information is excellent and greatly facilitates recasting the search, it is a challenge to collect the relevant information due to the fact that the additional material is not referenced in the document.

The benchmark model used in this search is the direct pair production of stops that decay through small lepton-flavor-universal RPV $\lambda'_{ijk} L_i Q_j D_k^c$ couplings ($\lambda'_{133} = \lambda'_{233} = \lambda'_{333}$) to yield displaced $\tilde{t} \rightarrow eb, \mu b$, and τb decays. The signal is simple to generate, where the only challenge is in handling the displacement properly. The most identifying pre-selection requirement of this search is that the transverse impact parameter, $|d_0|$, is required to be larger than $100 \mu\text{m}$ for both the electron and muon. The impact parameter is not the point where the parent object (e.g. τ or b) decays, i.e. the v mentioned above, but rather the distance to the point of closest approach of the lepton’s track relative to the center of the beampipe in the transverse plane. Backgrounds in this search from $Z \rightarrow \tau\tau$ or heavy flavor tend to result in leptons that are nearly collinear with the parent due to the small mass-to-momentum ratio, and yield a small impact parameter even for decays well on the lifetime tail of the parent. Events

Table 8. Left: preselection cuts in [275] (see also [506]). Right: the transverse impact parameter bins that define the exclusive signal regions. Table and figure taken from [205].





Figure 77. Validation of the CMS displaced $e\mu$ search [275] for the displaced supersymmetry benchmark model [507]. Figure taken from [205].

are binned across three exclusive signal regions: SR3, where both leptons have transverse impact parameters $|d_0|$ between 0.1 and 2.0 cm; SR2, with $|d_0|$ of both leptons between 0.05 and 2.0 cm, but not satisfying the requirement of SR3; and SR1, with $|d_0|$ between 0.02 and 2.0 cm, but not within SR2 or SR3. All selection requirements are summarized in table 8.




Figure 78. Number of expected events in the signal regions defined based on $|d_0|$ for the CMS displaced $e\mu$ search. The green points refer to the expected signal published by the analysis. Left: validation for 8 TeV analysis. Production cross section assumed NLO+NLL value 85.6 fb for $M_{\tilde{t}_1} = 500$ GeV, BR = 0.33 in each ℓ -channel. Right: validation for 13 TeV analysis. Production cross section assumed NLO+NLL value 67.0 fb for $M_{\tilde{t}_1} = 700$ GeV. The 13 TeV numbers are made using efficiency maps published for the 8 TeV search, as the 13 TeV maps are not yet public. Figures taken from [504].

In figure 77, we present the validation of the CMS displaced $e\mu$ search [275] from the study performed in [205]. For this search, we show the recommended 25% modeling uncertainty¹⁹⁰. The recast agrees very well with the results from the CMS displaced $e\mu$ search in the region of highest sensitivity, $300 \mu\text{m} \lesssim c\tau \lesssim 50 \text{ cm}$, but exhibits a moderate deviation on the tails. As this extremely low efficiency region is overly sensitive to the tails of kinematic distributions, it may be the case that the sensitivity is slightly under-estimated for lifetimes near 1 m or $100 \mu\text{m}$, but this discrepancy typically has no qualitative impact on any application of the results.

6.4.2.1. Extrapolation to 13 TeV. We now show another reinterpretation example of the CMS displaced $e\mu$ in order to highlight the comparison between 8 TeV [275] and 13 TeV [197] analyses. We compare in figure 78 our reproduction of expected signal events with the published validation material for the 8 TeV version, and the partially-available validation material for the 13 TeV search. Information on efficiency maps from the 8 TeV analysis was needed to obtain an extrapolation to 13 TeV, as the 13 TeV efficiency maps are not yet public. As we can see, the 8 TeV recast for the CMS displaced lepton search [275] agrees very well in the region of highest sensitivity. The 13 TeV recasting, however, under-estimates the CMS values by a factor of two or more. This is likely due to the fact that the lepton efficiencies can not be directly extrapolated from 8 to 13 TeV, as assumed in making figure 78. Also, with the absence of a cut-flow table, it is impossible to verify where the mis-match arises, whether it is due to mis-modeling of the signal region cuts or due to genuine changes in the efficiencies.

Lessons learned. The selection and trigger efficiencies provided by CMS are very useful for recasting the 8 TeV CMS search for displaced leptons [275] and allow for a very good level of agreement. The main challenge, however, consisted in collecting all the available

¹⁹⁰ A similar validation was done in [96] with details provided in its appendix D. After applying a flat 80% efficiency, a 20% modeling uncertainty is found to be appropriate, as shown in figure 14 of the reference.

information, which was not provided by the main document in [275]. Furthermore, the corresponding information for the 13 TeV search is not publicly available and an extrapolation of the 8 TeV efficiencies was shown to be inadequate.

6.4.3. Displaced jets. Searches for displaced jets are less straightforward to reinterpret than displaced leptons. Interest in accurate reinterpretation is increasing, as many new physics models give rise to this particular signature. The CMS search for displaced di-jets [508] was reinterpreted in [26] to explore LLP signatures for certain weak-scale models of baryogenesis [24, 26, 509], as well as a study [96] to understand current limits on long-lived signatures in supersymmetry.

The CMS search [26] uses a multivariate discriminant composed of observables that are challenging to model in MC, such as the track-multiplicity of the vertex and the root mean square of a cluster track-multiplicity variable. The reinterpretation approach in [26, 96] was to construct track information at truth level based on the output of a parton shower program (such as PYTHIA 8), and then use the truth-level information to construct the various vertex, cluster, and track-level observables for each event. As it is difficult to adequately account for inefficiencies of track and vertex reconstruction, the efficiency of passing the cuts with truth-level observables was considered and then it was normalized to the results from CMS. To do so, the authors of [26] simulated identical signal models to those with efficiencies reported by the CMS collaboration, assumed that the MC truth-level reconstruction gave an adequate description of kinematics but *not* track and vertex reconstruction, and so computed a ratio of truth-level efficiencies to those reported by CMS. The resulting efficiency ratios were used to re-scale the truth-level results of other models, leading to a reinterpretation of the CMS search for different models beyond the ones they considered. The details can be found in [26]; a more sophisticated approach in which object efficiencies were estimated and applied to tracks and DV in [96].

To validate this approach, truth-level quantities for the models constrained in [508] were computed and compared to the numbers and distributions reported. For example, comparisons of the distributions of the observables going into the multivariate reinterpretation discriminant, as well as the output of the multivariate discriminant itself, could be performed. While the truth-level distributions disagreed with those of CMS for individual observables, the actual multivariate discriminant output agreed with that of CMS at better than 25%. The ratio of truth-level efficiencies to CMS efficiencies are also compared for different LLP masses and kinematics, and these typically agree with one another at the factor-of-two level [26]. This suggests that this reinterpretation of the CMS results in terms of cross-section limits is likely accurate to the factor-of-two level.

Lessons learned. We find that a rather naïve truth-level reconstruction of the event could give a reinterpretation of cross-section limits to agree within a factor of two, provided the efficiencies were normalized to the experimental values using an overlapping set of signal models. One of the major obstructions to improving the accuracy of the estimate was the model dependence observed among the ratios of efficiencies between truth-level information and CMS. For instance, it was found that highly boosted models showed a much lower relative reconstruction rate in data versus truth-level MC than less-boosted LLPs. Since pair production of LLPs was considered near threshold in [26], this degradation in the performance of highly boosted LLPs does not greatly affect the confidence in the final result. However, it does suggest that characterizing the effects of the particle boost are important for reinterpretation.

In addition, with a larger and more diverse set of signal benchmarks available, the prospects for the reinterpretation of search results are better. The reasons are twofold:

- Increasing the number of presented signal models by the collaboration allows for more cross-checks between MC and the results in data. This allows for more sophisticated tuning of efficiencies as applied to truth-level events.
- Having a more diverse set of benchmark signal models means that it is easier to disentangle various kinematic effects on the efficiency (such as the LLP mass, boost, etc) and find a signal benchmark that most closely matches the model for which one wants to derive sensitivity.

6.4.4. Displaced LJs. A variety of scenarios predict the existence of LLPs decaying to a pair of highly collimated leptons, also known as LJs [127]. Models giving rise to LJ signatures include heavy RHNs, exotic Higgs decays, and dark gauge bosons. The relevant signature is one or more LJs emerging from a DV. In many cases, there can also be associated prompt objects, such as a prompt lepton produced in conjunction with the RHN (corresponding to CC production in the simplified model framework of section 2.4).

Existing searches. Current search results are outlined in section 3.2. The search in [273] was interpreted in the framework of the Falkowski–Ruderman–Volansky–Zupan (FRVZ) model [155] for the Higgs boson interacting with a hidden sector containing a massive dark photon (γ_d), massive neutralinos, and three hidden scalars. Displaced LJs are produced at the end of Higgs cascade decay that also yields two hidden light stable particles (HLSPs). Depending on the hidden-sector spectrum, the cascade decay may yield two or more γ_d that each decay to pairs of SM charged particles via kinetic mixing with the hypercharge gauge boson. The γ_d decay products are highly collimated. For $m_{\gamma_d} < 500$ MeV, LJs are the dominant decay mode, while for larger dark photon masses, displaced hadronic jets can also be significant.

Results are presented as limits on $\sigma \times \text{BR}(H \rightarrow n\gamma_d + X)$ ($n = 2, 4$) as a function of the $\gamma_d c\tau$. The strongest limits from the 8 TeV dataset arise from events that require at least one LJ with muons. For $n = 2$, a $\sigma \times \text{BR}(H \rightarrow 2\gamma_d + X)$ of $\gtrsim 1$ pb is excluded for $c\tau \sim 50$ mm.

Recast: dark photon with non-abelian kinetic mixing. The FRVZ model on which the ATLAS analysis was based implies the presence of at least two displaced LJs in the final state as well as two additional unobserved HLSPs that correspond to missing energy. The ATLAS analysis did not impose any \not{E}_T cuts. The presence of the two HLSPs affects the kinematics and topology of the signal event but does not explicitly enter the event selection or reconstruction. Thus, one should be able to reinterpret the ATLAS analysis in terms of any other model containing two or more displaced LJ, along with potentially additional, unobserved objects.

We consider a scenario for a dark photon X_μ that mixes with the neutral SM $SU(2)_L$ gauge boson W_μ^3 via a higher dimensional operator. Two possibilities have recently been considered: a dimension six operator involving the SM Higgs doublet, the $SU(2)_L$ field strength $W_{\mu\nu}^a$, and the corresponding $U(1)'$ field strength $X_{\mu\nu}$ [511]; and a dimension five operator involving $W_{\mu\nu}^a$, $X_{\mu\nu}$, and a real scalar triplet $\Sigma \sim (1, 3, 0)$ [512]. We consider the latter since it can yield an event topology similar to that of the FRVZ model and since it is the leading operator that may generate non-abelian kinetic mixing (NAKM) of the $U(1)'$ gauge boson with the SM. We will henceforth refer to this model as the NAKM5 scenario. The corresponding operator is

$$\mathcal{O}_{WX}^{(5)} = -\frac{\beta}{\Lambda} \text{Tr}(W_{\mu\nu} \Sigma) X^{\mu\nu}, \quad (6.6)$$




Figure 79. Constraints on the NAKM5 scenario, recast from the ATLAS search reported in [272]. The left panel gives the exclusion in the $(c\tau, \sigma \times \text{BR})$ plane, where the region above the parabola is excluded. The diagonal lines indicate the dependence of $\sigma \times \text{BR}$ on $c\tau$ for different representative choices of v_Σ . The right panel gives the exclusion region in the $(v_\Sigma, \Lambda/\beta)$ plane for $m_X = 0.4 \text{ GeV}$ (red region) and $m_X = 1.5 \text{ GeV}$ (yellow region).

where Λ is the associated effective field theory mass scale and β is a dimensionless coupling that is nominally $\mathcal{O}(1)$. When the neutral component of the Σ obtains a vacuum expectation value (vev) v_Σ , one has the ratio of the dark photon and SM photon couplings

$$\epsilon = \beta \sin \theta_W \left(\frac{v_\Sigma}{\Lambda} \right). \quad (6.7)$$

Note that the ρ -parameter constrains v_Σ to be smaller than about 4 GeV. Consequently, for Λ of order one TeV, ϵ is small (and consistent with present experimental constraints) without requiring the presence of a suppressed coupling in the Lagrangian. This feature distinguishes the dimension-five non-abelian kinetic mixing from the dimension-four kinetic mixing between X_μ and the SM hypercharge gauge boson.

The final state with two γ_d (resulting from the $X_\mu - W_\mu^3$ mixing) can be produced in one of two ways: (a) electroweak Drell-Yan pair production of triplet scalars that subsequently decay to a γ_d and SM gauge boson via the operator in equation (6.6); (b) production of the γ_d and a triplet scalar via $\mathcal{O}_{WX}^{(5)}$ with a subsequent decay of the triplet to a second γ_d plus a SM gauge boson via the same operator. In each case, one would expect the presence of two γ_d plus one or more unobserved massive prompt bosons in the final state. The ATLAS DV plus LJ analysis can be recast in a straightforward way for these event topologies, as no information about the unobserved prompt object has been used. It is worth emphasizing that both the production processes as well as the Σ decay rate are independent of the triplet vev, v_Σ . The latter only enters ϵ and, thus, only affects the dark photon $c\tau$.

The corresponding implications of the ATLAS 8 TeV results are indicated in figure 79. The first panel shows the limits on the $\sigma \times \text{Br}$ as a function of $c\tau$. The corresponding model sensitivity is shown for different choices of v_Σ by the diagonal lines. Note that for fixed v_Σ both the $\sigma \times \text{Br}$ and $c\tau$ depend on the operator coefficient β/Λ , leading to a non-trivial relationship between the two experimental quantities. This situation differs from the FRVZ

model, where $\sigma \times \text{Br}$ is independent of $c\tau$ since the mixing parameter ϵ does not depend on the parameters governing production of the hidden sector particles or their cascade decays. The intersections of the model lines with the ATLAS limits can then be translated into bounds on Λ/β as a function of v_Σ for different choices of the dark photon mass (denoted here as m_X), as shown in the second panel of figure 79. For a 1.5 GeV dark photon, the excluded region reaches 600 GeV for $v_\Sigma = 1 \text{ MeV}$.

The aforementioned recast does not require detailed information on event topology, other than the dark photon decay length. Consequently, the 13 TeV limits [273] translate rather straightforwardly into stronger bounds on the model parameter space, reaching to somewhat larger Λ/β .

Lessons learned. The triggering requirements used in the ATLAS analysis thus far limit the reach of displaced LJ searches. For models having a signal with only the displaced LJ and no other observable objects, such as the FRVZ model, triggering solely on MS tracks not associated with ID tracks is appropriate, though even here the 3mu6 trigger may not be sufficiently inclusive, as it requires at least three ROIs in the MS (see section 3.2.1). Events with LJ pairs for which neither LJ can be resolved into two separate ROIs will be missed.

It is clear that triggering on associated prompt objects, such as the lepton from one of the final state vector bosons in the NAKM5 scenario or from the W boson in heavy RHN models (CC LLP production), could significantly enhance the triggering efficiency and extend the reach of displaced LJ searches to a wider class of models and to a broader range of model parameter space. Inclusion of an associated prompt object in triggering may also enhance background rejection and relax the requirement on $\Delta\phi$ between LJs.

In addition, the implications of the mass scale of the intermediate BSM particles and the number of final state prompt objects remains to be investigated. The ATLAS 8 TeV search assumed the hidden sector particles are light compared to the mass of the SM Higgs boson, whose decay chain leads to the final states involving multiple γ_d and HLSPs. For the NAKM5 scenario and models with heavy RHNs, these assumptions about mass hierarchy may not apply. It is also not clear what impact the DV LJ isolation requirements have when there is an associated prompt lepton in the signal event.

6.4.5. Non-pointing photons. The search for non-pointing photons produced in association with missing transverse energy (\cancel{E}_T) [301] plays an important role in probing BSM particles that decay to a SM photon and an invisible particle through a highly suppressed coupling. Besides the GMSB models [513], which were the main motivation for the non-pointing photon search, this type of signal can also appear in many hidden-sector models. For example, in the dipole-mediated DM model (the ‘Dark Penguin’) [156], the production of two heavier dark fermions $pp \rightarrow Z^*/\gamma^* \rightarrow \chi_h \bar{\chi}_h$ is followed by the decays $\chi_h \rightarrow \chi_l + \gamma$. If the flavor structures of the DM mass and coupling are aligned, χ_h can be long-lived and give rise to non-pointing photons. Another example is provided by the dark-shower scenario [514, 515] that explains the galactic center gamma-ray excess. In this model, many hidden pions can be produced in the same LHC event. Some of these have displaced decays to a pair of SM photons, while others decay outside of the detector, yielding \cancel{E}_T . Notice that in this case the topology of the events is different from the previous examples, as the non-pointing photons and \cancel{E}_T originate from separate particles (and for more discussion on dark showers, see chapter 7).

Here, we describe a method used to recast the bounds of [301] to a BSM scenario that is different from the GMSB model, using the Dark Penguin signal [156] as an example.




Figure 80. Left: the Δz_γ distribution of the non-pointing photon signals measured by ATLAS. The background reported by ATLAS (blue histogram) was obtained from a data-driven analysis, using diphoton events with $E_T < 20$ GeV. Also shown, stacked on top of the background (red histogram), is the signal distribution from the dipole-mediated DM scenario with $(m_{\chi_h}, m_{\chi_l}, M) = (300, 10, 300)$ GeV, $\lambda\sqrt{N\bar{Y}} = 6$, and $\varepsilon = 10^{-7}$. See [156] for more details. Right: the geometry of the displaced signals.

6.4.5.1. Calculation of the signal efficiency for the non-pointing photon search. We follow the non-pointing photon analysis in [301], performed by the ATLAS collaboration on about 20 fb^{-1} of 8 TeV data. In [301] delayed photons were also considered, but here we focus only on the measurement of the Δz_γ of non-pointing photons (see figure 80). For DM signals given by the long-lived $\chi_h \rightarrow \chi_l \gamma$ decay, Δz_γ can be related to the χ_h decay length ℓ_d in the lab frame:

$$\begin{aligned} \Delta z_\gamma &= \ell_d \left(\hat{r}_{\chi_h, z} - \frac{\hat{r}_{\chi_h, T} \cdot \hat{r}_{\gamma, T} \hat{r}_{\gamma, z}}{1 - (\hat{r}_{\gamma, z})^2} \right) \\ &= \ell_d [\cos \theta_{\chi_h} - \cos(\phi_{\chi_h} - \phi_\gamma) \cot \theta_\gamma \sin \theta_{\chi_h}], \end{aligned} \quad (6.8)$$

where $\hat{r}_{T,z}$ represent the transverse and longitudinal components of the unit vector \hat{r} , respectively, as shown in the right panel of figure 80. To obtain the Δz_γ distribution of the DM decay, we first simulate the prompt process, $p p \rightarrow \chi_h \bar{\chi}_h$, $\chi_h \rightarrow \chi_l \gamma$, $\bar{\chi}_h \rightarrow \bar{\chi}_l \gamma$ in MadGraph5, then we apply the cuts performed in the ATLAS analysis, and finally reweight the events using the dark penguin form factors. Then we calculate the proper lifetime of χ_h and boost it to the lab frame using the momenta of each parton-level event. The angular information of the photon and χ_h allow us to calculate Δz_γ in equation (6.8) as a function of the decay length. Using this, each simulated MC event contributes to the differential cross section in Δz_γ as

$$\frac{d\sigma_{\text{displaced}}}{d\Delta z_\gamma} = \sigma_{\text{prompt}} \frac{dP}{d\Delta z_\gamma} = \sigma_{\text{prompt}} \frac{|\mu|}{2} e^{-\mu \Delta z_\gamma}, \quad (6.9)$$

where the μ characterizing the probability distribution $dP/d\Delta z_\gamma$ of the decay is defined as

$$\mu \equiv \frac{\Gamma_{\chi_h} m_{\chi_h}}{p_{\chi_h}} \left(\hat{r}_{\chi_h, z} - \frac{\hat{r}_{\chi_h, T} \cdot \hat{r}_{\gamma, T} \hat{r}_{\gamma, z}}{1 - (\hat{r}_{\gamma, z})^2} \right)^{-1}. \quad (6.10)$$

Summing the distributions derived from all the simulated events we obtain the differential cross section in Δz_γ shown in figure 80.

The ATLAS search requires at least two loose photons with $|\eta| < 2.37$ and $E_T > 50$ GeV. At least one photon is required to be in the barrel region $|\eta| < 1.37$. To

avoid collisions due to satellite bunches, both photons are required to have an arrival time at the ECAL t_γ smaller than 4 ns, with zero defined as the expected time of arrival for a prompt photon from the PV. We approximate t_γ with the TOF of the χ_h , requiring it to be smaller than 4 ns. In our sensitivity estimate, we do not include the detailed isolation cuts on the photon. We also neglect the effect of the displaced decay on the angular acceptance of the photons, simply imposing the requirements on $|\eta|$ at the level of the prompt event. The signal region also requires $\mathcal{E}_T > 75$ GeV.

Finally, to simplify the discussion we assume that every event has a reconstructed PV in the geometrical center of the detector.

For events where only one photon satisfies $|\eta| < 1.37$ (i.e. it is in the barrel calorimeter), this photon is used for the measurement of Δz_γ . For events where both photons are in the barrel, the photon with larger t_γ is used. We approximate this timing condition by taking the photon emitted by the more boosted χ_h , in which case the average decay is more delayed. In figure 80 the generated Δz_γ signal distribution is shown on top of the expected background. The latter is taken from figure 4 of the ATLAS paper [301]. Because we are focusing on the non-pointing photon signals, to set constraints on the DM couplings in [156] we remove events with $|\Delta z_\gamma| < 30$ mm. In our exploratory analysis we only consider the statistical uncertainty on the background, neglecting the effect of systematics.

Lessons learned. The ATLAS paper gives detailed descriptions of the cuts and background analysis, which makes an approximate estimation of the signal efficiency quite straightforward.

The background analysis in [301] is based on a data-driven study, for which events passing the diphoton selection with $\mathcal{E}_T < 20$ GeV are used as control region sample. It is challenging for theorists to simulate the background for different energy cuts. This is particularly true for the \mathcal{E}_T cut that plays a vital role in the DM and HV searches.

To obtain a more precise result, it would be useful if the ATLAS collaboration could provide the reconstruction efficiency of non-pointing photons as function of Δz_γ , or of the angle between the photon and the surface of the ECAL, a variable that may be relevant to the efficiency. The paper does provide a table of signal acceptance times efficiency for the GMSB model used therein. However, the numbers depend on details of the particular model used, and it is hard to extract the efficiency that is associated to the non-pointing photon reconstruction. Therefore, when estimating the signal efficiency, we only consider efficiency from the selection cuts and do not include possible suppressions from photon reconstruction.

It would also be very useful to have a table of background events for different cuts on \mathcal{E}_T . For instance, the Dark Penguin signature has \mathcal{E}_T from the decay of $\chi_h \bar{\chi}_h$ with electroweak-scale m_{χ_h} , and \mathcal{E}_T can easily be higher than 100 GeV. By contrast, in the dark shower scenario where soft hidden pions decay to two SM photons, the \mathcal{E}_T originating from additional late-decaying pions can be much lower than the 75 GeV cut used in the ATLAS analysis. Knowing the background and systematic uncertainty for different \mathcal{E}_T cuts would be very important to constrain different models with potentially very different kinematics.

6.4.6. Displaced vertices. Displaced-vertex searches differ from displaced jets and displaced leptons due to the requirement of an actual secondary vertex from the displaced objects. These searches are sensitive to LLP lifetimes that allow it to decay in the inner trackers or MS of the LHC detectors, where vertexing is possible [196, 237, 248, 251, 253, 265, 279]. These searches have extremely low backgrounds as there are no irreducible contributions from the SM, making them sensitive to very small signals of new physics (for more information, see

chapter 4). Moreover, the identification of displaced decays can be used to extract kinematical information in a decay, such as (invisible) particle masses [516, 517].

In this section we review some reinterpretations of DV searches, differentiating between reinterpretations making use of only truth-level information to identify displaced decays and reinterpretations in which an attempt is made to reconstruct DV from displaced tracks (with an approximate detector response).

6.4.6.1. Truth level DV. The work in [26] reinterpreted the 8 TeV ATLAS search for a displaced muon and a multi-track vertex ($DV + \mu$) [518], where LLP signatures for certain weak-scale models of baryogenesis [24, 26, 509] were explored. For reinterpreting this search, a similar procedure described in section 6.4.3, on displaced jets, of constructing ratios of truth-level versus ATLAS efficiencies for the ATLAS multi-track vertex analysis [518] was performed, with similar results for the validation being correct within approximately a factor of two.

This $DV + \mu$ analysis has since been superseded by [196], in which a DV is searched for at 8 TeV in association with either a muon, electrons, jets, or missing transverse momenta. Recently, an updated ATLAS analysis [237], which looks for multi-track DV at 13 TeV in association with large \mathcal{E}_T , was made public. This search now includes a prescription using parametrized efficiencies as a function of vertex radial distance, number of tracks and mass. Their prescription can be applied to vertices and events passing certain particle level acceptance requirements using the truth MC event record.

Here we validate the prescription with parametrized selection efficiencies in [237]¹⁹¹. The results of this search are interpreted by ATLAS in a split-SUSY simplified model with a long-lived gluino that hadronizes, forming an R -hadron before decaying as $\tilde{g} \rightarrow q\bar{q}\tilde{\chi}_1^0$. Event samples are generated with PYTHIA 8.2 [220]. We use truth-level \mathcal{E}_T and we identify the truth R -hadron decay position and decay products, as the ATLAS collaboration provides selection efficiencies that can be directly applied to these truth-level quantities. These efficiencies can be found in the auxiliary material in [519], and are given at the event-level as a function of the truth \mathcal{E}_T and DV radial distance, and at the vertex level parametrized as functions of vertex invariant mass and number of tracks. The efficiencies are given for different detector regions, encapsulating also the effect of the material veto cut.

The selection of events used for the signal region requires:

- truth level $\mathcal{E}_T > 200$ GeV;
- one trackless jet with $p_T > 70$ GeV, or two trackless jets with $p_T > 25$ GeV. A trackless jet is defined as a jet with $\sum_{\text{tracks}} p_T < 5$ GeV.

In addition, events must have at least one DV with:

- transverse distance between the IP and the decay position > 4 mm;
- the decay position must lie in the fiducial region $r_{\text{DV}} < 300$ mm and $|z_{\text{DV}}| < 300$ mm;
- the number of selected decay products must be at least 5, where selected decay products are charged and stable, with $p_T > 1$ GeV and $|d_0| > 2$ mm;
- the invariant mass of the truth vertex must be larger than 10 GeV, and is constructed assuming all decay products have the mass of the pion.

Applying these cuts and efficiencies, we get event-level efficiencies for two of the benchmarks (with gluino masses of 1400 GeV or 2000 TeV, and the neutralino mass is fixed to 100 GeV). Based on the efficiencies obtained and the estimated number of background

¹⁹¹ This prescription is also validated in [504].




Figure 81. Comparison of our recast and ATLAS [237] on the upper limit on the gluino production cross section versus its proper lifetime [504].

vertices, we can extract 95% C.L. upper limits on the total visible cross section for the two gluino masses. For reference, assuming 100% efficiency, we get an upper limit of 0.091 fb. The curves in figure 81 show our recasting results compared to ATLAS. The level of agreement is very good, within 20% for most of the lifetime values. We also point out that the recasting limits agree well even for regions where the efficiency is very low ($\tau > 10$ ns and $\tau < 10^{-2}$ ns). This 13 TeV ATLAS multitrack analysis [237] has also been reinterpreted in the context of long-lived sterile neutrinos [48, 264].

6.4.6.2. Displaced-vertex reconstruction. Before parametrized efficiencies applicable for truth-level DV were made public, attempts to recast DV searches were made by performing their reconstruction from charged tracks only, with an approximate detector response. In [206] the ATLAS DV+jets multitrack analysis [196] was recast. Reinterpretation was performed using generator-level events and the detector fiducial region was reproduced as well as possible. A description of the criteria from the ATLAS analysis that were implemented in the recast is shown in table 9. The jets are clustered according to the anti- k_T prescription [366] in the analysis with momentum smearing, validated by reproducing the jets+ E_T exclusion curve of prompt searches. The selection of events used for the signal region (and the approximations to real detector simulation) were as follows:

A tracking efficiency of the form

$$\begin{aligned} \varepsilon_{\text{trk}} = 0.5 \times & \left(1 - \exp\left(\frac{-p_T}{4.0 \text{ GeV}}\right) \right) \times \exp\left(\frac{-z_{\text{DV}}}{270 \text{ mm}}\right) \\ & \times \max\left(-0.0022 \times \frac{r_{\text{DV}}}{1 \text{ mm}} + 0.8, 0\right), \end{aligned} \quad (6.11)$$

is used, where r_{DV} and z_{DV} are the transverse and longitudinal distance of the track's production vertex (which is the same as the DV origin when using truth-level generator information). This functional form is designed to take into account the size of the detector (i.e. it imposes a linear dependence on r_{DV} , and an exponential dependence on z_{DV}), as well as a turn-on like feature dependent on the p_T of the track. It reproduces the overall behavior of efficiency falling off with vertex displacement. The parameters were determined by fitting the




Figure 82. Validation of the DV + jets search for the ATLAS benchmark of a simplified RPV model with a 700 GeV squark decaying to a neutralino, $\tilde{q} \rightarrow q\tilde{N}_1 \rightarrow q\mu qg$. Left: without any tracking efficiency. Right: with a tracking efficiency function given by equation (6.11), taken from [206].

Table 9. Implementation of cuts applied in the ATLAS multitrack DV + jets search, from [196]. Cuts denoted by an asterisk (*) are approximations to the experimental analysis in the absence of the full detector simulation.

DV jets	4 or 5 or 6 jets with $ \eta < 2.8$ and $p_T > 90, 65, 55$ GeV, each
DV reconstruction*	DV made from tracks with $p_T > 1$ GeV, $ \eta < 2.5$ and $ d_0 > 2$ mm. Vertices within 1 mm are merged. Note: a tracking efficiency needed here; we assume a functional form given by equation (6.11)
DV fiducial	DV within $4 \text{ mm} < r_{\text{DV}} < 300 \text{ mm}$ and $ z_{\text{DV}} < 300 \text{ mm}$
DV material*	No DV in regions near beampipe or within pixel layers. Discard tracks with $r_{\text{DV}}/\text{mm} \in \{[25, 38], [45, 60], [85, 95], [120, 130]\}$.
N_{trk}	DV track multiplicity ≥ 5
m_{DV}	DV mass > 10 GeV

efficiency curve (with lifetime dependence), for three benchmarks in the analyses. We find that fitting only one benchmark does not correctly reproduce the efficiency curve for any of the others. This is most likely due to insufficient dimensionality of the efficiency map. We expect that a full tracking-efficiency parametrization depends not only on r_{DV} , z_{DV} and p_T , but also on transverse and longitudinal impact parameters (d_0, z_0), and on the charge and pseudorapidity of the track. Furthermore, we expect a vertex efficiency that depends on the topology of the event and the nature of the particles forming the vertex. The fit for the event efficiencies from this tracking function can be seen in figure 82.

Lessons learned. With a larger and more diverse set of signal benchmarks, the prospects for reinterpretation are better. For example, the ATLAS analysis examined in [518] only showed limits for three signal model benchmark points for which efficiencies were shown,

making it challenging to find a benchmark whose kinematics matched the desired signal models for the reinterpretations of other models (for example, in [26]). Because the efficiencies and limits were shown for either a high-mass, low-boost LLP or a low-mass, high-boost LLP, this made it more challenging to reinterpret the results for other types of kinematics.

The new parametrized efficiencies presented by ATLAS in [519] are extremely useful. They constitute an optimal efficiency map for recasting these type of analyses, as they can be applied in a straightforward way to truth-level quantities. Before this information was made public, efficiency tables (for vertex-level efficiency) in terms of r_{DV} were only available for few channels and for a single benchmark. It was not clear how to translate this information to other channels, or to parent particles of a different mass. In this case, a functional parametrization for track efficiency was needed (as derived in [206]) to be able to reproduce the experimental results. Finding this kind of parametrization is not easy, as it needs to be validated across different benchmarks. For example, [206] found that fitting only one benchmark did not correctly reproduce the event-level efficiency curve for any of the others.

6.5. Handling LLPs in DELPHES-based detector simulations

6.5.1. Long-lived particle simulation in DELPHES 3.4.1. The DELPHES package [342] allows for the generic simulation of the response of a typical detector used in high-energy physics experiments. It is widely used for simulating the effects of the ATLAS and CMS detectors or the hypothetical detectors that could be used for the future FCC and CLIC projects. The architecture of DELPHES is composed of distinct and specialized modules that interact with each other. The detector is described by the user through an input card, where the modules to be used in the simulation are sequentially enumerated and their input parameters are specified.

The detector simulation relies on a mix of parametric and algorithmic modules. More precisely, tracking is simulated through efficiency and smearing functions that are applied to the properties of the electrically-charged stable particles. The particles are then propagated to the calorimeters and dedicated modules simulate the energy deposits in the electromagnetic and HCALs. The output of such a step consists of a list of calorimetric towers. Moreover, DELPHES includes a particle-flow-like algorithm that combines tracking and calorimetric data in order to improve the identification of the final-state objects and the resolution on their reconstructed momenta.

Jet clustering is performed via an internal call to the FASTJET package [367], which takes as input the list of calorimeter towers or, alternatively, the particle-flow candidates, and outputs jet objects. Lepton and photon isolation is then handled through a specific isolation module. Finally, DELPHES takes care of removing the double-counting of objects that could be simultaneously identified as elements of different collections. The final output is stored in a ROOT file.

In addition, DELPHES allows for the simulation of PU effects by superimposing minimum-bias events attached to displaced interaction vertices along the beam direction to the hard scattering event. Procedures mimicking PU removal can then be configured in the input card. The subtraction of the charged particles belonging to PU vertices is performed at the tracker level. Neutral particles are removed by applying the jet area method [520] supplied within the FASTJET package. More advanced methods are also available in DELPHES, such as the PUPPI [521] or SOFTKILLER [522] techniques, and they can easily be added to the input card. The loss of performance originating from PU, in particular relative to the isolation, is automatically accounted for.

The official DELPHES package (version 3.4.1) with the default detector cards needs to be adapted for the proper handling of LLPs. By default, the decay products of a LLP enter the simulation as if the corresponding decay would have occurred within the tracker volume. However, the user has the possibility to define a volume in which the particle can decay and still be detected outside the tracker volume. This is achieved in practice by making use of the `RadiusMax` parameter of the `ParticlePropagator` module, that is by default set to the tracker radius stored in the `Radius` parameter. When setting the `Radius` and `RadiusMax` parameters to different values, the particles decaying outside the tracker volume, but inside the ‘decay volume’ of radius `RadiusMax`, are included in the collection of stable particles stored in the output ROOT file. They can in this way be used for an offline, more correct, treatment.

Moreover, several modules that are not used in the default ATLAS and CMS cards could serve for a better simulation of the LLPs in DELPHES. For instance, the `TrackSmearing` module allows the user to smear the track momentum according to the impact parameter in the transverse plane (i.e. the d_0 parameter) and in the longitudinal plane (i.e. the d_z parameter).

By default, the detector simulation in DELPHES totally ignores the presence of any LLP. While this is convenient for neutral particles like a neutralino which could be considered as invisible from the detector standpoint, a charged particle leaves tracks in the tracker and would interact with the calorimeters if its lifetime is large enough. In this case, if the LLP decays inside the tracker, its trajectory is properly propagated to the calorimeters and the DV is correctly accounted for. However, if the LLP decays outside the tracker, its decay products are ignored in DELPHES, unless the `RadiusMax` parameter has been specified to be larger than the tracker `Radius` parameter. In this case the decay products can be found in the `Delphes/stableParticles` output collection and treated with adequate smearing functions and efficiency directly from the DELPHES output.

Finally, DTs are simply treated as missing energy in DELPHES. Emerging tracks or tracks containing kinks are not treated appropriately, in the sense that the parameterizations required for a proper description of such signatures has not been implemented yet. Also, DELPHES does not include any trigger simulation, and the latter is in general complex in the case of LLPs.

6.5.2. Displaced tracks with the MADANALYSIS 5 tune of DELPHES. The DELPHES-LLP package can be installed from the version v1.6 of MADANALYSIS 5 [488, 523] and contains improvements of DELPHES specific to LLPs. It leads to new possibilities for phenomenological investigations of LLPs and the recasting of related LHC analyses¹⁹².

This new package was designed to handle neutral LLPs that decay into leptons within the tracker volume. Realistic efficiencies are applied to the displaced tracks and several parameters specific to this kind of analysis have been made available within MADANALYSIS 5. An extension to the case of neutral LLPs decaying into muons outside the tracker volume can be easily implemented, the muons being thus reconstructed only through their hits in the muon chambers. The simulation of the displaced leptons is performed through efficiencies and resolution functions to be specified by the user. Furthermore, another extension allowing the user to handle long-lived charged particles that decay into leptons could be implemented. A similar dedicated tune of DELPHES and MADANALYSIS was devised for the studies of neutral particles decaying to displaced leptons and jets in [44, 49, 169]¹⁹³.

¹⁹² More information is available at the following website: <https://madanalysis.irmp.ucl.ac.be/wiki/MA5LongLivedParticle>.

¹⁹³ The code is available to download at <https://sites.google.com/site/leftrighthep/>.

The DELPHES-LLP package contains a new module called `MA5EfficiencyD0` that allows for the definition of a track reconstruction efficiency parameterized as a function of the $|d_0|$ and d_z parameters (named `d0` and `dz` in the DELPHES input card). The default efficiency function, specified via a `DelphesFormula` is taken from the 8 TeV tracking performance of CMS [275]¹⁹⁴,

```
setEfficiencyFormula
{ (d0 <= 20) * (-5.061 07e - 7 * d0**6
+ 0.000 027 275 6 * d0**5 - 0.000 493 21 * d0**4
+ 0.002 871 89 * d0**3 + 0.005 220 07 * d0**2
- 0.091 795 7 * d0 + 0.924 921)
+(d0 > 20) * (0.00)}
```

In addition, the data-format of DELPHES has been extended so that the Muon and Electron classes now include the transverse ($|d_0|$) and longitudinal (d_z) impact parameters relative to the closest approach point (encoded in the `d0` and `dz` variables), the coordinates of the closest approach point (x_d , y_d , z_d) (encoded in the `xd`, `yd` and `zd` variables), and the four-vector of the vertex from which the lepton is originating from (t_p , x_p , y_p , z_p) (encoded within the `tp`, `xp`, `yp` and `zp` variables), the latter quantity being evaluated from MC information.

Consequently, the data-format of MADANALYSIS 5 has been extended, so that the Muon and Electron classes now contain the `d0()` and `dz()` methods allowing to access the value of the $|d_0|$ and d_z parameters, the `closestPoint()` method that returns the coordinates of the closest approach point (through the `X()`, `Y()` and `Z()` daughter methods) and the `vertexProd()` method that returns coordinates of the DV from which the lepton originates (through the `X()`, `Y()` and `Z()` daughter methods). An analysis example [494] can be found in the public analysis database of MADANALYSIS 5¹⁹⁵, where information about the re-implementation in MADANALYSIS 5 of [197], an analysis of 2.6 fb^{-1} of 13 TeV LHC data, is available. This is a search for LLPs decaying into electrons and muons, where signal events are selected by requiring the presence of either an electron or a muon whose transverse impact parameter lies between $200 \mu\text{m}$ and 10 cm . For a given benchmark signal where a pair of long-lived stops is produced through QCD interactions and where each stop further decays into a displaced b -jet and a displaced lepton. In table 10, we present the number of events surviving the selection of the three different signal regions of the CMS analysis. Our event generation has been performed with PYTHIA 8 [219] and for the benchmark scenario used. We observe that a good agreement is obtained, except in the case of long stop lifetimes, $c\tau \gtrsim 1 \text{ m}$. This is however the region in which no public information on the CMS reconstruction efficiencies is available.

The MADANALYSIS 5 tune of DELPHES has neither been designed for disappearing (or appearing) tracks nor for track kinks. Concerning the disappearing (or appearing) tracks, the only missing experimental ingredients are the track reconstruction efficiency and resolution as a function of the number of missing hits in the (inner) outer layers of the tracker. There is to this date no public material on the tracking performance description related to track kinks.

6.5.3. What about other LLP signatures? In this section, we briefly discuss how DELPHES could be improved for a better handling of LLP signatures.

¹⁹⁴ <https://twiki.cern.ch/twiki/bin/view/CMSPublic/DisplacedSusyParametrisationStudyForUser>

¹⁹⁵ <http://madanalysis.irmp.ucl.ac.be/wiki/PublicAnalysisDatabase>

Table 10. Number of events populating the three signal regions (SR-1, SR-2 and SR-3) of [197] for different stop decay lengths ($c\tau_{\tilde{t}}$). We compare the CMS and MADANALYSIS 5 (MA5) results in the second and third column of the table, respectively, and the difference taken relatively to CMS is shown in the last column.

Region	$c\tau_{\tilde{t}}$ (cm)	MA5	CMS	Difference (%)
SR-1	0.1	3.89	3.8	2
	1	4.44	5.2	15
	10	0.697	0.8	15
	100	0.0610	0.009	>100%
SR-2	0.1	0.924	0.94	2
	1	3.87	4.1	5
	10	0.854	1.0	15
	100	0.0662	0.03	~100%
SR-3	0.1	0.139	0.16	15
	1	6.19	7.0	10
	10	4.45	5.8	25
	100	0.497	0.27	85

6.5.3.1. Displaced jets. Displaced jets are jets that are reconstructed either from SA calorimeter information, or from the particle-flow input with a minimum requirement on the multiplicity of tracks with high transverse displacement (see [251]). Conceptually, such jets can be handled in DELPHES provided that the displaced tracks are properly parametrized. As described above, a module designed to smear the full set of track properties, including their transverse and longitudinal displacement, exists (i.e. the TrackSmearing module). In addition, efficiencies based on displacement parameters have already been implemented in MADANALYSIS 5 (see above) and a module that performs the matching of an existing jet collection with a track collection based on track displacements is very similar to the already existing TrackCountingBTagging module. Minor modifications to this module are hence needed to be able to select tracks based on an absolute displacement instead of the impact parameter significance. Finally, in order to be able to perform a displaced jet selection, one would need a (not yet existing) module that performs jet clustering on the basis of the secondary vertices and the displaced tracks matched with these jets. Alternatively, a module that includes a vertex reconstruction efficiency and a module including a vertex position smearing could be implemented.

6.5.3.2. Displaced vertices. The missing modules described in section 6.5.3.1 could perfectly serve the purpose of a DV analysis. Provided that tracking efficiencies and resolutions are available as a function of the full set of tracking parameters (d_0 , d_z , p_T , ϕ , and θ) and, eventually, of the MC truth vertex position, a simple vertexing algorithm can be implemented in DELPHES.

6.5.3.3. Discussion: DELPHES versus a specific parametric simulation. In a fully parametric simulation of the detector, the detector effects are encoded in terms of efficiencies and resolution functions. Such a simulation is typically very fast, but is likely to suffer from a lack of accuracy in the modeling of complex observables such as jet properties and missing energy. The DELPHES simulation is an admixture of such a parametric simulation and of an algorithmic one. This is slower, but has the clear advantage of correctly treating in particular the reconstruction of the jets and the missing energy.

In order to be able to answer whether DELPHES should be used in place of a fully parametric simulation in recasting LLP analyses, further studies are needed. In the meantime, the following guidelines could be used. If the signal selection is based only on displaced tracks, a simple parametric simulation should in principle be sufficient. This simulation could encapsulate the track reconstruction efficiency and resolution, including PU effects. DELPHES could then optionally be used to mix the resulting ‘reconstructed’ tracks with the additional tracks originating from the PU vertices. On the other hand, if the analysis under consideration additionally uses calorimetric information (i.e. jets or missing energy), DELPHES should be preferred to a fully parametric framework. However, a precise quantification of these effects cannot be assessed without detailed comparative studies between the two approaches. Finally, it should be pointed out that neither of these techniques can be used to correctly simulate the instrumental background, which is challenging to simulate and is discussed in chapter 4.

6.6. Recasting inside the experimental collaborations

Reinterpretations performed within the experiments themselves present unique advantages and disadvantages. They allow for thorough and consistent treatment of detector effects and geometry, object reconstruction, and systematic uncertainties in a way which is impossible through external recasting. Groups can share resources and easily communicate all necessary details. On the other hand, they are of course limited to the model(s) chosen for reinterpretation. In the ideal situation, reinterpretations which provide meaningful results can be performed with minimal overhead to a given analysis.

As the LHC enters an era of decelerating luminosity growth and analysis techniques are becoming more sophisticated, the LHC analyses become harder to re-implement with sufficient accuracy outside of the experiments compared to cut-based analyses. Analyses increasingly utilize machine-learning algorithms that transform a large number of event-level and particle-level observables into higher-level discriminants which are not easily characterized by low-dimensional efficiency tables and may require inputs that third-party detector simulations are not able to reproduce. In particular non-prompt searches may depend on non-traditional reconstruction objects and details of the detector simulation and geometry in ways that require a more detailed simulation than is achievable by, e.g. third-party simulators. Hence, experiments are investigating approaches that enable internal reinterpretation using the full set of available information.

Full-fidelity reinterpretations are especially relevant for LLPs, since the signal simulation may depend more heavily on details not well captured by third-party simulation tools. For example, for sufficiently high lifetime, the decays must be handled by a full detector simulation such as GEANT (or some complex interaction between GEANT and MC packages such as Pythia). Such decays are not well-covered by tools like Delphes as the response of such in-detector decays may require access to a more detailed geometry description.

6.6.1. The RECAST framework. The RECAST Framework [496] is a developing platform for experiments as well as researchers external to the collaborations who wish to reinterpret LHC analyses. RECAST enables cloud-based analysis execution and common presentation of reinterpretation results. The framework consists of two components:

The RECAST front end. This is a web-based service in which reinterpretation of analyses can be suggested by interested authors that provide necessary inputs such as UFO model files [225], process and parameter card templates, or suggested scan grids. Responses to such requests, possibly by more than one analysis implementation, can then be uploaded. Such a




Figure 83. Left: web-based interface for the RECAST back end for an experiment that presents the requested parameter points and color-coded results. Right: an experiment-internal reinterpretation executing on the distributed infrastructure built by CAP and RECAST.

web service, interfaced with services such as HepData, may then serve as a resource for the LHC community to organize and share reinterpretation results obtained by the various analysis implementations.

The *RECAST back end*. An important objective of the framework is to enable a full-fidelity reinterpretation of an LHC analysis using the original analysis code developed within the experiment that can be approved by the collaboration and be placed on an equal footing with the original publication. In contrast to third-party recasting tools, in which multiple analyses are implemented using a single, common framework that is executed on a single computing element, such an exact analysis re-execution often necessitates a distributed data analysis using a number of different frameworks in use within the experiments. Therefore, RECAST has developed a flexible graph-based analysis description and execution back end [524] that enables a faithful re-execution of nearly arbitrary analysis code on cloud platforms such as those offered by CERN¹⁹⁶. The back-end provides experiments with an access-controlled interface to view reinterpretation requests, retrieve the necessary analysis description from repositories such as the CERN analysis preservation portal (CAP) [525], execute the analysis on datasets for the new model and—if approved—upload the results to the public front-end. Figure 84 shows a screenshot of the current prototype user interface, giving collaboration members an overview over requested points as well as controls to steer processing and submission.

These services are being developed in close collaboration with the CERN Analysis Preservation project, which is a common project supported by the four major LHC experiments. While this integration work is on-going, the computing back end for RECAST

¹⁹⁶ Thanks to this flexibility, the popular third-party recasting tools can be easily integrated into this back-end as well, with current integrations being available for CheckMate and Rivet.




Figure 84. Left: pMSSM exclusion in the gluino-neutralino mass-plane. Results partially provided by RECAST. From [528]. Right: follow-up dark-matter reinterpretation. Exclusions presented in chargino-neutralino mass-plane. From [530].

has been successfully used for a number of Runs 1 and 2 reinterpretations published by the ATLAS Experiment.

6.6.2. Analysis preservation as a driver for reinterpretation. Within the LHC experiments, the ability to reinterpret analyses is, perhaps unintuitively, mostly limited by the internal availability of the analysis routines to the wider collaboration as opposed to, e.g. availability of computing resource constraints. The large number of measurements and searches, the heterogeneity and complexity of the analysis software, as well as the size of the collaboration, all lead to a situation in which very often only a small number of analyzers of the original analysis team is able to execute any given analysis. Furthermore, due to the collaborative development model, analyzers are typically responsible for only a subset of the analysis, which results in knowledge fragmentation. Therefore, both ATLAS and CMS are now designing an interface to store analysis-relevant information (the CAP [525]) to mitigate this problem. In the context of RECAST, the software and analysis-workflow preservation aspects of this effort are most relevant. The former is mainly implemented through archiving of source-code repositories and the archival Linux Containers, which now enjoy wide-spread industry support, while internal structure of the analysis workflow is archived in CAP in the form of declarative workflow specifications such as yadage [526], which has been developed for RECAST, and the common workflow language [527]. It is planned that CAP and RECAST will utilize a common computing back-end in order to re-execute the analyses that have been preserved in the portal. As the preservation is enabled by recent technological advances and the process of archival is increasingly streamlined, it is expected that a higher number of experiment-internal analysis codes will be available for RECAST.

6.6.3. RECAST examples. *ATLAS-internal analysis examples and results:* A number of reinterpretation publications have been supported by the back end underpinning RECAST. After Run 1, ATLAS has conducted a thorough reinterpretation of the SUSY landscape in the context of the phenomenological pMSSM [528], a study involving 20 SUSY analyses and 50 000 fully simulated pMSSM parameter points. While at that time, most analyses had to be reinterpreted manually, the 2L electroweak analysis [529] included in that paper served as a prototype analysis and provided results using the highly automated RECAST back-ends.




Figure 85. Exclusion limits as a function of λ''_{323} and $m(\tilde{g})$ as a result of an internal reinterpretation of several ATLAS SUSY results. Expected limits are shown with dashed lines, and observed as solid. Taken from [243].

The analysis was then later re-used with minimal additional effort in two further publications that focused on more domain-specific SUSY realizations: a five-dimensional DM reinterpretation of electroweak searches [530], as well as a reinterpretations in the context of general gauge-mediated models [531].

Recently ATLAS has reinterpreted ten analyses [243] in terms of models of supersymmetry with non-vanishing baryon number-violating coupling strength λ'' partly through the use of the new analysis preservation infrastructure. In such models the lightest neutralino is unstable with a decay length of

$$L(\text{cm}) = \frac{0.9\beta\gamma}{\lambda''^2} \left(\frac{m(\tilde{g})}{100 \text{ GeV}} \right)^4 \left(\frac{1 \text{ GeV}}{m(\tilde{\chi}_1^0)} \right)^5. \quad (6.12)$$

This reinterpretation required a joint re-execution of a mix of analyses originally designed for R -parity-conserving and RPV models and special systematics have been added into the statistical analyses to account for the detector response and flavor-tagging rate of displaced jets. Results were presented in a two-dimensional parameter space of gluino mass and neutralino lifetime (or analogously the λ'' coupling) as shown in figure 85. Such reinterpretations are difficult to perform outside of the experiments, as the publicly available information lack details of the detector and analyses.

Third-party tool integration. Both the CheckMate [486, 487] and the Rivet [490] analysis catalogues have been implemented in the analysis execution framework. Both are configured to analyze events that are provided in the HepMC format [532]. Due to the modular approach of the analysis back end, a number of MC generation workflows, such as Herwig [533], SHERPA [534] or MadGraph [238], can be used depending on their ability to correctly model the desired signal.

For analyses where multiple implementations exists, e.g. from multiple third-party tools such as Rivet BSM, CheckMate or MadAnalysis [488, 489] as well as from multiple

experiment-internal configurations (fast simulation, full simulation), RECAST will allow the community to compare and contrast reinterpretation results.

6.6.4. Outlook. Thanks to industry-backed technological advances, a realistic technical solution to the problem of analysis preservation for the LHC experiments and the original RECAST proposal has come into view. The initial use of such infrastructure for the reinterpretation for prompt SUSY searches is generalized easily for LLP searches as the tools used for signal simulation, analysis preservation, and execution do not make simplifying assumptions on the nature of the BSM signal or analysis structure. As such, RECAST may cover reinterpretation use cases, where either third-party reinterpretations are impossible due to missing public information or limitations of third-party tools or accurate, experiment-approved results are desired.

6.7. Reinterpretation with prompt analyses

Since the decay time probability of an unstable particle follows an exponential decay law (dependent upon the mean lifetime), some percentage of the decays of the LLP will occur outside the detector, leading to an \not{E}_T signature if the LLP is electrically and color neutral. Likewise, some part of the LLP decays can appear ‘promptly’. Prompt searches with and without \not{E}_T can therefore provide additional, corroborating constraints on models with LLPs, especially for short lifetimes. Therefore, it is important to understand the sensitivity of prompt searches to displaced objects.

Reinterpreting prompt searches in the context of LLPs is, however, quite nontrivial, because (a) prompt searches may or may not make explicit requirements on the PV and (b) it is currently not documented how reconstruction efficiencies drop as a function of small displacement. Thus, the reinterpretation of prompt searches in the context of LLPs is currently best done within the collaborations themselves.

An example of such an experiment-internal reinterpretation can be found in a CMS search for an RPV SUSY model where pair-production of stops each proceed through an RPV decay to a b quark and a lepton. A dedicated long-lived search for this model exists in the $e\mu$ channel [197]. This search includes selection criteria which require the transverse impact parameter to the IP be greater than $10\ \mu\text{m}$. This maximizes sensitivity to the long-lived model and greatly reduces SM backgrounds. It also necessarily highly reduces the sensitivity of the search at low stop lifetime. The exclusion curve in the stop lifetime ($c\tau$) versus top squark mass is shown in the left frame of figure 86.

A reinterpretation was performed of a search for pair-production of second generation LQs [535, 536]. In this model, massive LQs are pair-produced. Each of these bosons then decays to a muon and c quark, leading to a final state with two muons and two jets from c quarks. In the prompt limit of the RPV SUSY model with a final state of two muons and two jets from b quarks, the kinematics of the LQs are nearly identical. In the LQ search, no selection is made on jet flavor. The LQ search uses final selections which are optimized to the event kinematics for each LQ mass hypothesis, but the search in general strives to remain as model-independent as possible. In this case, the reinterpretation was simply performed using the original LQ analysis, and replacing the signal samples with the long-lived RPV samples, only taking into account the reduced branching fraction to the final state with two muons and two jets. The expected and observed exclusion curves of the reinterpretation are shown in the right frame of figure 86.

The reinterpretation gives large improvements for lifetimes $\leq 1\ \text{mm}$, and as expected, contributes little in the large lifetime limit. This type of reinterpretation is valuable not only




Figure 86. Left: expected and observed 95% confidence level limits in the $c\tau$ – $M_{\tilde{t}}$ plane from the displaced lepton search [197] for pair-production of long-lived stops decaying to b quarks and leptons. Right: exclusion on the pair-production of long-lived stops decaying to b quarks and leptons in the $c\tau$ – $M_{\tilde{t}}$ plane, from the prompt reinterpretation of the second generation LQ search [536].

because it extends coverage of a given model, but also because it helps guide the analysts performing the dedicated search to focus their efforts in areas which are truly uncovered.

Reinterpretations like this one, which provide meaningful results without placing a large burden on analysis teams, should be highly encouraged. Other reinterpretations along these lines can be found in [537–539]. Another relatively simple reinterpretation is [243], although in this case, it should be noted that the original analysis was modified in a simple way, that is, the reliance on tracking information to identify jets and suppress non-collision background was removed, in order to be sensitive to long-lived gluinos.

6.8. Our proposals for the presentation of results

Here we summarize the recommendations for the presentation of searches involving LLPs. These recommendations follow from the detailed examples presented in sections 6.3 and 6.4.

Our primary recommendation is that the experiments provide as detailed information as possible to make a general recasting feasible. We therefore encourage the experiments to:

- A.1. Provide LLP reconstruction and selection efficiencies at the signature or object level. Although the parametrization of efficiencies is strongly analysis dependent, it is advantageous if they are given as a function of model-independent variables (such as functions of DV d_0 , p_T , η , etc), so they do not rely on a specific LLP decay or production mode.
- A.2. Present results for at least two distinct benchmark models, with different event topologies, since it greatly helps to validate the recasting. For clarity, the input cards for the benchmark points should also be provided.
- A.3. Present cut-flow tables, for both the signal benchmarks and the background, since these are very useful for validating the recasting.

- A.4. When an analysis is superseded, differences and commonalities with previous versions of the same analysis should be made clear, especially if the amount of information presented in both analyses differs. The understanding as the extent to which the information presented in an old version can be used directly in a later version greatly helps the recasting procedure, and also highlights ways in which the new search gains or loses sensitivity relative to the superseded analysis.
- A.5. Provide all this material in numerical form, preferably on HEPdata, or on the collaboration wiki page. A very useful resource we also highly encourage is a truth-code snippet illustrating the event and object selections, such as the one from the ATLAS disappearing-track search [218] provided in HEPdata under ‘Common Resources’.

We realize that implementing the above recommendations requires an enormous amount of time and effort by the collaborations and may not always be feasible to the full extent. However, good examples of presentations are already available, such as the parametrized efficiencies provided by the ATLAS 13 TeV DV [237] (see the auxiliary material to [519]) and the CMS 8 TeV HSCP [495] analyses.

When the object- or signature-level efficiency maps are not feasible, providing efficiencies for an extensive, diverse array of simplified models can be useful for reinterpretation. Concerning simplified-model results, we recommend that the experiments:

- B.1. Provide signal efficiencies (acceptance times efficiency) for simplified models and not only upper limits or exclusion curves. Note that efficiencies for *all signal regions*, not only the best one, are necessary for reinterpretation.
- B.2. Present efficiency maps as a function of the relevant simplified-model parameters, such as the LLP mass and lifetime, with sufficient coverage of the simplified-model parameter space. While for direct production of LLPs the parameter space is two-dimensional (LLP mass and lifetime), simplified models with cascade decays have a higher-dimensional parameter space. In these cases we strongly recommend efficiencies to be provided for a significant range of *all* the parameters.
- B.3. Release the efficiencies in digital format (on HEPdata or the collaboration wiki page), going beyond the two-dimensional parameterization suitable for paper plots whenever necessary. In particular, for auxiliary material, we recommend multidimensional data tables instead of a proliferation of two-dimensional projections of the parameter space;
- B.4. Consider in each analysis a range of simplified models which aim to encompass:
 - (a) Different decay modes, including distinct final-state particles and multiplicities.
 - (b) Different LLP boosts (for example, provide efficiencies and limits for distinct parent particle masses, which decay to the LLP).

Although extensive, the above recommendations for the choice of simplified models allow for a thorough comparison between the range of validity of the LLP analyses and a detailed test of recasting methods. Furthermore, when an MC-based recasting is not available, one can use the ‘nearest’ simplified model or a combination of them to estimate the constraints on a theory of interest. Finally, if a sufficiently broad spectrum of simplified models is covered, this can be useful for quickly testing complex models which feature a large variety of signatures, and rapidly finding the interesting region in a model scan before going to more precise but computationally more expensive MC simulation.

We hope that our recommendations, in particular, points A.1–A.5, will serve as a guide for best practices and help establish a reliable and robust reinterpretation of LLP searches. The added value for the experiments and the whole HEP community will be the immediate

and more precise feedback on the implications of the LLP results for a broad range of theoretical scenarios, including gaps in coverage.

7. New frontiers: dark showers

Chapter editors: Simon Knapen, Jessie Shelton

Contributors: Michael Adersberger, James Beacham, Malte Buschmann, Cari Cesarotti, Marat Freytsis, Gregor Kasieczka, Dylan Linthorne, Sascha Mehlhase, Siddharth Mishra-Sharma, Matt Reece, Sophie Renner, Jakub Scholtz, Pedro Schwaller, Daniel Stolarski, Yuhsin Tsai

7.1. *Introduction: the anatomy of a dark shower*

Hidden sectors are increasingly common features in many models that address mysteries of particle physics such as the hierarchy problem, the origins of DM, baryogenesis, and neutrino masses, in addition to being a generic possibility for physics BSM; ‘hidden valleys’ are one such broad class of hidden-sector scenarios [66, 69]. Given the complexity of the SM, such hidden sectors may very well have rich dynamics of their own, with numerous far-reaching implications for their phenomenology [67, 68, 70, 71, 540]. The main LHC signatures predicted by such hidden valley models are characterized by an injection of a large amount of energy into a hidden sector, which is then shared among a large number of relatively soft particles [70]. We will refer to this class of signatures, where rich dynamics internal to the hidden sector yields a high multiplicity of dark states, as ‘dark showers’. Given that the particles emerging at the end of this process are necessarily both comparatively light and secluded from the SM, their lifetimes can easily become long, thus giving rise to displaced signatures at the LHC.

Long-lived particles are especially generic predictions of hidden valleys with confining gauge groups, similar to QCD in the SM [66]. It is worth noting that QCD already provides many familiar examples of LLPs, realizing macroscopic lifetimes through a hierarchy of scales (Λ_{QCD}/m_W) combined with approximately preserved symmetries (K_L , B and D hadrons) or restricted phase space (n). Also provided by QCD are numerous examples of particles that have a hierarchy of lifetimes. For instance, charged pions experience a slow decay through a very off-shell W boson, while neutral pions can decay much faster through their anomalous coupling to photons. The neutral pion lifetime is thus orders of magnitude shorter than the lifetime for charged pions. Both long-lived species and a hierarchy of lifetimes between species are generic predictions and nearly unavoidable consequences of confining theories which produce dark showers. However, LLPs with a hierarchy of lifetimes also arise naturally in non-confining hidden sectors, especially in theories with multiple species. A familiar example from the literature is theories with dark photons. Here, a small kinetic mixing can make the dark photon long-lived, but a simple, well-motivated extension of the model by a dark Higgs boson introduces additional dark species and production mechanisms for dark states that are independent of the small coupling controlling the macroscopic lifetime. This naturally yields high-multiplicity events featuring particles with a hierarchy of lifetimes [32, 230, 456].

More generally, a dark shower topology can be broken down in three components (see figure 87), each of which allows for a large degree of variation between models:

1. *Production.* A dark shower event begins with the production of one or more heavy states which decay into the dark sector. These heavy states could be part of the SM, most




Figure 87. Schematic representation of a dark shower event, in this case from hidden valley model with a Z' production portal. Figure adapted from [71].

notably the Higgs boson, or could be a new particle from the menu of BSM states we have become accustomed to (Z' 's, color triplets/octets, electroweak doublets/triplets, etc). In some cases the production mechanism provides an important trigger handle (e.g. VH production for the Higgs), but this is not universal (e.g. Z' production). The options for production are laid out in section 7.2.

2. *Showering and hadronization.* If the dark sector contains an asymptotically free gauge group, the originally produced particles will shower and possibly hadronize within the dark sector. This yields a final state with a potentially large number of dark states, similarly to how quarks and gluons undergo showering and hadronization to yield a jet of hadrons. The shape of the shower and the p_T spectrum depends on the coupling of the dark gauge group: the shower may be pencil-like, as in QCD, completely spherically symmetric, or something in between. Alternatively, it is possible that the hidden sector does not contain a gauge sector but instead features a perturbative cascade decay over a large number of states. Indeed, in certain cases the perturbative picture is dual to the strongly coupled showers. In general, showering and hadronization are the sources of greatest uncertainty from a theory perspective; the current status and some new results are discussed in section 7.3.
3. *Decay.* After the dark degrees of freedom hadronize (or reach the bottom of the cascade in perturbative models), they can decay back to the SM. The decay may occur through the same (off-shell) portal as the production, but this is not essential, and one may expect multiple species with a range of lifetimes. The specifics of the decay step (e.g. muon multiplicity) are particularly important if there is no good trigger handle from the production topology. Decays are frequently interconnected with showering and especially hadronization, however, and it is not tractable to enumerate all possibilities without making simplifying assumptions and/or inserting additional theory prejudice. For this reason it is often useful to focus on the species with the largest multiplicity and/or shortest (macroscopic) lifetime; this frequently provides a reasonable guide to the overall signatures, just as one may obtain a reasonable $O(1)$ picture of QCD jets by considering only their pions. We survey a (non-exhaustive) list of popular decay portals in section 7.4.

A priori it is typically possible to construct a model by choosing an ingredient from the menu of options for each of the three components outlined above. This is an enormous model space, and it may appear daunting to construct searches capable of capturing all possibilities.

On the other hand, the signatures of these models are often so striking that they enable powerful, inclusive searches, sensitive to a very large portion of this overall model space, provided that triggers allow the event to be recorded. Toward this end, it is useful to observe that dark shower events have the following generic features:

1. Events have a *variable and potentially large multiplicity of LLPs*. The number of produced particles of various dark species depends on the details of the parton shower and/or hadronization, as in QCD, and varies from event to event. Typically there are more than two LLPs per event.
2. The BSM species produced in dark showers exhibit a *hierarchy of proper lifetimes*. This could result in production of, e.g. mostly prompt particles with a few displaced decays; mostly invisible detector-stable particles with a few displaced decays within the detector; or anything in between.
3. *LLPs are not generically isolated*, i.e. they often appear within $\Delta R \lesssim 0.4$ of other LLPs and/or prompt objects (such as the decay products of short-lived species originating from the same shower).
4. The *energy flow in the event reflects the evolution of the BSM parton shower and hadronization*, and thus looks non-SM-like. For instance, hidden sector jets may be either narrower or broader than QCD jets, depending on the hidden sector gauge group, gauge coupling, and particle spectrum. Additionally, SM particle multiplicity—i.e. the relative fractions of pions, kaons, etc, produced in a jet—differs from QCD. While energy flow can be a powerful discriminant to separate dark showers from the SM at any stage from the trigger level forward, it is model dependent and must be assessed on a case-by-case basis. Note that this also implies that traditional (displaced) jet triggers do not suffice to cover the general model space, and alternative trigger strategies (e.g. track or muon multiplicity) are needed as well.

The existence of $\gtrsim 2$ LLPs per event, and indeed frequently $N_{\text{LLP}} \gg 2$, generally ensures that these events can be easily distinguished from background if they can be recorded on tape and subsequently reconstructed. However, the unique features of events with dark showers require new strategies to ensure that this class of theories is actually captured by the trigger. In particular, non-isolation and the hierarchy of proper lifetimes can result in qualitatively novel collider signatures that require new triggering approaches. Meanwhile, event-level observables such as non-SM-like energy flow or particle multiplicity are potentially powerful but highly model-dependent discriminants, and must be considered on a case-by-case basis. We discuss the strengths and shortcomings of existing triggers and off-line strategies in sections 7.5 and 7.6, as well as a number of new ideas. An executive summary of our main points and recommendations is provided in section 7.7. The chapter concludes with a collection of example models for which MC event samples are currently available (see section 7.8).

7.2. Production

Events with dark showers generically begin with the pair production of dark partons Q_D . In most cases the two produced partons are of the same species, but this does not necessarily need to be true. For clarity and simplicity, we confine ourselves to the case where Q_D is a SM singlet. Then the production modes for Q_D can be simply related to the production modes discussed for neutral LLPs in section 2.4, in the chapter on Simplified Models. In particular, the most relevant production modes are:

- *Heavy parent*: Pair production of a SM-charged HP X_D , which subsequently decays via $X_D \rightarrow Q_D + (\text{SM})$. The SM quantum numbers of the parent X_D , together with its mass, control both the overall production cross section and the typical prompt AOs in the event. Depending on the lifetime of X_D , showering can begin before or after its decay. The model of [331] features this production mechanism, where the parent X_D is a heavy scalar carrying both color and dark color charges. After QCD pair-production, the mediators each decay into a visible jet and a dark shower.
- *Higgs (HIG)*: Production in exotic Higgs decays, $h \rightarrow Q_D \bar{Q}_D$. As the Higgs boson provides an especially sensitive window into low-mass dark sectors, this production mechanism is one of the best-motivated at the LHC [67, 117]. In particular, Higgs portal production is the dominant production mechanism in many Twin Higgs and related models of neutral naturalness [118, 263, 541]. The Higgs boson determines the overall mass scale of the event, often awkwardly low for LHC triggers, while the branching fraction remains a free parameter. As discussed in section 2.3.1, the SM Higgs has a characteristic set of accompanying prompt objects, which can extend trigger options. As also emphasized in section 2.3.1, this category additionally encompasses production through parent Higgs-portal scalars, which may be either heavier or lighter than the SM Higgs.
- Z' (*ZP*): Here a new Z' boson couples to both SM and hidden sector states, allowing for production through $q\bar{q} \rightarrow Z' \rightarrow Q_D \bar{Q}_D$. This scenario was developed in the original hidden valley models [66, 69], and considered again recently in [372]. The Z' mass determines the overall mass scale, while production cross sections depend on its couplings to both SM and hidden sector states. The only typical AOs are ISR radiation.
- *Charged current*: Here the parent dark states couple to the SM through the neutrino portal, $\mathcal{O}_{\text{int}} = HL\mathcal{O}_{\text{dark}}$. When the dark states carry charge under a dark gauge group, this coupling requires (at least) two dark states, one fermionic and one bosonic. For instance, if the dark sector contains both scalar and fermionic fundamentals, ϕ and ψ , respectively, then one can construct the dimension five interaction $\mathcal{O}_{\text{int}} = HL\phi^*\psi$. This production mode has not been well studied in the literature. It is worth noting that given this dark field content, it is generically possible to also construct Higgs portal couplings, $\mathcal{O}_H = |H|^2|\phi|^2$, $|H|^2\bar{\psi}\psi$, which can generally include lower dimensional operators.

As we have taken Q_D to be SM singlets, the first process discussed in section 2.4, direct pair production, is typically negligible. Of course, if BSM mediators connecting the dark sector to the SM, like the X_D and Z' in the examples discussed above, are heavy enough to be integrated out at LHC scales, then the resulting higher-dimension operators can mediate direct pair production of Q_D . Single production of Q_D is generally suppressed, and in many cases impossible, in particular when Q_D transforms nontrivially under an unbroken gauge group. In Abelian theories, single production of a dark gauge boson is possible, e.g. through a loop of heavy bi-charged matter; this gauge boson can then subsequently shower. In perturbative cascades, a single BSM state Q_D may certainly be produced, but whether Q_D goes on to produce one or more ‘showers’ is highly dependent on the detailed kinematics of the event.

In contrast with most of the simplified models in section 2.4, after a parton Q_D is produced it undergoes extensive evolution in the hidden sector, so that there is no one-to-one connection between the initial Q_D and relevant detector objects. Thus, an event that begins with two partons may result in a final state that contains two pencil-like jets, each containing more than one displaced object; a spherical distribution of displaced objects; or anything in between.

The most important consequences of production modes for our purposes are twofold. First and most importantly, the production mode informs the types of prompt AOs in an event, as well as determining overall event rate and energy scale. These AOs can be useful levers to distinguish signals from SM background, beginning at the trigger level. Secondarily, production modes rely on a mediator that couples the dark sector to the SM, which may be a BSM particle like a bifundamental X_D or a Z' , or a SM particle like the Higgs boson. In many models, this mediator-SM coupling also ultimately governs the decay of the dark sector states back into the SM. This is the case for the examples of [66, 67, 69, 118, 263, 331, 372, 541, 542], and is certainly the minimal possibility. It is worth mentioning, however, that decays may be governed by a different interaction than production. As a simple perturbative example, consider the hidden abelian Higgs model. This theory can realize dark showers in a variety of ways; consider for concreteness the perturbative cascade decay chain $h \rightarrow ss, s \rightarrow Z_D Z_D$, which can yield collimated pairs of dark photons when $m_s \ll m_h$ (see, e.g. [543]). In this case a Higgs portal interaction governs production, but the long-lived Z_D decays back to the SM through a separate vector portal interaction.

7.3. Shower

7.3.1. Motivation. A familiar feature of QCD is the formation of jets, sprays of approximately collinear hadrons arising from a parton emitted in a hard scattering process. The physics of jet formation is independent of hadronization or confinement, originating in the singularities of the weakly coupled theory at short distances [544], where the 't Hooft coupling is $\lambda \equiv g_s^2 N_c \ll 1$. In perturbation theory, the differential probability for a quark to radiate a gluon carrying a small energy fraction z at small angle θ is

$$P(z, \theta) dz d\theta \sim \frac{\lambda}{4\pi^2} \frac{dz}{z} \frac{d\theta}{\theta}, \quad (7.1)$$

independent of the underlying hard process. The logarithmic divergences at $z, \theta \rightarrow 0$ indicate that perturbative theories favor radiation that is *soft* (low energy) and *collinear*. This enhanced emission of collinear radiation is the source of jets, even in theories that do not confine at all, such as the perturbative, conformal Banks-Zaks gauge theories [545]¹⁹⁷. The large logarithms appearing in these calculations can be numerically resummed through Markov Chain algorithms, leading to the *parton showers* widely used to model jets in MC simulations.

In the limit of strong 't Hooft coupling¹⁹⁸ ($\lambda \gg 1$) perturbation theory breaks down; soft and collinear radiation are no longer enhanced over more general radiation at wide angles. In QCD, the coupling gradually runs strong in the infrared, but other gauge theories (which could be realized in nature as hidden valleys [66]) exist that have an intrinsically large 't Hooft coupling persisting over a wide energy range. Such theories can be understood with the use of gauge/gravity duality. Examples may be conformal, e.g. strongly coupled $\mathcal{N} = 4$ super-Yang-Mills, or confining, e.g. those detailed in [547, 548]. Such large- λ theories were conjectured to lead to spherical event shapes [70], a result that has been directly proven for strongly-coupled large- N_c CFTs [549]. At colliders, these spherical events would lead to

¹⁹⁷ See [546] for a brief and clear recent introduction to jet physics emphasizing the role of scale invariance.

¹⁹⁸ When making this and similar statements below, we always implicitly mean the coupling at the energy scales being probed by a given measurement. For example, while QCD does become strongly coupled near its confinement scale, treating the theory as weakly coupled is justified provided one only ever talks about jet masses and energies at much larger scales.




Figure 88. Illustrations of event shapes, from jetty (left) to spherical (right) with an intermediate event in the middle. The values of Sphericity and Thrust are: (left) Sphericity = 0.006 36, Thrust = 0.991; (center) Sphericity = 0.530, Thrust = 0.706; (right) Sphericity = 0.940, Thrust = 0.521.

characteristic soft unclustered energy patterns (SUEPs) [182, 242, 550]. An illustration of the range of event shapes from jets to SUEPs is shown in figure 88.

Thus we have two well-understood regimes, jets at $\lambda \ll 1$ and SUEPs at $\lambda \gg 1$. There is a gap in the middle when $\lambda \sim 1$ over a wide range of energies. Both perturbative QCD and gravitational duals approach strong coupling in this regime (from different sides) and cannot be trusted to give accurate predictions; for some questions, the predictions may not even be qualitatively accurate.

This regime is of interest because we want to ensure that LHC experiments do not miss a hidden valley signal simply because it looks different than expected. We should aim to be able to trigger on and analyze these events. While sufficiently long lifetimes may provide useful trigger handles, theories at intermediate 't Hooft coupling could also occur with prompt decays while still failing to provide the types of trigger handles typically associated with prompt hard production and decay. Given our inability to reliably calculate the predictions of this scenario, in this section we will take a pragmatic approach: we push the two tools we have, perturbation theory and gauge/gravity duality, into a regime where we do not fully trust them. To the extent that their predictions overlap, one could gain confidence in the qualitative picture we obtain. Where they differ, one would hope that experimental strategies broad enough to encompass the range of possibilities would also be sensitive to poorly modeled scenarios. In the following we present some initial results, which we aim to improve and expand upon in future published work.

7.3.2. Phenomenological models

7.3.2.1. Parton showers and their limits. One can approach the intermediate regime from the direction of weak 't Hooft coupling using parton shower methods. The coupling is then a direct parameter of the model whose value is easy to vary. However, naively setting the coupling to large values would quickly lead to unphysical results. The reason is that the simplest derivation of parton shower evolution equations assume that the final state can be organized as a series of splittings where $z, \theta \ll 1$ at each iteration. Since the full matrix element for a given final state has no uniquely-defined splitting history associated with it, this soft, collinear condition needs to hold as a function of the final-state kinematics alone. As the




Figure 89. Schematic view of a $2 \rightarrow 3$ branching process, with the initial two-particle state approximated as a localized dipole. For more details and further discussion see [551].

coupling is increased, the increased showering probability means that there is an increased chance of populating the phase space while violating these conditions. Regions of phase space without nominal soft-collinear enhancements (but in truth populated with comparable probability) will end up under-populated, leading to more jet-like events than the underlying theory actually predicts.

Things can be somewhat improved by a more careful consideration of parton-shower methods. It is currently known how to implement iterated corrections allowing one to relax one, but not both, of the inequalities given above. Implementations which allow for the correct inclusion of finite- z (but small- θ) effects can be accomplished by extending the approximation of equation (7.1) with terms non-singular in z .

If we have reason to believe that the wide-angle structure of dark showers will be more critical to detecting models in the transition region, sacrificing finite- z corrections for a better understanding of the finite- θ region is preferable. This requires going beyond the $1 \rightarrow 2$ splitting picture of equation (7.1) as information about multiple particles in the event must be encoded. A known solution is to phrase the shower in the language of $2 \rightarrow 3$ evolution kernels [552], where radiation is treated as coming from dipoles rather than individual charges, as in figure 89. Necessarily more complicated, in this approach the splitting function describing the emission of a gluon with momentum p^μ is then expressed as

$$P(p_r^\mu) d\Phi_r \sim \frac{\lambda}{s_{ij}} \left[\left(1 - \frac{s_{ir}}{s_{ij}} - \frac{s_{rj}}{s_{ij}} \right) \left(\frac{2s_{ij}^2}{s_{ir}s_{rj}} + \frac{s_{ir}}{s_{rj}} + \frac{s_{rj}}{s_{ir}} \right) \right. \quad (7.2)$$

$$\left. + \text{non-singular} \right] d\Phi_r, \quad (7.3)$$

$$, s_{ij} = (p_i + p_j)^2. \quad (7.4)$$

for the process $a + b \rightarrow a + r + b$. Features related to the expected broadening of jets at larger coupling will be well-modeled deeper into the transition region with a dipole shower. Angular structure will be correctly reproduced even at larger couplings while final-state energy sharing of collinear particles becomes increasingly untrustworthy. When some sort of angular localization of energy flow can be expected (i.e. soft-collinear enhancements are still present) the corrections at finite splitting angle may still provide a good approximation.

7.3.2.2. Gauge/gravity duality: spheres to jets. Gauge/gravity duality allows calculation in $\lambda \gg 1$ theories with spherical events, from which we extrapolate toward the $\lambda \sim 1$ regime. The

simplest case is Anti-de Sitter (AdS)/CFT duality, where events are perfectly spherical [549]. AdS space is a warped product of a $(3 + 1)$ D Minkowski space with an infinite fifth dimension. To obtain a theory with a discrete spectrum, we cut off the fifth dimension of AdS space at a hard wall or *IR brane* as in Randall–Sundrum (RS) models [553]. 5D fields then decompose into a tower of 4D Kaluza–Klein (KK) modes (dual to hadrons in the gauge theory), each with an associated wave function in the fifth dimension. This wave function, as well as the 4D mass, is calculated by solving the equations of motion up to quadratic order. The couplings between 4D mass eigenstates are proportional to the overlap of their wave functions in the fifth dimension.

Heavy KK modes decay to lighter modes, which decay to still lighter modes, populating a cascade of particles. The case of a flat (unwarped) extra dimension yields sinusoidal wave functions and modes with linearly spaced masses. In this case, the *KK number* is a conserved discrete momentum in the fifth dimension, and KK modes decay only to daughters at threshold. The RS model breaks translation invariance in the fifth dimension, so KK number is no longer exactly conserved and a variety of decays are possible. Nonetheless, for the simple model of a bulk cubic coupling, KK number is still approximately conserved, as illustrated in figure 90, top panel, and the sum of the two daughter modes is approximately equal to the KK mode of the parent. At each step of the cascade, daughter particles have small momentum. This results in spherical event shapes, with no highly boosted daughters [554]. Here we have studied a scalar field with a Φ^3 interaction for simplicity, with the expectation that qualitative features of the event shapes persist in more realistic models.

To push our toy model into the jetty regime, we introduce interactions that explicitly break KK number themselves. We continue to study the simple case of a scalar field, but now, to move toward jettier events, we include the Φ^3 interaction as a boundary term, restricted to the IR brane. As shown in figure 90, right panel, this opens up a much wider range of possible decays with greater phase space, which lead to more variety in the resulting event shapes. We emphasize that moving the Φ^3 term to the boundary is a simple toy model for the purpose of this section that accomplishes the goal of altering event shapes. However, it is not expected to be a good approximation of the dual of an actual confining gauge theory with smaller λ .

Although the IR brane interaction is used here as an ad hoc tool for generating less spherical events, we expect that similar results can be achieved in a more principled way. As λ decreases toward $O(1)$ values, the expectation from gauge/gravity duality is that more bulk fields, dual to the many single-trace operators in the gauge theory, become light. This suggests that models with several interacting bulk fields are a more faithful toy model of the physics of the dual gauge theories at intermediate λ . Such models could be of interest for a wider set of questions than event shapes, as they move the AdS/QCD toolkit—always a toy model at best—somewhat qualitatively closer to QCD. Such models and their consequences will be studied in a forthcoming publication.

7.3.3. Results. We wish to characterize to what extent our available methods allow us to model the range of behaviors we might expect from new showering sectors. Since we do not expect collimated sprays of final-state particles to be a sensible way of organizing information in the event in all cases, for our purposes here we focus on observables that can be defined globally. In particular, we study a pair of event-shape variables that have proven useful to both establish and provide precision data on the non-abelian nature of QCD.





Figure 90. The branching ratio from a parent with KK-number 80 into two daughters for a cubic interaction on the bulk (top) or on the boundary (bottom). The nonzero probabilities occur along the line where KK-number is conserved for the bulk interaction, whereas the boundary interaction has non-zero probability to decay with large jumps in phase space.

Sphericity is defined as the scalar sum of the two smaller eigenvalues of the sphericity tensor

$$S^{ab} = \frac{\sum_i p_i^a p_i^b}{\sum_i |\mathbf{p}_i|^2}, \quad (7.5)$$

where the sum is over all final-state particles in the event [555]. With the eigenvalues λ_i defined in decreasing order, we have $S = \frac{3}{2}(\lambda_2 + \lambda_3)$, which can take on the values $0 \leq S \leq 1$. Thrust is instead defined via a maximization procedure with respect to all possible axes in the event [556],

$$T = \max_{|\mathbf{n}|=1} \frac{\sum_i |\mathbf{n} \cdot \mathbf{p}_i|}{\sum_i |\mathbf{p}_i|}. \quad (7.6)$$

Both observables essentially measure the divergence of an event from the pencil-like final-state structure of a $2 \rightarrow 2$ scattering process without making any direct reference to jets.

Historically, thrust had the advantage of being infrared and collinear safe. In a given event, the change in the thrust due to an additional radiated parton vanishes as the parton becomes soft or collinear to the thrust axis. Singular regions of phase space thus do not contribute to finite values of the thrust, and its measured distribution in QCD is well described by a perturbative calculation up to corrections that scale as $O(\Lambda_{\text{QCD}}/Q)^2$, where Q is a high-energy scale associated with the total system being probed by the thrust. This is not the case for sphericity—specifically a perfectly collinear splitting still changes the value of the sphericity tensor. A perturbative calculation of the sphericity is then divergent for finite values of sphericity, a divergence which can be tamed by either an explicit cutoff at the hadronization scale supplemented by a phenomenological hadronization model (the approach taken by MC generators) or by absorbing it into a form factor (as done in analytic calculations). In either case $O(1)$ sensitivity is induced to the region of phase space dominated by non-perturbative hadronization effects. For our concerns, this difference can be turned into an advantage, as any apparent difference between the two observables can act as a diagnostic of the sensitivity of our predictions to the non-perturbative parameters in the parton shower.

We generated events from models expected to yield a range of behaviors. We considered extra-dimensional models with both bulk and boundary interactions, with the former expected to yield very isotropic events. For the parton shower, we used a modified version of the VINCIA dipole-antenna parton shower [557] in which an $SU(N)$ gauge theory with only light quarks showers and hadronizes into light mesons with masses $m_{\pi_v}/\Lambda_v \sim m_\pi/\Lambda_{\text{QCD}}$. We then varied the Coupling boundary conditions and one-loop running while adjusting shower cutoffs to ensure that couplings remain perturbative throughout the parton shower.

We summarize the results in figure 91, with the uncertainty in the parton shower distributions coming from considering both transverse-momentum and dipole-virtuality shower ordering. The similar behavior of the two distributions indicates that sensitivity to non-perturbative effects from hadronization are not large. The lowest sphericity/highest thrust distribution is fairly close to that expected from QCD, but a wide range of non-QCD-like behaviors is observable. A significant fraction of the allowed range for these observables is populated by a combination of the extra-dimensional and the parton shower approaches, and the boundary interaction KK model and parton shower models give similar results.

Examining the behavior of observables that are perturbatively incalculable in the weakly coupled theory indicates that a degree of caution is warranted, however, with the two approaches giving qualitatively different results to certain questions. As an example, we look at the correlation between sphericity and total event multiplicity for the two closest parton




Figure 91. Comparison between accessible ranges of parton shower and AdS/CFT-inspired models (labeled ‘KK model’) for sphericity and thrust. In the parton shower, the curves correspond to events produced at $\sqrt{s} = 100\Lambda$, with confinement scale Λ ; the shading indicates the uncertainty resulting from comparing transverse momentum and dipole virtuality shower ordering. The β functions of the theories are tuned such that $g^2(100\Lambda)/4\pi$ is 0.06, 0.12, 0.24 for the 3 distributions. In the KK model case, the dashed curve corresponds to a bulk interaction while the solid curve corresponds to a boundary interaction. We also show expectation values for all distributions.

shower and extra-dimensional models. Displayed in figure 92, we see that multiplicity is broadly correlated with sphericity in the parton shower, while being nearly sphericity independent for the KK model. Such qualitative differences between the parton shower and the extra-dimensional models warrant more detailed study of event behavior in the transition region, while both approaches would benefit from further consideration of how jet-level observables, whether physically well-justified or not, vary over their accessible ranges.

7.4. Decay

In general, the mass spectrum of confining hidden sectors is poorly known, with the exception of a few special cases, like pure glue confining theories [558] and the SM QCD sector. The




Figure 92. Correlation between sphericity and hadron multiplicity for events generated with parton shower (top) and AdS/CFT-inspired models with boundary interaction (bottom).

dynamical process of hadronization is even less well understood, and one must typically resort to uncontrolled extrapolations from the measured fragmentation functions in the SM. Moreover, states with different CP and spin quantum numbers can have greatly different lifetimes even in well-understood examples [66, 540, 559], a problem that is further exacerbated if the hidden sector contains one or more approximate flavor symmetries. For concreteness we here primarily focus on the case where only one dark species has a detector-scale lifetime; other dark states may either decay promptly to the SM or appear as MET. A useful special case realizing this scenario is one in which all internal hidden sector decays are sufficiently rapid for heavier hidden states to promptly decay down to the lightest state in the spectrum, which in turn decays to the SM. The discussion below is focused on the possible properties of this lightest hidden state which decays to the SM. We refer to it as ‘the’ LLP, keeping in mind that in complete models there may very well be multiple species of LLPs. It is therefore important to emphasize that *experimental strategies should not be overly optimized* towards the naïve assumption of a single LLP giving rise to all visible decays, as it may not be generically true. In particular, an inclusive experimental strategy should avoid making detailed assumptions about the distributions of proper lifetimes of long-lived states in the event.

7.4.1. Portals and branching ratios. The available decay channels and branching ratios of the LLP are critical both for the trigger strategy and the off-line analysis. In particular, various (multi-)muon triggers can be very effective for models where the LLP branching ratio to muons is not too small. At the same time, displaced hadronic decays tend to give more discriminating power in the off-line analysis, since they produce a larger number of tracks as compared to the leptonic modes. Since each event contains multiple LLPs, a single event can contain *both* leptonic and hadronic vertices, where one uses the former for the trigger and the latter for off-line background rejection. For this reason it is not straightforward to interpolate sensitivity between lepton-rich and hadron-rich hidden showers, since the optimal search strategy for the intermediate cases is qualitatively different from the strategies for the two extremes.

As a starting point for our exploration, we therefore recommend a small number of theoretically motivated decay portals, which cover the fully leptonic and fully hadronic cases, as well as two other intermediate scenarios. We focus our attention on operators of dimension four or five which do not induce additional flavor violation in the quark sector:

- *Neutrino portal:* If the LLP is a neutral, possibly composite, fermion X , it may decay through a small mixing with the SM neutrinos. This state predominantly decays through the $X \rightarrow \ell^+ W^{-(*)}$ and $X \rightarrow \nu Z^{(*)}$ channels and its decays tend to be rich in leptons. The muon fraction in the final states depends on the mixing angle of the X with the muon neutrino, which is model dependent. However, even if this angle is accidentally small as compared to the mixing angles with electron and tau neutrinos, one still expects a muon in roughly $\sim 10\%$ of all decays through the $W^{(*)}$ channel. In τ -rich scenarios, muons also originate from leptonic τ decays.
- *Hypercharge portal:* It is plausible for a dark sector to contain a vector particle, whether it is an elementary $U(1)$ gauge boson or a composite (the analogue of the ρ in the SM). An elementary vector boson can be copiously produced through decays of a hidden sector meson through a chiral anomaly, analogous to π^0 decay in the SM. Whether elementary or composite, such a vector state generically mixes with SM hypercharge through the kinetic mixing operator $\epsilon B^{\mu\nu} F'_{\mu\nu}$, where $F'_{\mu\nu}$ is the dark vector boson field strength. The

Table 11. Probability of finding exactly one muon/two muons in one LLP decay through the photon and Higgs portals. For the lowest mass points, the branching ratios for the photon and Higgs portals were taken from [458, 440], respectively. For the 8 and 15 GeV benchmark the hadronization was performed with Pythia 8 [219, 501].

Mass (GeV)	Photon	Higgs
0.5	0/0.4	0/0.09
1.2	0/0.35	/
8	0.08/0.16	0.25/0.02
15	0.1/0.15	0.3/0.05

branching ratios of such a state then depend on its mass and can be extracted from data, as shown in table 11 (see, e.g. [230, 458]).

- *Higgs portal*: In the same spirit, it is possible that a hidden sector scalar S mixes with the SM Higgs boson through the $SH^\dagger H$ operator. In this case, its branching ratios to SM fermions are proportional to m_f^2 , with the caveat that non-perturbative effects modify the story substantially for $m_S \lesssim 5$ GeV. For $m_S \lesssim 1$ GeV, hadronic branching ratios can be obtained through chiral perturbation theory; however, in the intermediate range 1 GeV $\lesssim m_S \lesssim 5$ GeV, the theory uncertainties are substantial and we do not attempt to make any quantitative statements in this regime (see [560] and references therein). As shown in table 11, the muon branching fraction predicted by Higgs-portal decays is smaller than for the previous two portals, but it can still be relevant if the (non-isolated) muons from B -meson decays are taken into account.
- *Gluon portal*: The hidden scalar (S) or pseudoscalar (a) could also decay to the SM through a coupling to gluons of the form $S \text{Tr} G^{\mu\nu} G_{\mu\nu}$ or $a \text{Tr} G^{\mu\nu} \tilde{G}_{\mu\nu}$. In this case the direct leptonic branching ratio is zero, although a small number of muons may still be produced in the hadronization process of the gluons.
- *Photon portal*: Similar to the gluon portal, the LLP could decay to two photons through the $S F^{\mu\nu} F_{\mu\nu}$ or $a F^{\mu\nu} \tilde{F}_{\mu\nu}$ operators. The signature for this case is qualitatively different from the previous four, since there are many fewer tracks. In particular, tracks only originate from photon conversions in the detector and suppressed Dalitz decays to $e^+e^-\gamma$. The signal is therefore a trackless jet with a high rate of energy depositions in the ECAL relative to the HCAL.

If the lifetime of the LLP is $\lesssim \mathcal{O}(50)$ m, dark showers generically give rise to multiple decays within the detector volume. In table 12 we show the probability for the Higgs and hypercharge portals to produce an event which contains at least three or four muons, as a function of the number of decaying LLPs in the event. Assuming the event has at least four LLP decays, a multi-muon trigger has good efficiency for the Higgs portal above the $b\bar{b}$ threshold¹⁹⁹ and for all masses for the hypercharge portal²⁰⁰. This trigger strategy should also have good efficiency for neutrino-portal decays, although this scenario is more model dependent. For the gluon and photon portals a different set of triggers is needed, as discussed below. The lesson is that multi-muon triggers do not suffice to cover all possible options, but could provide a reasonably generic trigger path for an important subset of the relevant models.

¹⁹⁹ The importance of muons in decays which are rich in heavy flavor was emphasized in [71].

²⁰⁰ Again, this statement is LLP-lifetime dependent, and for sufficiently long-lived LLPs this strategy will break down as the probability of getting sufficiently large numbers of LLP decays within the detector becomes small.

Table 12. Probability of finding at least 3 muons/at least 4 muons in an event for the photon and Higgs portals, as a function of the number of LLP decays. Branching ratios and hadronization were determined as in table 11.

	Mass (GeV)	3 decays	4 decays	5 decays	6 decays	7 decays
Photon	0.5	0.36/0.36	0.53/0.53	0.67/0.67	0.77/0.77	0.85/0.85
	1.2	0.29/0.29	0.44/0.44	0.57/0.57	0.68/0.68	0.77/0.77
	8	0.13/0.07	0.23/0.13	0.33/0.21	0.42/0.28	0.5 /0.35
	15	0.14/0.07	0.23/0.13	0.33/0.2	0.42/0.27	0.51/0.35
Higgs	0.5	0.02/0.02	0.04/0.04	0.07/0.07	0.1/0.1	0.13/0.13
	8	0.04/0.0	0.09/0.02	0.15/0.04	0.23/0.07	0.31/0.12
	15	0.11/0.02	0.21/0.07	0.33/0.13	0.44/0.21	0.55/0.31

7.4.2. Lifetime. Without additional model assumptions, the theory prior on the lifetime of the LLP is rather weak. We can, however, extract some insight from the generic scaling of the width of the LLP as a function of its mass (m): for the Higgs and hypercharge portals, $\Gamma \sim m$; for the photon and gluon portals, $\Gamma \sim m^3/f^2$ with f the decay constant of the LLP (in this case, the LLP behaves similar to an axion-like particle, or ALP); for the neutrino portal, $\Gamma \sim m^5/m_W^4$. The obvious trend in all cases is that the lifetime rises as the mass decreases, steeply so for the case of the neutrino portal. It is furthermore important to note that the above scalings are lower bounds, and in many models the leading decay portal involves a higher-dimensional operator, leading to a stronger scaling with mass. This is especially relevant in confining models where the hidden states are composite particles. For example, in a pure-glue hidden valley coupled through the Higgs portal [559], the lightest glueball can decay through its mixing with the Higgs, but as the portal coupling contains the dimension-four combination of dark gluon field strengths $G_{\mu\nu}^a G^{a\mu\nu}$, its width scales as $\Gamma \sim m^7 \times (v_h/M^2 m_h^2)^2$, where we have taken $m \sim \Lambda_D$, the scale of dark QCD, and M is the suppression scale of the portal operator.

An additional consideration is that the LLP can be discovered directly in (low energy) collider or beam-dump experiments. Since the LLP has an irreducible production cross section through the same coupling that governs its decay, collider experiments effectively impose a lower bound on the lifetime of these states. Beam-dump and supernova constraints on the other hand rely on a displaced signal and constrain the lifetime from above. These bounds are, however, not typically applicable for masses above a few hundreds of MeV (and for a summary of some of the latest constraints on the hypercharge, photon and gluon portals see, e.g. [561, 562]). However, particles that can be produced in sizable numbers at the LHC are in general coupled sufficiently strongly to the SM that they face upper limits on their lifetimes from BBN. In most cases, this upper bound is essentially mass independent and requires $\tau \lesssim 1$ s, but this upper bound can be much weaker for lighter particles and/or particles that decay to hadron-poor final states. For the photon, gluon and hypercharge portals, prompt decays are currently allowed for masses all the way down to the electron threshold, though for the hypercharge portal LHCb is expected to completely close this window below ~ 100 MeV [341, 382]. For the Higgs portal, current LHCb results from b to s transitions [267, 268] already require $c\tau \gtrsim 0.5$ mm for masses below 4 GeV. For the neutrino portal, lifetimes $c\tau \lesssim 10$ m are excluded if the mass is below 1 GeV and the mixing is predominantly with the electron and muon neutrino [563].

As argued above, if the mass of the LLP is larger than a few hundred MeV, their lifetime is only constrained from above by BBN. In the context of a dark-shower topology, there is, however, a much more stringent upper bound to lifetimes that can be probed at the LHC. In particular, to observe at least a handful of decays in the detector volume, the lifetime in the lab frame should satisfy

$$c\tau_{\text{lab}} \lesssim 10 \text{ m} \times \frac{N_{\text{LLP}}}{N_{\text{decays}}} \quad (7.7)$$

with N_{LLP} the typical LLP multiplicity and N_{decays} the number of observed decays in the detector²⁰¹. For instance, for $N_{\text{decays}} \sim \text{few}$ and $N_{\text{LLP}} \sim \mathcal{O}(10)$ this implies $c\tau_{\text{lab}} \lesssim 50 \text{ m}$.

Near this heuristic upper limit in the LHC's sensitivity to lifetime, the shower effectively gets stretched out over the detector elements, and it is useful to study how the decays are distributed over the different detector elements. We can use a toy MC to make some simple estimates of this effect by making the following simplifying assumptions: (1) the number of LLPs is Poisson distributed; (2) their decay lengths follow the usual exponential distribution with uniform average lifetime in the lab frame $c\tau_{\text{lab}}$; (3) the angular distribution of the vertices is approximately spherical in the lab frame. Figure 93 shows the distribution of the number of decays occurring in the tracker, calorimeter and muon chamber for two benchmark points. We hereby assume for concreteness the approximate CMS geometry with the ‘tracker’, ‘calorimeter’ and ‘muon chambers’ defined by $r < 100 \text{ cm}$, $100 \text{ cm} < r < 200 \text{ cm}$ and $400 \text{ cm} < r < 700 \text{ cm}$, respectively. We find that the number of decays in all detector components is roughly equal, with slightly more decays in the muon chamber for a higher number of LLP decays in the shower. The latter result is easily understood from the larger size of the MS as compared to the tracker and the calorimeter. The qualitative conclusion is that tracker-based searches for multiple DV are likely to have reasonable sensitivity over the whole relevant lifetime range. For the longer lifetimes, the sensitivity necessarily degrades although it can be partially recovered by incorporating the muon chamber in the analysis. On the other hand, searches relying exclusively on the muon chamber or the calorimeter are less sensitive to a wide range of lifetimes. To quantify this effect more accurately, detailed simulations for specific benchmark models accounting for more realistic boost and η distributions are needed, although the relative importance of the tracker is likely to continue to hold. The displaced SUEP scenario is an interesting exception, as discussed in section 7.6.2.2.

7.5. Trigger strategies

7.5.1. General considerations. Due to the immense diversity of hidden sector models, there is no trigger strategy which can comprehensively cover all options, nor is it straightforward to compile a list of such strategies. This problem is further compounded by the lack of available MC simulations, except in a handful of special cases. Many of the considerations informing the trigger strategies discussed for singly and doubly produced LLPs in chapters 3 and 5 apply to dark shower events as well. However, the four characteristic features of dark showers laid out in section 7.1 can pose both additional opportunities and challenges:

1. *The particle multiplicity can range from high to very high.* This in itself can provide a very powerful trigger strategy, especially when the muon fraction in final states is appreciable, as discussed in section 7.4.1. For models where this is not the case or where the muons are too soft to fire the trigger, triggers relying on a large multiplicity of tracks

²⁰¹ If only one or two LLPs are expected in the detector, then the dark shower maps onto the topologies with one or two DV discussed earlier in this document.




Figure 93. Distribution of decays occurring in the various detector elements, for those events with more than three decays occurring in the fiducial volume.

or an over-density of pixel hits could provide an alternative strategy. We will return to this discussion later in this section.

2. *The hidden sector may contain states with vastly different lifetimes.* While this makes it challenging to get a comprehensive grasp on the full space of possible event topologies, this feature provides interesting opportunities from the point of view of the trigger. For instance, often the most striking off-line signature is a set of DV in the tracker, which are notoriously challenging targets for the trigger. It is completely plausible that these vertices would originate from the decay of one species of hidden sector particles, while a second species decays either promptly or very displaced. This means that the event may very well come with prompt leptons, H_T , an appreciable amount of \cancel{E}_T , and/or activity in the muon chamber.
3. *The decay products may not be isolated.* Combining the first two features, we should prepare for the possibility that the objects we aim to use as trigger handles may not be

isolated. Especially for triggers on muons and (displaced) tracks, it is advisable to relax isolation criteria as much as possible and instead rely on the presence of one or more additional objects to reduce rates as necessary. For example, one would expect that a trigger on two isolated muons would be less effective than a trigger on three or four non-isolated muons or a trigger with two non-isolated muons and a moderate amount of \mathcal{E}_T . Non-isolation is likely to be a limiting factor for the acceptance of dark shower events in specialized triggers designed for singly- or doubly-produced LLPs.

4. *The energy flow in the dark shower may be non-standard.* This implies that a dark shower may be broader or narrower than SM QCD jets, and the momentum distribution of the particles in the shower may also differ substantially from QCD predictions. This means that triggers for (trackless) jets could be effective for some models, but would by no means capture the full range of possibilities. Since such p_T distributions are very model dependent, it is moreover sensible to keep object p_T thresholds as low as possible, and instead rely on the high multiplicity by requiring multiple low p_T objects, rather than a few high p_T objects.

In summary, while some dark shower models will be very challenging to trigger on, this is not true for the entire class of models. For example, some (although certainly not all) models can readily be captured on the traditional \mathcal{E}_T , H_T or lepton triggers. This means that these traditional triggers *are an excellent place to start, but a bad place to stop*. Indeed, in the context of limited resources, it makes sense to first pursue the scenarios where the trigger is not a major challenge and use these scenarios to develop off-line reconstruction techniques. It is, however, crucial to follow up with more innovative trigger strategies. While this is largely still a topic of study, we can identify two categories of non-traditional triggers:

1. *Triggers on displaced objects:* This type of trigger is notoriously challenging due to bandwidth and online computing limitations, but nevertheless a lot of progress has been made in the last few years. This is discussed in chapters 3 and 5. For our purposes here, it suffices to note that high multiplicity of LLPs in dark shower events can help provide excellent acceptance in such triggers, provided that any isolation criteria in the trigger are not spoiled by shorter-lived objects in the event.
2. *Specialized triggers for dark showers:* There are a number of (speculative) ideas for triggers which exploit the high multiplicity and/or energy features of a dark shower, as described above, although this area is much less developed than the triggers on displaced objects. We will discuss these ideas for specialized triggers briefly in the remainder of this section, and comment on the unique capabilities of LHCb when it comes to triggering on dark showers.

7.5.2. Specialized trigger ideas. In this subsection we briefly highlight some proposals for specialized trigger strategies that inherently depend on the unique features of dark shower events.

7.5.2.1. Multiplicity triggers. As emphasized throughout this chapter, dark shower models can produce a high multiplicity of dark particles, ranging from several tens up to thousands of particles in the most extreme case. This motivates triggers that aim to minimize p_T thresholds and instead exploit high particle multiplicity. The most obvious option is a suitable multi-muon trigger (three, four, or more muons), where the priority should be to avoid tight isolation criteria, if needed at the expense of increasing the muon multiplicity and/or p_T thresholds. Prompt multi-muon triggers would, however, only capture (nearly) prompt

decays, and it is critical to pursue the design and implementation of a trigger with good efficiency for displaced muons; see section 5.1.3 for more on planned displaced muon triggers.

Another way of exploiting the particle multiplicity is to trigger on an anomalously large number of tracks originating from the PV. This would likely only work for prompt or nearly prompt decays, in models which produce a rather low muon fraction and/or a very soft p_T spectrum. For this strategy to work, one must first pass the L1 trigger with MET, which can be provided by the shower itself or by initial state radiation (ISR). At the HLT it may be possible to count the number of tracks with the planned ATLAS HTT system or the future CMS HTTs. Alternatively, if the tracks are too soft and the multiplicity high enough, it is possible to pass the HLT by looking for an over-density of hits in the innermost tracking layer, centered around the z -coordinate of the PV [242].

7.5.2.2. Substructure triggers. Substructure triggers may open up the possibility of using substructure techniques to distinguish between QCD jets and other types of showers already at the trigger level. CMS currently uses jet trimming at HLT level to enhance acceptance for jets with hard splittings, optimized for boosted electroweak gauge bosons, (see, e.g. [564, 565]), and we are aware of an effort within ATLAS to investigate the inclusion of substructure routines at the HLT level. This strategy could be helpful in cases such as photon-jets (dark showers that decay into SM photons), which can otherwise pose significant trigger challenges: substructure variables such as n -subjettiness [566] and energy-energy correlations can be very effective at separating QCD jets, photons and photon-jets [567, 568]. In practical terms, this means that one may be able to use the relatively low threshold jet/tau trigger for the L1 trigger. This could be followed by a HLT trigger that uses, for instance, the ratio of the energy deposited in the ECAL/HCAL or substructure variables.

An advantage of substructure methods is that they are fairly robust with respect to the lifetime of the dark shower final states, provided the lifetime is not too long compared to the size of the relevant detector sub-system. As long as the showering process itself happens quickly (and provided that hadronization represents a small correction to the typical particle momenta), the separation between energy depositions in the calorimeter is already pre-determined. In particular, under these assumptions it does not matter if the SM decay products of the lightest dark states are themselves collimated. A clear difficulty with substructure triggers is that the actual dark showering and hadronization process, and therefore the substructure, is theoretically only understood in a handful of example models as discussed in section 7.3. While such a trigger is clearly valuable for a number of models, it is at the moment difficult to evaluate how broadly applicable it could be to different shower shapes.

7.5.2.3. Non-isolated DV trigger in the ATLAS MS. For sufficiently LLPs, fallback trigger strategies can be provided by existing ATLAS triggers that require a cluster of hits in the MS [241, 257]. As these triggers look directly for the presence of displaced objects and do not rely on any features of the rest of the event, they offer potential sensitivity to signals dominantly produced at low mass scales. The existing MS DV trigger requires at least three (barrel) or four (endcap) tracks in the MS [257]; a dedicated vertex-finding algorithm is run on these events at the analysis level. Thus this trigger option offers good sensitivity to hadronic decays of particles with mass $\gtrsim 5$ GeV, but cannot provide sensitivity to individual di-leptonic or photophilic decays, nor to decays of LLPs with GeV-scale masses.

In Run 1, LLP searches with DV in the MS relied on a trigger that recorded events where the DV additionally passed isolation requirements. Specifically, events were required to have no tracks with $p_T > 5$ GeV in the ID within $\Delta R < 0.4$ of the DV, and the distance to the

nearest jet (with $E_T > 30$ GeV) was required to be $\Delta R_j > 0.7$. However, this isolation cut was not imposed for jets with anomalously large energy deposition in the HCAL, $\log_{10}(E_{\text{had}}/E_{\text{EM}}) > 0.5$ in order to improve acceptance for LLPs decaying in the outer edges of the HCAL [241]. This isolation requirement can limit acceptance for models with dark showers where multiple BSM species can spoil each other's isolation, especially in cases where the shower contains both short-lived as well as long-lived dark particles. It is an open question whether tracks from sufficiently DV in the ID would contribute to the track veto, but particles that decay within the ID would be likely to spoil the isolation requirements through calorimeter deposits alone, unless their decays involved only muons.

New in Run 2 is a trigger on *non-isolated* clusters of hits in the MS [258], which uses the same triggering algorithm for the MS hits but does not impose isolation criteria. The primary aim of this trigger is to enable background estimation for single DV searches [262], but it would also efficiently record dark shower signal events, regardless of event shape or event energy scale, provided at least one particle species in the shower has $m \gtrsim 5$ GeV, decays to final states with at least three charged particles, and has reasonable probability to decay in the MS. This trigger may be especially useful for high-multiplicity signals, where the probability for getting one particle to decay in the MS can become appreciable even for shorter-lived LLPs.

7.5.2.4. Capabilities at LHCb. The LHCb detector is designed for studying displaced decays of SM mesons. This provides a good environment to search for dark shower events containing LLPs with lifetimes below the meter scale [68]. Although LHCb does not yet have a dedicated trigger for dark shower events, many of the trigger strategies used for SM meson searches may be applied to dark shower signals. The LHCb detector is a single-arm forward spectrometer covering the pseudorapidity range $2 < \eta < 5$, which means the signals for which it is optimal have typical p_T that is much lower than the signals targeted by ATLAS and CMS. Hence, the LHCb signal trigger usually has much weaker p_T requirements compared to an ATLAS or CMS search. This allows for better observation of dark shower events that come from the decay of light parent particles boosted in the forward direction, such as dark mesons from the decay of a light Z' .

For example, [387] searches for a pair of LLPs decaying into jet pairs. Although the signal is different from a dark shower, the trigger strategy in the search can be useful for dark shower events. The hardware trigger (L0) requires a single hadron, electron, muon, or photon with object-dependent p_T thresholds. For muons (hadrons), the thresholds are $p_T > 1.48$ (3.5) GeV, and given the rapidity range of the signal, this corresponds to momenta $p > 6$ (13) GeV for muons (hadrons). These low p_T requirements can thus keep soft and low-mass dark meson decays that can be hard to trigger on at ATLAS and CMS. The software trigger contains algorithms that run a simplified version of track reconstruction and an identification of displaced tracks and vertices, and the first software stage (HLT1) requires a high-quality displaced track satisfying the above p_T requirement. In [387], the final trigger stage (HLT2) further requires a DV with ≥ 4 charged tracks, and either a > 2 mm decay length in the transverse direction or the reconstructed vertex mass > 10 GeV. Since the search focuses on b -quark final states, the HLT2 trigger also contains a multivariate algorithm to identify b -hadron decays. Although the trigger is designed for LLPs heavier than 10 GeV scale, it is important for LHCb to determine if the same strategy can be applied to even lighter particles or different decay final states.

For specific decay channels of dark mesons, we can also adapt triggers from existing SM hadron searches that look for the same final states [383]. For example, if dark mesons decay into muon pairs or c -quarks, we can use similar triggers as in the $K_s \rightarrow \mu^+\mu^-$ [569] or

$B^0 \rightarrow D^+ D^-$ [570] searches to study the signal. A more recent muon trigger is used for the 13 TeV dark photon $A' \rightarrow \mu^+ \mu^-$ search [269], where a muon with $p_T > 1.8$ GeV is required at the hardware trigger level, and further quality cuts on the DV and muon identification are required. This search can be useful for dark mesons that decay through a kinetic mixing, in which muon final states have a sizable branching ratio.

Planned upgrades at LHCb are further discussed in section 5.2. These upgrades will make LHCb an even more powerful facility for studying LLPs, particularly in the low-mass and short-lifetime regime, and further study of LHCb's capabilities for dark showers is well warranted.

7.5.2.5. Low PU data. While Run 2 of the LHC has brought unprecedented opportunities for discovery of new physics with a large ~ 150 fb^{-1} dataset of 13 TeV pp collisions, the ability to explore the energy frontier comes at a cost: high trigger thresholds and challenging experimental conditions may limit the sensitivity of LHC experiments to BSM models such as SUEPs with soft and diffuse signatures. During Run 2 data-taking, the typical $\langle \mu \rangle$ value has already reached ~ 60 interactions per bunch crossing.

The low PU datasets provided by the LHC during Run 2 therefore present an interesting opportunity for dark showers. There are two such data sets: one with 0.5 fb^{-1} at 13 TeV and $\langle \mu \rangle = 2$ and one with 0.3 fb^{-1} at 4 TeV and $\langle \mu \rangle = 5$. While the size of these datasets is much lower than the 13 TeV high- μ pp dataset, the change in data-taking conditions is amenable to searches for BSM scenarios which would normally be difficult to distinguish from PU noise, and low-background searches may be performed under circumstances that would otherwise be impossible.

In low PU data, object-multiplicity (track- or cluster-based) triggers (section 7.5.2.1) are typically run with lower thresholds and higher rates than otherwise may be possible. The ability to cleanly reconstruct low- p_T tracks and vertices or soft calorimeter clusters could provide the only way to experimentally access some low-mass LLP scenarios. Simple analyses in this modest dataset may provide the first limits on some models, and these results could help direct the more advanced developments discussed throughout this report.

7.5.2.6. Zero bias strategy. If all other strategies fail to capture the signal, the zero bias strategy is an ultimate fallback [571]. In this case one would simply rely on passing the L1 trigger due to an object in an unrelated PU event. The effective dataset one can ultimately obtain this way is only $\sim 0.5 \text{ fb}^{-1}$, though it might increase to $\sim 50 \text{ fb}^{-1}$ if a specialized selection could be made at the high level trigger.

7.6. Off-line analysis

In this section we discuss aspects of the off-line analyses needed to discover dark showers. As has been a primary theme throughout this chapter, hidden sector states with displaced decays are generic, but prompt decays are possible as well. Naturally cases with and without promptly decaying species require different strategies. Though the main focus of this document is on long-lived phenomena, showers where *all* species are short-lived are closely related to the signals of our main interest and provide valuable illustrations of tools and techniques that can ultimately shed light on the underlying hidden sector dynamics itself, and we accordingly provide a discussion of prompt showers as well.

7.6.1. Prompt decays. We begin by discussing dark showers with promptly decaying BSM states, i.e. with no reconstructable displaced decays in the event. Promptly-decaying dark

showers are substantially more challenging to separate from backgrounds than showers containing LLPs, as the presence of multiple displaced objects is a very powerful background-suppression tool. However, the techniques that have been proposed and/or used for prompt showers are important for several reasons. First, they provide a useful illustration of how the unique properties of showering events have been approached in analyses to date without introducing the separate complication of displacement. Second, dark showers that produce prompt SM particles may very well also produce LLPs, thanks to the hierarchies of lifetimes that are generic in confining theories. Such events can thereby produce *semi-visible jets* [372], which contain detector-stable invisible states as well as promptly-decaying states. These semi-visible jets pose some specific challenges in analysis and reconstruction, as we briefly review below.

If SM particles resulting from the dark shower are produced promptly, there are two possible experimental handles: the portal through which the hidden sector couples to the SM, and the structure of the dark shower itself. The former case depends on the operators mediating the production and decay of showering states, as described in sections 7.2 and 7.4. For example, for decays governed by the Higgs, photon, and neutrino portals, one may expect an unusually muon-rich jet (section 7.4). Similarly, the jet may be semi-visible (section 7.6.1.1) if some of the states do not decay in the detector volume. On the other hand, if the SM final states are almost purely hadronic, as when decays are governed by the gluon portal, no such obvious handles are available and one must look at the substructure of the jets themselves to find differences from those in typical QCD events. The discussion below is organized according to the typical size of the dark shower, going from narrow QCD-like jets to large-radius jets and finally to fully spherical topologies (SUEPs).

7.6.1.1. QCD-like jets ($R \sim 0.4$). The case of LJs, originally motivated by DM considerations, provides an example of showering with noticeably different particle content than a QCD shower. Here, visible decays of the dark states are primarily to leptons [155, 572], which can occur either due to other decays being kinematically inaccessible or due to selection rules. The most striking signature of the model is the presence of collimated sprays of leptons rather than hadrons in the final state. Reconstructing these objects requires the elimination of lepton isolation criteria typically imposed on leptonic final states, together with variables such as layer-specific cuts on energy deposited in the ECAL, and a fraction of HT TRT hits and/or activity in the muon chambers, selecting for electromagnetically dominated radiation. Searches for prompt lepton jets have been performed by the ATLAS and CMS collaborations [274, 573].

A more elaborate prompt scenario is the case of *semi-visible jets* [372]. Here, the visible final states are hadrons, as in typical QCD jets, but missing energy is interleaved with the visible final states as only some fraction of the states produced in the dark shower decay back to SM particles, while others escape the detector. Semi-visible jets fail to provide the dramatic handles of the lepton jet scenario. However, by considering combinations of large amounts of \cancel{E}_T and differing cuts on $\Delta\phi$, the angular separation between the \cancel{E}_T and the closest jet, control of both SM backgrounds and discrimination from more typical SUSY-like high \cancel{E}_T scenarios can be achieved. Bounds on showers with a fraction of between 0.2 and 0.9 of produced dark hadrons decaying invisibly should allow for bounds comparable to conventional resonance searches [574].

If event-level handles like lepton (or heavy flavor) multiplicity or \cancel{E}_T are not present, sensitivity might still be achieved by looking toward the internal structure of the jets themselves. While this topic is presently relatively uncovered, with only a few exploratory results available [575], similar techniques to those that have proven useful for quark/gluon

discrimination seem promising. In both cases the lack of a perturbatively-generated hard scale means that parametric separation of signal and background is challenging. Infrared and collinear (IRC) safe observables related to jet mass, such as girth and two-point energy correlations are reasonably well studied in QCD, with results readily generalizable to other gauge theories [576–578]. For quark/gluon discrimination, observables characterized by Poisson-like distributions such as particle/track multiplicities or production ratios of particular SM particles tend to yield better discrimination, but suffer from being IRC unsafe and thus subject to large non-perturbative modeling uncertainties [579], a significant point of concern when extending their use beyond QCD. Here, some recently developed IRC-safe generalizations of multiplicity [580] may prove useful. Another approach might be to explore machine-learning techniques that do not require fully labelled data for training (so-called weak supervision) [581–583]. These can be trained directly on data, here at least in the case of the QCD background, so that modeling concerns about the shower can be partially alleviated. While achieving a sizable signal-to-background ratio is likely to be the major challenge in using such jet observables to separate dark showers from QCD, it is vital to note that if signal/background separation can be achieved through other means, e.g. by reconstructing DV, these same tools offer an enormously powerful window onto the underlying dynamics of the dark shower.

7.6.1.2. Large- R jets ($R \sim 1.0$). As discussed in section 7.3, dark showers at larger couplings are currently poorly understood. If couplings controlling the shower can no longer be treated as small at any scale, the QCD-inspired picture of pencil-like jets starts to break down, and we might expect that showers radiate more copiously at ever larger angles, so that large- R jets can become necessary to adequately capture the underlying hard process. This is first and foremost a triggering challenge, since it is no longer clear that triggers designed for local hard depositions can maintain their sensitivity. It is likely that a viable trigger path either needs to rely on a prompt associated object from the production mechanism (ISR jet, lepton, etc) or on a high multiplicity of leptons in the shower itself. Of course, H_T triggers continue to be useful for those models with sufficiently high event energy scales, provided that a sufficient amount of the radiation in the event clusters into trigger objects of sufficiently high p_T (\sim few $\times 10$ GeV) to be counted in the H_T trigger. As long as the event can be triggered on, observables similar to those at the end of section 7.6.1.1 should still prove useful to separate signal from background, possibly even more so, since the radiation pattern in the dark shower can now be expected to be substantially different from QCD.

7.6.1.3. Spherical showers. In the large-’t Hooft coupling regime, in which we expect much more showering at large angles than in QCD (see section 7.3), the final states are soft and spherically distributed in the rest frame of the shower. This leads to so-called SUEPs in the detector. In the rest frame of the event, the momenta of the final-state SM particles are on the order of the hadronization scale of the dark strong dynamics. If this scale is much lower than 100 MeV, the decay products are too soft to be reconstructed as tracks and the entire shower is effectively invisible. In this case existing jets+MET searches can apply. For a hadronization scale around $\mathcal{O}(100 \text{ MeV})$, some amount of ISR can be needed to boost the particles enough to render a fraction of them reconstructable. Finally, if the decay products are on average harder than $\mathcal{O}(100 \text{ MeV})$, the tracks associated to the SUEP vertex can typically be reconstructed off-line, subject to momentum-dependent reconstruction efficiencies. The main parameters of interest are accordingly the number of charged particles produced and the corresponding p_T spectrum. Also important are the fraction of invisible particles and the composition of SM particles in the final state. In particular, since the momenta of the invisible




Figure 94. Charged particle multiplicities measured in minimum bias events with the ATLAS detector for events with at least one track with a minimum p_T of 500 MeV and $|\eta| < 2.5$ [584].

particles balance on average, they are not likely to give rise to a substantial \mathcal{E}_T signal unless the dark shower is recoiling against a relatively hard ISR jet. A large fraction of invisible particles will instead degrade the sensitivity of a strategy relying on track multiplicity. The muon fraction, on the other hand, can be a powerful handle: even though the average muon p_T may be very low, due to the high particle multiplicity a handful of muons may still be hard enough to pass the trigger and reconstruction thresholds of the MSs.

A major background to SUEP-like signals comes from PU, which also yields a large number of isotropically distributed soft tracks, though in contrast to SUEP signatures, PU tracks arise from multiple vertices. Studies on the multiplicity of charged tracks from minimum-bias interactions like single, double or non-diffractive collisions are described in [584]. As shown in figure 94, the fraction of 13 TeV pp collisions having 80 or more associated charged tracks is $O(10^{-3})$, as measured in a PU-free environment. Some benchmark SUEP models are described in [242], one with a low-mass Higgs-mediator and two higher-mass scalar models, and charged particle multiplicities for these benchmarks are shown in figure 95; for these models, it is assumed that SUEP particles decay only to electrons and muons. Reference [242] demonstrates that counting tracks associated to the PV should provide a very powerful discriminant against PU background during off-line reconstruction for high-mass mediators, while discrimination for the Higgs-mediator model is extremely challenging. The crucial discriminating factor is the number of particle tracks that can actually be reconstructed, and both ATLAS and CMS can reconstruct tracks with a minimum momentum of roughly 400 MeV [585]. The multiplicity of tracks fulfilling this minimum- p_T requirement is shown in figure 95(c) for the three SUEP benchmark models. It was shown in [242] that, even for the highest-mass mediator, no significant losses of the




Figure 95. Event multiplicities of charged particle tracks fulfilling the respective indicated p_T requirements for three benchmark SUEP models. Figure from [242].

tracking efficiency are expected in spite of the hundreds of tracks present per event. It is moreover worth noting that, while the track multiplicity is large compared to what is generated by SM proton–proton collisions, it is still relatively small compared to the multiplicities that can be reconstructed in heavy ion collisions.

A large fraction of electrically neutral hadrons produced in the SUEP can further lower the number of associated tracks, which makes signal/background discrimination very difficult for the lowest-mass-mediator signals. A possible additional discriminant for those cases can be the hemisphere mass. This mass is estimated by dividing the event into a hemisphere associated to the ISR jet and one associated to the SUEP, while calculating the ‘jet’ mass, as in section 7.6.1.1, from the tracks in the SUEP hemisphere. This variable should be significantly higher for the SUEP events, where the heavy mediators are decaying, than for PU events.

Another possible background to SUEP signals is QCD multi-jet events. As described in section 7.5.2.1, triggering SUEP events based on \mathcal{E}_T relies on the emission of an ISR jet recoiling against the system. QCD scenarios similar to SUEP signatures could arise when one hard jet is recoiling against a system of multiple soft jets. The mean number of charged tracks per jet is less than ten for jets $O(100 \text{ GeV})$ [586]. To get to $O(100)$ charged tracks, as expected from a SUEP signature, a very high number of jets is needed, hence the perturbative cross section is heavily suppressed by many orders in α_s . This background can nevertheless play a significant role for the low-mass mediator models, where significantly fewer tracks are

expected. A veto against large calorimeter deposits can be a powerful handle to reject those events, as rather hard jets are needed to get many associated charged tracks.

7.6.2. Displaced decays. Outside the prompt regime, the lifetimes of the various dark states and the composition of SM final states produced in their decays drive the phenomenology. LLPs result when there are one or more species of dark states in the hidden sector spectrum that would be stable in the absence of couplings to the SM. In particular, different species naturally come with vastly different lifetimes, as is the case for the SM π^0 and π^\pm mesons. In what follows, we consider the single-species case in most detail, and treat the case of primarily leptonic decays (LJs) separately from the cases with substantial branching fractions to hadrons (emerging jets). We finally comment on the multi-species case, and relate it to the semi-visible jet scenario mentioned in the previous section.

Even in high multiplicity events, the isolation of displaced decays from each other is not likely to pose difficulties in reconstructing DVs in the inner tracker. Two nearby DVs can be separately resolved down to separations of \sim mm. Even in the case where two DVs are closer than 1 mm, all of the tracks associated to both vertices will simply be reconstructed as a single DV with larger track multiplicity.

7.6.2.1. Single species leptonic dark showers (displaced LJs). One signature predicted by dark shower models is LJs [155, 572], whose prompt decays are discussed in section 7.6.1. Depending on the lifetime of the decaying state, these leptons can easily be produced with a measurable displacement. Explicit searches for this signature have been performed by the ATLAS Collaboration at both 8 and 13 TeV [272, 273], as detailed in section 3.2. It is worth remarking that the experimental signature of a displaced decay to electrons can be very similar to the signature of a displaced decay to a pair of charged pions, and thus searches for LJs frequently cover pionic final states as well. Searches to date have targeted LJs containing up to two lepton (pion) pairs.

To further extend coverage into low-mass regimes, it may be of interest here to investigate the samples stored through data scouting for, e.g. di-muon resonances, although the limited information retained in scouted events likely makes this a promising avenue only for sufficiently short lifetimes that dedicated displaced track reconstruction is not necessary.

7.6.2.2. Single species hadronic dark showers (emerging jets). The ATLAS, CMS, and LHCb collaborations have developed powerful searches for pairs of DVs in the trackers, typically in association with missing energy or large H_T , as described in section 3.1. These searches are nearly background-free and inclusive in the number of vertices, and thus have good sensitivity for any model of dark showers that has a large enough signal acceptance in such searches, regardless of the detailed features of the shower shape.

While these searches demonstrate the power and flexibility of inclusive low-background searches, it is unfortunately very easy for signal acceptances in existing DV searches to be prohibitively small. Primary drivers for the loss of acceptance are the requirement of associated objects (leptons, MET, sizable visible H_T , etc), and/or cuts on the invariant mass of or number of tracks belonging to the DV. Given that a hidden shower tends to produce a multitude of LLPs, it should be possible to maintain very low background levels—and therefore the power and inclusivity of the search—by relaxing many of these requirements and demanding a larger number of DV or even displaced tracks instead.

When the lightest decaying dark state has mass $\lesssim 10$ GeV, backgrounds to DV searches do become more important, since the background DV rate rises rapidly as the number of tracks associated to the vertex falls; also, irreducible heavy-flavor backgrounds are important

in this mass range. The number of associated tracks is also crucial for the vertex reconstruction in both the ID and the MS. Hence, the composition of hadronic and leptonic particles in the final states may have a significant impact on the ability to reconstruct the associated vertices and/or the size of the expected backgrounds. Moreover, as the track momentum and multiplicity in the vertices drop, the odds increase that one or more tracks belonging to a particular vertex are not reconstructed. The efficiency of reconstructing displaced tracks in the inner tracker falls off substantially with displacement, as particles traverse fewer layers of the tracker: for particles produced at $r = 30$ cm from the IP, this efficiency is ~ 0.35 at ATLAS [255] and ~ 0.25 at CMS [587]. Explicitly requiring vertex reconstruction can thus come at a large cost in signal acceptance, especially for LLP decays that produce only two tracks per vertex. However, a large number of unassociated tracks with a large impact parameter is still a striking signature, despite the larger backgrounds that come with relaxing requirements on vertex reconstruction. This is the idea behind a recent CMS search [316], based on the model of [331] (see below). Whether or not explicit DV reconstruction is helpful depends on relative signal and background rates, and can be model dependent.

Backgrounds also increase as the lifetime of the decaying dark state becomes shorter; as track reconstruction efficiencies are very good for particles at small production radii, explicit vertex reconstruction is likely to be useful here to help keep backgrounds under control. Even in the short lifetime regime, a dark shower offers many additional handles for signal/background discrimination beyond the number of vertices, such as a common mass scale for reconstructed vertices and non-SM-like particle multiplicity distributions.

Given the striking nature of these high multiplicity signal events, it is typically not challenging to separate them from backgrounds once the event is on tape, provided that sufficiently many displaced tracks (and possibly vertices) can be reconstructed. At ATLAS, reconstructing displaced tracks can require running dedicated re-tracking algorithms in order to identify highly displaced tracks. This re-tracking can be computationally expensive, and necessitates the preselection of at most $\sim 5\%$ of the total event sample on tape. In this case, the preselection criteria are likely to be the limiting factor in signal acceptance at the analysis level. At CMS, the standard tracking algorithm is iterative and automatically reconstructs highly-displaced tracks, so preselection is not necessary.

Dark showers produced through mediators carrying SM charges, such as the scenario of [331], generally provide ample preselection criteria through associated objects and relatively high overall event H_T scales. For instance, typical events for the model in [331] contain two emerging jets and two QCD jets, though the additional two QCD jets can be absent in other models. The emerging jets can be reconstructed using default anti- k_t [366] $R = 0.4$ jets. A baseline preselection requires each jet to have $p_T > 200$ GeV, $|\eta| < 2.5$ and H_T exceeding 1000 GeV. These criteria assume that jets can be reconstructed using the calorimeters, which should be very efficient for the considered lifetimes in the model. CMS has recently published a search for this model [316], finding excellent sensitivity for the benchmark model developed in [331]; this analysis is an excellent demonstration that searches of this type are feasible despite the challenging nature of the signal.

When dark showers originate from a SM singlet, such as the Higgs boson or a Z' , the problem of preselection becomes more acute. The overall mass scale of signal events can easily be small, making event H_T useless for signal separation. Here perhaps one of the most robust avenues for preselection is muon multiplicity: as discussed in section 7.4, many of the operators governing LLP decays tend to predict muon-rich final states. With $\gg 2$ LLPs in an event, muon number becomes a useful and inclusive preselection criterion that places no demands on the possible presence of associated objects, event H_T , or detailed shower shape.

For muon-poor, low-mass dark shower events, pre-selection criteria become more model dependent. When such dark showers originate from an exotic Higgs decay, SM Higgs production provides a suite of associated prompt objects that offer pre-selection handles at some acceptance cost; see section 2.3.1. However it is also important to study to what extent features reflecting the presence of (multiple) LLPs in the event itself could be used as pre-selection criteria, e.g. anomalously track-poor jets, unusual ECAL/HCAL ratios (as realized by, for instance, LLPs with $c\tau \lesssim m$ decaying dominantly to electrons and/or photons), anomalously large numbers of poorly reconstructed or high impact parameter tracks, etc.

Off-line reconstruction for events with low-mass and/or very soft vertices, as occurs in SUEP-style models, poses additional challenges. For the reconstruction of secondary vertices in the ID in ATLAS and CMS, currently a minimum track p_T of 1 GeV is required [237, 588]. Having vertices with several tracks fulfilling this minimum p_T requirement might be rare if the hidden sector hadronization scale is low, and this is particularly true for low mediator masses. In this case, dedicated search strategies might be needed (particularly for short lifetimes). For example, it may be possible to look for increased multiplicity of hits in the outer layers of the ID compared to the inner parts, though unassociated hits from secondaries may be a significant background. A subtraction of hits from tracks stemming from vertices close to the beam pipe might be helpful, but dedicated studies are needed to assess the viability of this approach. In addition, the calorimeter may be another handle on this type of event: a large collection of soft particles could collectively contribute a non-standard energy pattern in the calorimeters provided the LLP lifetime is sufficiently long and if the calorimeter segmentation in radial direction is exploited. However, if the LLPs can reach the calorimeter, their decays typically take place over a distance longer than or comparable to the size of the calorimeter itself. This effect could wash out the signal to the extent that it may be difficult to observe above background, though a full truth-level simulation of this effect is certainly warranted.

Finally, one of the great advantages of an inclusive, low-background search strategy is that it enables a single search to apply robustly to a vast class of models. To fully realize the power and impact of such searches, it will be critical to make it possible to reinterpretable searches with publicly available material. In searches for DVs in a particular detector element, it may be advantageous to bin signal and background in terms of the number of reconstructed vertices with very loose η and ϕ requirements. While more difficult experimentally, it would also be highly valuable to eventually supply the transverse distance of each vertex from the IP on an event-by-event basis, combining the different detector elements. This would allow for straightforward reinterpretation of searches to models with different shower shapes, lifetimes and p_T spectra, allowing a single search to transparently apply to a very broad model space. If possible, it would also be useful to provide the distribution of background vertices as a function of the number of tracks, as this would help theorists estimate the sensitivity to models with different masses and LLP decay modes.

For searches that rely directly on displaced tracks, it is especially important to publish information on displaced track reconstruction that allows for reliable recasting since this is challenging to model accurately using only public data. For instance, the analysis of [316] relies on the transverse impact parameter significance of a track to construct discriminating variables, but theorists cannot reconstruct impact parameter significance without clear examples of the detector response. For these searches, it would again be useful to bin signal and background in terms of the number of reconstructed objects and their geometric location within the detector, as well as their p_T .

We conclude by reiterating that, while inclusive searches are naturally desirable, there are substantial practical obstacles for such a program, some of which have been discussed above.

In reality, it is likely that a number of different semi-inclusive searches with partially overlapping acceptance are necessary. To that end, further theoretical studies are needed to map out which selection cuts are most robust against the varying of the specifics of particular models.

7.6.2.3. Multi-species dark showers. If there are multiple species in the dark shower that decay to SM final states, the lifetimes of these species will in general be very different. The intermediate case where most LLP species decay with a macroscopic lifetime produces similar phenomenology to the single-species models discussed above. Another regime occurs when at least one species decays promptly, while the others are either stable or decay outside of the detector, which gives rise to a semi-visible jet [372] (see section 7.6.1.1 above). The intermediate scenario, where some dark states decay promptly while others have macroscopic detector-scale lifetimes, is equally generic. A combination of semi-visible techniques with DV reconstruction can be used to boost sensitivity in this case. It is especially important to pay attention to potential isolation criteria on DVs, as the presence of potentially large numbers of promptly decaying dark particles may necessitate their relaxation.

7.7. Executive summary

Dark showers are a common prediction of a wide range of hidden sector theories. Although the model space is dauntingly enormous, it is possible to make some very general statements about the *signature* space of interest: dark shower events can be characterized by (1) a variable and frequently large multiplicity of particles per event; (2) BSM states with a hierarchy of proper lifetimes, making the existence of at least one LLP species generic; (3) frequently, non-isolation of LLPs from other objects in the event; and (4) non-SM-like energy flow and particle multiplicity. These features typically ensure that dark shower events are very distinctive signatures in the generic regime where at least one species has a detector-scale lifetime. Thus, it should be possible to design powerful and inclusive searches which would be sensitive to a very large portion of the vast model space without needing to rely on poorly predicted (and model-dependent) properties such as shower shape. Toward this end we have identified several promising directions for future study for both theory and experiment, and provide some recommendations here.

As always, one of the primary challenges in searches for dark showers is ensuring the events are recorded on tape. In some models, dark showers may be accompanied by \cancel{E}_T , H_T , or a number of associated leptons that can be triggered on, but there are many scenarios where this is not the case, or where these handles come at the expense of a large reduction in signal rate. The foremost example of this is when the dark shower is initiated by an exotic Higgs decay. It is thus critical to pursue dedicated trigger strategies, and we present a few ideas here. We expect that displaced triggers designed for singly- or doubly-produced LLPs will typically have reasonable acceptance for dark shower events, though we caution that in some cases the non-isolation of LLPs in dark shower events may limit their acceptance. Another promising avenue for triggering on dark showers exploits the high particle multiplicity typical of such events, and we recommend study of triggers on (displaced) multi-muons in particular.

The off-line analysis of dark shower events poses several challenges as well. At this time, only the limiting cases of pencil-like or fully spherical showers are under good theoretical control, and hadronization can introduce significant additional uncertainties, especially for models with a spectrum substantially different from QCD. New studies of showering in the intermediate regime were performed for the purpose of this document (presented in section 7.3), which revealed that different approaches can yield qualitatively different

phenomenology. On the one hand, this provides interesting opportunities for searches to make use of event shape and/or jet substructure variables, in particular if the decays of hidden sector states occur promptly or with small displacements. On the other hand, for larger displacements it implies that one should be careful not to heavily bias the selection choices of a search towards a particular shower shape, and rely instead on the displaced LLP decays to separate signal from background. We expect that the most inclusive, most broadly applicable, and most readily reinterpreted searches will be those in which the data is binned in terms of the number of reconstructed displaced objects and their locations within the detector. For recasting purposes, and to assist with unraveling the underlying physics in the event of a discovery, it would moreover be important to supply information concerning distributions of the number of tracks per vertex and/or the vertex mass whenever possible.

The presence of $N_{\text{DV}} \gg 2$ DV in an event will be enormously powerful for background suppression, provided those vertices can be reconstructed. To identify highly displaced tracks, it may be necessary to run dedicated and computationally expensive re-tracking algorithms, requiring the imposition of some pre-selection criteria to identify events of interest. As these pre-selection criteria are likely to be the limiting factor controlling post-trigger signal acceptance in many models, we recommend developing criteria for dark showers that rely on particle multiplicity and, if possible, the presence of multiple LLPs in the event, while keeping p_{T} thresholds as low as possible.

7.8. Appendix: example models

In this appendix we survey the models for which MC simulations are currently available. Please also see the [appendix](#) as some of the models are included in simplified model library.

7.8.1. Lepton-jets from radiating DM. One of the simplest and most widely discussed types of low-mass hidden sector particle is a dark photon A' , i.e. a new gauge boson associated with a local $U(1)$ symmetry in the dark sector. By kinetically mixing with the SM hypercharge gauge boson, a dark photon can act as the mediator of DM–SM interactions, in addition to being responsible for DM self-interactions. The range of dark photon masses of relevance to dark showers is between MeV and GeV. The dark sector Lagrangian in such a scenario reads

$$\mathcal{L}_{\text{dark}} \equiv \bar{\chi}(i\cancel{D} - m_{\chi} + ig_{A'}\cancel{A}')\chi - \frac{1}{4}F'_{\mu\nu}F'^{\mu\nu} + \frac{1}{2}m_{A'}^2A'_{\mu}A'^{\mu} - \frac{\epsilon}{2}F'_{\mu\nu}F^{\mu\nu}. \quad (7.8)$$

Here, χ is the fermionic DM particle with mass m_{χ} , and $g_{A'} = \sqrt{4\pi\alpha'}$ is the $U(1)'$ gauge coupling. From the point of view of dark showers, interesting values for the coupling strength are $\alpha' \gtrsim 0.01$. If α' is much smaller, there is too little radiation to form a dark shower. The dark photon mass is denoted by $m_{A'}$, and the kinetic mixing parameter by ϵ . Typically, ϵ is constrained by current limits to be $\lesssim 10^{-3}$. The particular mechanism by which the dark photon mass arises is not important in this context—it could originate from a dark sector Higgs mechanism or from the Stückelberg mechanism.

If $m_{\chi} \ll 100$ GeV, DM particles may be produced at the LHC with a large boost. This entails a large probability for radiating additional collinear A' bosons (see figure 96). A detailed analytical and numerical description of such dark photon showers is presented in [589]. The A' bosons eventually decay to observable SM particles through the kinetic mixing term in equation (7.8). Depending on the value of ϵ , the decays can be either prompt or displaced. Phenomenologically, the final state of the process $pp \rightarrow \bar{\chi}\chi + nA'$ thus consists of two ‘jets’ of collimated A' decay products, plus missing energy.




Figure 96. The process that gives radiating dark matter its name: production of two DM particles χ , followed by the emission of several soft or collinear dark photons A' [589].




Figure 97. Branching ratios of a kinetically mixed dark photon A' as a function of $m_{A'}$ [589].

The A' branching ratios into different SM final states depend sensitively on $m_{A'}$ (see figure 97). At $m_{A'} \lesssim 400$ MeV, the dominant decay modes are $A' \rightarrow e^+e^-$ and $A' \rightarrow \mu^+\mu^-$. The decay rate into each lepton flavor ℓ is

$$\Gamma(A' \rightarrow \ell^+ \ell^-) = \frac{1}{3} \alpha \epsilon^2 m_{A'} = \frac{1}{8 \times 10^{-6} \text{ cm}} \left(\frac{\epsilon}{10^{-3}} \right)^2 \left(\frac{m_{A'}}{\text{GeV}} \right), \quad (7.9)$$

where $m_\ell \ll m_{A'}$ has been assumed for simplicity. For purely leptonic decays, the A' shower thus corresponds to a ‘LJ’, i.e. a set of collimated leptons. This signature has been previously discussed, for instance, in [30, 32, 117, 127, 184, 572, 590–593]. Experimental searches for LJs have been presented in [270, 272, 273, 274], where [272, 510] focus on LJs with DV.

Lepton-jet searches may also be sensitive to dark photons with masses $\gtrsim 400$ MeV, even though the leptonic branching ratio is reduced to between 20% and 70% in this regime. The mix of leptons and hadrons expected from A' decays at $m_{A'} \gtrsim 400$ MeV implies, however, that most leptons will not be isolated, but occur in conjunction with hadronic activity in the same detector region. In view of this, dedicated trigger and analysis strategies may significantly boost the sensitivity (see section 7.6.2.1).

In certain parameter regions, the decays of radiated A' bosons may closely resemble a purely hadronic QCD jet. This will happen in particular when $m_{A'}$ is close to a QCD resonance, or when the average A' multiplicity in each shower is low, and the hadronic

branching ratio is sizeable. In this case, separation of the signal from the QCD background will most likely be possible only if ϵ is so small that A' decays are displaced.

Let us summarize several important considerations to take into account when devising a search for dark sector radiation:

- *There may or may not be a signal in the tracking detector.* Prompt A' decays will typically lead to such signals, except for specific A' decay modes to neutral particles, for instance $A' \rightarrow K^0 \bar{K}^0$ and $A' \rightarrow \pi^0 \gamma$. Displaced A' decays can also leave a signal in the tracker, however if the lifetime is sufficiently long more decays will occur in the muon chamber (see section 7.4.2).
- *There may or may not be a signal in the calorimeters.* Most A' decay modes will be visible to the calorimeters. However, it is important to realize that a signal in the HCAL can arise not only from hadronic activity, but also from displaced decays to leptons occurring inside the calorimeter.
- *There will be missing energy contained within the LJ.* As the decaying A' bosons are aligned with the DM particle from which they were radiated, the corresponding missing momentum vector points in the same direction. However, unless there is significant ISR, the missing momentum will typically be balanced between the two showers shown in figure 96, and may therefore be small.

In figure 98, we show for illustration the exclusion limits that past LHC searches place on the dark photon parameters.

7.8.2. Emerging and semi-visible jets. Hidden valley models [66] with QCD-like hidden sectors allow for interesting collider signatures. Thus we consider a confined dark gauge group $SU(N_d)$, where $N_d \geq 2$, with confinement scale Λ_d which sets the mass of the dark hadrons. There are also n_f flavors of dark quarks whose bare masses are lighter than Λ_d . The hidden valley comes equipped with a portal that couples the dark sector to the SM, and the mass is usually taken to be $M \gg \Lambda_d$. The portal can be an s -channel vector mediator, Z_d , which couples to SM quarks and dark quarks:

$$\mathcal{L} \supset -Z_{d,\mu} \sum_{i,a} (g_q \bar{q}_i \gamma^\mu q_i + g_{q_d} \bar{q}_{d,a} \gamma^\mu q_{d,a}). \quad (7.10)$$

Here g_{q/q_d} are coupling constants and i, a are flavor indices. One can also have a t -channel scalar bifundamental mediator, X , which carries color and dark color and can decay to a quark and a dark quark. In the case of the scalar mediator, the only allowable coupling is of the form

$$\mathcal{L} \supset \kappa_{ij} \bar{q}_i q_{d,j} X + \text{h.c.}, \quad (7.11)$$

where κ_{ij} is a $3 \times n_f$ matrix of Yukawa couplings. One could also add multiple flavors of X mediators, something that has also been implemented [542].

When dark quarks are produced, they shower and hadronize and the same tools that are familiar for QCD can be used to simulate these processes. Because of the large gap between the mediator mass and the confining scale, there will be large particle multiplicity and the dark hadrons will typically form into jet-like structures. In the large N_d limit, the fraction of dark baryons produced is suppressed, and in the case of QCD this fraction is $\mathcal{O}(0.1)$. Therefore the hadronization in these simulations is typically dominated by dark mesons. The lightest hadronic states are the dark pions π_d , acting as goldstone bosons of the $U(n_f) \times U(n_f)$ dark flavor symmetry. When a heavier mesonic state is produced it will promptly decay into dark pions if kinematically allowed, making the dark pions the dominant component of the dark showering process.




Figure 98. Limits, at the 95% C.L., on the dark photon mass $m_{A'}$ and the kinetic mixing parameter ϵ , for $m_\chi = 4$ GeV, $\alpha_{A'} = 0.2$. We have assumed χ production through a Z' portal with a mass of 1 TeV, with a production cross section of 0.58 pb at $\sqrt{s} = 7$ TeV and 0.85 pb at 8 TeV. We show exclusion limits from the ATLAS search for prompt lepton-jets in 5 fb^{-1} of 7 TeV data [594] (blue shaded region) and from the ATLAS displaced lepton-jet search in 20.3 fb^{-1} of 8 TeV data [272] (red shaded region). The lighter colored region around $m_{A'} = 2$ GeV corresponds to the transition region between an analysis in terms of hadron final states and an analysis in terms of quark final states and is based on interpolation. The computation was carried out in Pythia 8 [220, 595, 596]; see [589] for details. We also show the existing 90% C.L. exclusion limits from the electron and muon anomalous magnetic moment [19, 597, 598], HADES [599], KLOE 2013 [600] and 2014 [601], the test run results from APEX [602], BaBar 2009 [603] and 2014 [604], beam dump experiments E137, E141, and E774 [605–607], A1 [608], Orsay [609], U70 [610], CHARM [611], LSND [612], as well as constraints from astrophysical observations [613, 614] and π^0 decays [615]. Figure based on [589].

These events can be simulated with the HV [595] module/package that appears in Pythia8 [220]. Pythia8 hosts a hidden valley class which incorporates the $SU(N_d)$ model, allowing the user to vary the masses of the spectrum (π_d , ρ_d , etc), number of flavors n_f and parameters of the running coupling (Version ≥ 8.226). Pair production of X and resonant production of Z_d are implemented at tree level, and the decay of dark mesons to SM states is also present.

In [574], production of the heavy mediators with ISR/FSR was considered. This was done by interfacing with Madgraph5_aMC@NLO [226] using a modified version of the spin-1 DMsimp model²⁰² implemented through FeynRules [616]. The models are located in the repository²⁰³ folder DMsimp_s_spin1. The generation files for the t -channel exchange of the scalar X are located in the folder DMsimp_tchannel. The bi-fundamentals are denoted with s_{u11} , s_{u12} , s_{u21} , s_{u22} ..., where u explicitly specifies the QCD flavor index and the numbers are the explicit dark non-abelian group indices. Similarly, the dark quarks are labeled as q_{v11} , q_{v12} , q_{v21} , q_{v22} . A FeynRules model file (DMsimp_tchannel.fr) as well as the Mathematica notebook (DMsimp_tchannel.nb) used to generate the

²⁰² <http://feynrules.irmp.ucl.ac.be/wiki/DMsimp>

²⁰³ <https://github.com/smsharma/SemivisibleJets>

UFO output are also provided. The showering and hadronization in the dark sector can still be performed in `Pythia8`.

7.8.2.1. Semi-visible jets. Generically, some of the dark pions could decay promptly while some could be long-lived or even collider-stable, analogous to the appearance of neutral versus charged pions in the SM. If both stable and unstable hadrons are produced in a collision, the missing energy could be aligned along one of the jets, resulting in low acceptance for traditional monojet-style searches. This is the *semi-visible jet* scenario, and [372, 574, 617] introduce a simplified model-like parameterization to map the complicated dynamics of the underlying dark sector onto a limited number of physically-motivated variables—in particular, the fraction of stable versus decaying pions, the characteristic mass scale of dark pions and the dark coupling strength. The MC production described here along with the `Pythia8` HV module allows the user to vary these parameters for the *s*- and *t*-channel UV completions. See [372, 574, 617] for further details.

7.8.2.1.1. Emerging jets By contrast, the dark pions can be taken to have detector-scale lifetimes in the *emerging jets* scenario [331]. The expected lifetime of the dark pion can be quantified using equation (7.11). Under the assumption of universal couplings $\kappa_{ij} = \kappa$ and $m_q > \Lambda_d$, we can calculate the proper lifetime of the dark pions:

$$c\tau_0 \approx 80 \text{ mm} \times \frac{1}{\kappa} \times \left(\frac{2 \text{ GeV}}{f_{\pi_d}} \right)^2 \left(\frac{100 \text{ MeV}}{m_q} \right) \left(\frac{2 \text{ GeV}}{m_{\pi_d}} \right) \left(\frac{M_X}{1 \text{ TeV}} \right)^4. \quad (7.12)$$

Here f_{π_d} is the dark pion decay constant, and m_q is the mass of the SM quark in the final state. A similar formula applies for the Z_d mediator. Therefore a jet of dark hadrons will be created as dominantly invisible particles, but at long distance, the dark pions will decay back to SM particles and appear in shape like an ordinary jet, although it contains a large number of DV. The jet emerges as it travels through the detectors. In the case of purely *t*-channel interactions, having a non-trivial κ_{ij} in equation (7.11) will break the $U(n_f) \times U(n_f)$ dark flavor symmetry. With an appropriate n_f , the dark quark flavors can be exactly aligned with the SM down quark flavors. It is immediately clear that such alignment leads to dark pions with lifetimes that differ significantly from one another. In return, the flavor composition of the emerging jets will vary throughout the detector volume. This differs from the *s*-channel interaction, in which no breaking occurs, and n_f -stable versus $n_f(n_f - 1)$ -unstable dark pions exist.

7.8.3. SUEPs. In the strongly coupled regime, dark showers can be become spherical and almost arbitrarily soft, leading to soft, unclustered energy patterns, or SUEPs, discussed earlier in this chapter. In this case the sheer multiplicity of the final states becomes the most important experimental handle. It has been shown through the AdS/CFT correspondence that their spectrum should follow an approximately thermal distribution [618], which is what was assumed in the phenomenological study in [242]. The production mechanism for these studies was the Higgs or a heavy Higgs-like scalar. The MC code that was written for this study is currently not available publicly, though events in `hepmc` format can be obtained by contacting the authors. At this time only leptonic decays of the hidden mesons are implemented.

8. Conclusions

The research program carried out over the first nine years of the LHC at CERN has been an unqualified success. The discovery, in 2012, at a center-of-mass energy of 7 and 8 TeV, of a new particle thus far consistent with the SM Higgs boson has opened numerous new research

directions and has begun to shed light upon the source of electroweak symmetry breaking, vector boson scattering amplitudes, and the origin of particle masses. And the establishment of a wide range of searches for new physics at 7, 8, and 13 TeV with the ATLAS, CMS, and LHCb detectors—searches thus far consistent with SM expectations—has inspired new ideas and thinking about the most prominent open issues of physics, such as the nature of DM, the hierarchy problem, neutrino masses, and the possible existence of supersymmetry.

The overwhelming majority of searches for new physics have been performed under the assumption that the new particles decay promptly, i.e. very close to the proton–proton IP, leading to well-defined objects such as jets, leptons, photons, and missing transverse momentum. Such objects are constructed requiring information from all parts of the detector including hits close to the IP, calorimeter deposits known to be signatures of particles originating from the IP, and muons with tracks that traverse the entirety of the detector, moving out from the IP. However, given the large range of particle lifetimes in the SM—resulting from general concepts such as approximately preserved symmetries, scale hierarchies, or phase space restrictions—and the lack of clear, objective motivation related to any particular model or theory BSM, the lifetime of hypothetical new particles is best treated as a free parameter. This leads to a wide variety of spectacular signatures in the LHC detectors that would evade prompt searches, and which have received modest attention compared to searches for promptly decaying new particles. Because such signatures require significantly customized analysis techniques and are usually performed by a smaller number of physicists working on the experimental collaborations, a comprehensive overview and critical review of BSM LLPs at the LHC has been performed by a community of experimentalists, theorists, and phenomenologists. This effort ensures that such avenues of the possible discovery of new physics at the LHC are not overlooked. The results of this initiative have been presented in the current document.

We developed a set of simplified models and tools, in chapter 2, that can be used to parametrize the space of LLP signatures. The simplified models were organized around generic ways that various BSM LLPs can be produced and decay to displaced or non-standard objects in the LHC detectors, rather than emphasizing any one particular theory or physics motivation. These can serve as a useful grammar by which to compare coverage of LLP signature space and model classes among current and future experiments.

To that end, in chapter 3, we utilized these models and tools to assess the coverage of current LLP searches and we identified multiple avenues for improving and extending the existing LLP search program. Opportunities for new and improved triggering strategies, searches, and open questions for the experimental collaborations to explore centrally were presented as a list at the end of the chapter.

Moreover, due to the non-standard nature of LLP searches, many of them are performed under very low-background conditions. As a result, sources of backgrounds largely irrelevant to searches for promptly decaying BSM particles are important for LLP analyses and can be surprising and unexpected. In chapter 4 we discussed several sources of backgrounds for LLP searches, collecting the knowledge gained, often by trial-and-error, by experimentalists over many years of searches.

Also with an eye to the future, in chapter 5 we explored the potential for the expanded capabilities of proposed detector upgrades at ATLAS, CMS, LHCb, and related dedicated detectors, highlighting areas where new technologies can have a large impact on sensitivity to LLP signatures and suggesting several studies to be performed by the collaborations to ensure new physics potential is not missed for the upcoming era of the HL LHC.

Additionally, to ensure that current and future searches can be maximally useful in the future, in chapter 6 we explored how current searches can apply to new models and performed

a comprehensive overview of some of the challenges and pitfalls inherent in attempting to recast existing LLP analyses, leading to recommendations for the presentation of search results in the future.

Finally, in chapter 7, we looked toward the newest frontiers of LLP searches, namely high-multiplicity or ‘dark shower’ LLP signals that can, for example, be signatures of complex hidden sectors with strong dynamics and internal hadronization. In this chapter, we elaborated on the theoretical and experimental challenges and opportunities in expanding the LHC reach of these signals. Such dark shower signatures have a high potential for being overlooked with existing triggering strategies and analysis techniques. Moreover, the dark-QCD-like theoretical models from which they can arise are currently being explored in depth and the resulting LHC phenomenology is in the process of being understood. We discussed the current state of this work and we anticipate exciting independent developments in the near future.

This document is incomplete by design, since it is a record of the critical thinking and examination of the state of LLP signatures by a large number of independently organized members of the LHC LLP Community as it has evolved from 2016 to 2019, and a major component of the work has been the identification of several open questions and opportunities for discovery in such signatures. As these questions are addressed and new searches emerge from the experimental collaborations, so, too, will new ideas emerge and evolve from the community. We expect this document to be followed by future papers to record, review, and summarize the evolution of LLP signatures and searches, always with the intention of more effectively facilitating the discovery of new particles at the LHC and beyond.

Acknowledgments

We are very grateful to Andreas Albert, Tao Huang, Laura Jeanty, Joachim Kopp, Matt LeBlanc, Larry Lee, Haolin Li, Simone Pagan Griso, Michele Papucci, Matt Strassler, and Tien-Tien Yu for helpful conversations and comments on the draft.

Editors:

J Beacham acknowledges support from the US Department of Energy (DOE) and the US National Science Foundation (NSF). G Cottin acknowledges support from the Ministry of Science and Technology of Taiwan (MOST) under Grant No. MOST-107-2811-M-002-3120 and CONICYT-Chile FONDECYT Grant No. 3190051. N Desai was supported in part by the OCEVU Labex (ANR-11-LABX-0060) and the A *MIDEX project (ANR-11-IDEX-0001-02) funded by the ‘Investissements d’Avenir’ French government program managed by the ANR. J A Evans is supported by DOE grant DE-SC0011784. S Knapen is supported by DOE grant DE-SC0009988. A Lessa is supported by the Sao Paulo Research Foundation (FAPESP), project 2015/20570-1. Z Liu is supported in part by the NSF under Grant No. PHY-1620074, and by the Maryland Center for Fundamental Physics. S Mehlhase is supported by the BMBF, Germany. M J Ramsey-Musolf is supported by DOE grant DE-SC0011095. The work of P Schwaller has been supported by the Cluster of Excellence ‘Precision Physics, Fundamental Interactions, and Structure of Matter’ (PRISMA+ EXC 2118/1) funded by the German Research Foundation (DFG) within the German Excellence Strategy (Project ID 39083149). The work of J Shelton is supported in part by DOE under grant DE-SC0017840. The work of B Shuve is supported by NSF under grant PHY-1820770. X C Vidal is supported by MINECO through the Ramón y Cajal program RYC-2016-20073 and by XuntaGal under the ED431F 2018/01 project.

Contributors:

M Adersberger is supported by the BMBF, Germany. The work of C Alpigiani, A Kvam, E Torro-Pastor, M Profit, and G Watts is supported in part by the NSF. Y Cui is supported in part by DOE Grant DE-SC0008541. J L Feng is supported in part by Simons Investigator Award #376204 and by NSF Grant No. PHY-1620638. I. Galon is supported by DOE Grant DE-SC0010008. K Hahn is supported by DOE Grant DE-SC0015973. J Heisig acknowledges support from the F R S-FNRS, of which he is a postdoctoral researcher. The work of F Kling was supported by NSF under Grant No. PHY-1620638. H. Lubatti thanks the NSF for support. P Mermel was supported by Grant PP00P2_150583 of the Swiss National Science Foundation. S Mishra-Sharma is partially supported by the NSF CAREER Grant PHY-1554858 and NSF Grant PHY-1620727. V Mitsou acknowledges support by the Generalitat Valenciana (GV) through MoEDAL-supporting agreements and the GV Excellence Project PROMETEO-II/2017/033, by the Spanish MINECO under the project FPA2015-65652-C4-1-R, by the Severo Ochoa Excellence Centre Project SEV-2014-0398 and by a 2017 Leonardo Grant for Researchers and Cultural Creators, BBVA Foundation. J Prisciandaro is supported by funding from FNRS. M Reece is supported by DOE Grant DE-SC0013607. D Robinson is supported in part by NSF grant PHY-1720252. D Stolarski is supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC). S Trojanowski is supported by Lancaster-Manchester-Sheffield Consortium for Fundamental Physics under STFC grant ST/L000520/1. S Xie is supported by the California Institute of Technology High Energy Physics Contract DE-SC0011925 with the DOE. This manuscript has been partially authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the US Department of Energy, Office of Science, Office of High Energy Physics.

Appendix. Simplified model library

Chapter Editor: Brian Shuve

Contributors: Eshwen Bhal, David Curtin, Alessandro Davoli, Andrea De Simone, Jared Evans, Thomas Jacques, Zhen Liu, Siddharth Mishra-Sharma, Alessandro Morandini, Michael Ramsey-Musolf, Jessie Shelton, Jiang-Hao Yu.

A.1. Instructions for the Simplified Model Library

The simplified model library is available at the LHC LLP Community website²⁰⁴, hosted at CERN. We refer here extensively to the simplified models in tables 1–3. Because it is already quite an extensive task to come up with simplified models for so many (production)×(decay) modes, we for now restrict ourselves predominantly to the ‘filled’ entries in tables 1–3. If you are interested in performing an experimental search or developing a simplified model library entry for one of the ‘unfilled’ entries, please contact the chapter editor.

There are essentially two possible pipelines to simulate LLP events with the library:

1. *LLP decay as part of matrix-element calculation:* Using UFO models in the library, it is possible to generate the production and decay of LLPs at the parton level using calculations of the matrix element for production and decay. As a concrete example, we provide cards that allow the production and decay of LLPs using MadGraph5_aMC@NLO [226] and the accompanying MadSpin [227] package. This employs the narrow-width approximation, but otherwise gives rise to the correct angular

²⁰⁴ <http://cern.ch/longlivedparticles>

distribution of LLP decay products. The downside is that if a particular decay is not allowed from the interactions in the UFO model file, the UFO must be modified to include the new coupling. The output of `MadGraph5_aMC@NLO` is then fed into programs such as `Pythia 8` [219, 220] for showering, implementation of underlying event and other particle-level processes.

2. *Phenomenological LLP decays*: Using UFO models in the library, it is possible to generate the production of LLPs, leaving them stable as outputs of the matrix-element-level calculation recorded in LHE format [224]. The LLP can then be subsequently decayed in programs such as `Pythia 8`, which allows a particle to decay into any final state, albeit without correctly modeling the angular distribution. This could be convenient for models where the interactions leading to LLP decay are not included in the UFO, or where computational time is a concern and the angular distribution of LLP decays is irrelevant. We provide detailed instructions in the library files for how to implement decays of LHE files via `Pythia`.

In the final version of the library, we aim to provide example cards to direct the production and decay of LLPs in both pipelines. Note that in all of the simplified model proposals below, any particles *not* present in the production or decay chain should have their masses set to a very large value ($M \gtrsim 5$ TeV) to ensure they are sufficiently decoupled from direct production at the LHC.

Currently, we only provide simplified model libraries for neutral LLPs. The simplified models for LLPs with electric or color charges are equally compelling, but their simulation is more subtle. In particular, the simulation of the propagation and decay of the LLPs are more challenging if the LLP carries a SM gauge charge. Such effects can be included in detector simulations using `GEANT4` [231], but decay processes typically need to be hard-coded into `GEANT` or otherwise interfaced with other MC programs. This is an important issue for the community to address, as discussed in sections 2.4.2 and 2.4.3. Without implementing the decays, it is straightforward to use the SUSY model to simulate the production of any of the electrically or color charged LLPs.

Finally, we note that there is not currently a set of minimal simplified models to cover dark showers. However, we are collecting models used in studies of dark showers that could potentially be helpful for experimentalists and theorists alike. The existing models are included in the library, along with a very brief description in section A.1.2.

A.1.1. Neutral LLPs. The instructions for simulating the simplified model channels for neutral LLPs are given below. Note that it is often true that the same simplified model production and decay channel can be simulated using several simplified models. However, the philosophy of the simplified model approach is that the UV model used to simulate the process is not important when sensitivity is expressed in terms of physical masses and cross sections. As a result, we typically provide only one set of instructions for simulating each simplified model channel.

We begin by presenting the simplified model library instructions for double pair production (DPP) in table A1, Heavy Parent (HP, QCD-charged parent) in table A2, and Heavy Parent (HP, EW-charged parent) in table A3. We then proceed to the Higgs (HIG) production modes in tables A4–A6. For the Z' (ZP) production modes, we use a set of simplified models described in tables A7–A9. A relatively simple model file is provided for each table. In addition, a more adjustable ‘advanced’ model file is provided which includes all ZP production and decay modes, and allows for features such as individual couplings to each generation of quarks. This comes at the cost of a greatly increased set of parameters, and the

Table A1. Simplified model library process proposals for double pair production (DPP) production mode. Where a ‘wino’ LSP is specified, an admixture of Higgsino is required to lead to direct pair production of the neutral wino component. As an alternative, one could have $pp \rightarrow \tilde{\chi}^\pm \tilde{\chi}^0$, $\tilde{\chi}^\pm \rightarrow W^\pm \tilde{\chi}^0$ promptly, and take the $\tilde{\chi}^\pm$ to be degenerate with $\tilde{\chi}^0$ such that the additional charged decay products are essentially unobservable.

Decay mode	Simplified model library process
$X \rightarrow \gamma + \text{inv.}$	MSSM+GMSB. LLP is a bino ($\tilde{\chi}^0$) produced due to $pp \rightarrow \tilde{\chi}^0 \tilde{\chi}^0$ via t -channel squark exchange ($M_{\tilde{q}} > 5$ TeV). Bino decays to photon + gravitino, $\tilde{\chi}^0 \rightarrow \gamma + \tilde{G}$
$X \rightarrow jj$	MSSM+RPV. LLP is sneutrino LSP ($\tilde{\nu}$) that is pair-produced via weak gauge interactions. $\tilde{\nu} \rightarrow q\bar{q}$ via the $Q\bar{d}^c$ operator
$X \rightarrow jj + \text{inv.}$	MSSM. LLP is second neutralino (wino) LSP $\tilde{\chi}_2^0$ that is pair-produced via weak gauge interactions. $\tilde{\chi}_2^0 \rightarrow q\bar{q}\tilde{\chi}_1^0$ via an off-shell sfermion, and the $\tilde{\chi}_1^0$ is invisible with arbitrary mass
$X \rightarrow jjj$	MSSM+RPV. While this is partially covered by $jj + \text{inv.}$ in the case where the additional quark is not reconstructed, we include it here for completeness. LLP is wino LSP ($\tilde{\chi}^0$) that is pair-produced via weak interactions. $\tilde{\chi}^0 \rightarrow q_\alpha q_\alpha q_\beta$ via an off-shell sfermion and the $u_\alpha^c d_\alpha^c d_\beta^c$ operator
$X \rightarrow jj\ell_\alpha$	MSSM+RPV. LLP is wino LSP ($\tilde{\chi}^0$) that is pair-produced via weak interactions $\tilde{\chi}^0 \rightarrow \ell_\alpha q\bar{q}$ via an off-shell sfermion and $L_\alpha Q\bar{d}^c$ operator.
$X \rightarrow \ell_\alpha^+ \ell_\alpha^-$	MSSM+RPV. LLP is sneutrino $\tilde{\nu}_\beta$ of flavor β that is pair-produced via weak interactions $\tilde{\nu}_\beta \rightarrow \ell_\alpha^+ \ell_\alpha^-$ via the $L_\alpha L_\beta E_\alpha^c$ operator
$X \rightarrow \ell_\alpha^+ \ell_\alpha^- (+\text{inv.})$	MSSM. LLP is second neutralino $\tilde{\chi}_2^0$ that is pair-produced via weak interactions $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 \ell_\alpha^+ \ell_\alpha^-$ via an off-shell slepton
$X \rightarrow \ell_\alpha^+ \ell_\beta^- (+\text{inv.})$	MSSM+RPV. LLP is sneutrino $\tilde{\nu}_\alpha$ of flavor α that is pair-produced via weak interactions $\tilde{\nu}_\alpha \rightarrow \ell_\alpha^+ \ell_\beta^-$ via the $L_\alpha L_\beta E_\alpha^c$ operator. An additional massless invisible final state can be obtained with a wino LLP decaying into $\ell_\alpha^+ \ell_\beta^- \nu_\alpha$ through the same operator and an off-shell slepton. The massive invisible case is less motivated for $\alpha \neq \beta$

possibility of including unwanted diagrams if the process is not carefully specified. A simple python script is provided to generate the processes and set unwanted parameters to zero for those users wishing to use the advanced model files.

Finally, we provide instructions for the CC production modes in table A10. This production mode is most easily simulated using a left-right symmetric model or other right-handed-neutrino model.

A.1.2. Dark shower models. In the repository, we also include two models of dark showers. These are not meant to be representative in the same way that the simplified models are;

Table A2. Simplified model library process proposals for heavy parent (HP) production mode where the parent particle carries a QCD charge. In most of the above cases, a squark parent can be replaced by a gluino parent with an additional jet in its decay.

Decay mode	Simplified model library process
$X \rightarrow \gamma + \text{inv.}$	MSSM+GMSB. LLP is a bino ($\tilde{\chi}$) produced via $pp \rightarrow \tilde{q}\tilde{q}^*$, $\tilde{q} \rightarrow \tilde{\chi} + q$. Bino decays to photon+ gravitino, $\tilde{\chi} \rightarrow \gamma + \tilde{G}$
$X \rightarrow jj + \text{inv.}$	MSSM. LLP is wino LSP $\tilde{\chi}_2^0$ that is produced via $pp \rightarrow \tilde{q}\tilde{q}^*$, $\tilde{q} \rightarrow q\tilde{\chi}_2^0$. Then, $\tilde{\chi}_2^0 \rightarrow q\tilde{q}\tilde{\chi}_1^0$ via an off-shell quark
$X \rightarrow jjj$	MSSM+RPV. LLP is bino LSP $\tilde{\chi}$ that is produced via $pp \rightarrow \tilde{q}\tilde{q}^*$, $\tilde{q} \rightarrow q\tilde{\chi}$. Then, $\tilde{\chi} \rightarrow q_\alpha q_\alpha q_\beta$ via the $u_\alpha^c d_\alpha^c d_\beta^c$ operator
$X \rightarrow jj\ell_\alpha$	MSSM+RPV. LLP is bino LSP ($\tilde{\chi}$) that is produced via $pp \rightarrow \tilde{q}\tilde{q}^*$, $\tilde{q} \rightarrow q\tilde{\chi}$. $\tilde{\chi} \rightarrow \ell_\alpha q\bar{q}$ via an off-shell sfermion and $L_\alpha Qd^c$ operator
$X \rightarrow \ell_\alpha^+ \ell_\alpha^-$ or $\ell_\alpha^+ \ell_\beta^-$	MSSM+RPV. LLP is sneutrino ($\tilde{\nu}$) that is produced via $pp \rightarrow \tilde{g}\tilde{g}$, $\tilde{g} \rightarrow jj\tilde{\chi}$, $\tilde{\chi} \rightarrow \tilde{\nu}\bar{\nu}$. Then, $\tilde{\nu}_\alpha \rightarrow \ell_\alpha^+ \ell_\beta^-$ or $\tilde{\nu}_\beta \rightarrow \ell_\alpha^+ \ell_\alpha^-$ via the $L_\alpha L_\beta E_\alpha^c$ operator
$X \rightarrow \ell_\alpha^+ \ell_\alpha^- + \text{inv.}$	MSSM. LLP is second neutralino ($\tilde{\chi}_2^0$) that is produced via $pp \rightarrow \tilde{q}\tilde{q}^*$, $\tilde{q} \rightarrow q\tilde{\chi}_2^0$. Then, $\tilde{\chi}_2^0 \rightarrow \ell_\alpha^+ \ell_\alpha^- \tilde{\chi}_1^0$
$X \rightarrow \ell_\alpha^+ \ell_\beta^- + \text{inv.}$	MSSM+RPV. LLP is bino ($\tilde{\chi}^0$) that is produced via $pp \rightarrow \tilde{q}\tilde{q}^*$, $\tilde{q} \rightarrow q\tilde{\chi}^0$. Then, $\tilde{\chi}^0 \rightarrow \ell_\alpha^+ \ell_\beta^- \nu_\alpha$ via $L_\alpha L_\beta E_\alpha^c$ operator and off-shell slepton (massless invisible only)

Table A3. Simplified model library process proposals for heavy parent (HP) production mode where the parent particle carries electroweak charge.

Decay mode	Simplified model library process
$X \rightarrow \gamma + \text{inv.}$	MSSM+GMSB. LLP is a bino ($\tilde{\chi}^0$) produced via $pp \rightarrow \tilde{\chi}^+ \tilde{\chi}^-, \tilde{\chi}^+ \rightarrow W^+ \tilde{\chi}^0$ ($\tilde{\chi}^+$ is a wino). Bino decays to photon+ gravitino, $\tilde{\chi} \rightarrow \gamma + \tilde{G}$
$X \rightarrow jj$	MSSM+RPV. LLP is sneutrino LSP ($\tilde{\nu}$) produced via $pp \rightarrow \tilde{\chi}^+ \tilde{\chi}^-, \tilde{\chi}^+ \rightarrow \tilde{\nu} \ell^+$. The sneutrino decays via $\tilde{\nu} \rightarrow q\bar{q}$ via the $u_\alpha^c d_\alpha^c d_\beta^c$ operator
$X \rightarrow jj + \text{inv.}$	MSSM. LLP is wino $\tilde{\chi}_2^0$ that is produced via $pp \rightarrow \tilde{\chi}^+ \tilde{\chi}^-, \tilde{\chi}_2^+ \rightarrow W^+ \tilde{\chi}_2^0$. Then, $\tilde{\chi}_2^0 \rightarrow q\bar{q}\tilde{\chi}_1^0$ via an off-shell quark
$X \rightarrow jjj$	MSSM+RPV. LLP is bino LSP ($\tilde{\chi}^0$) that is produced via $pp \rightarrow \tilde{\chi}^+ \tilde{\chi}^-, \tilde{\chi}^+ \rightarrow W^+ \tilde{\chi}^0$. Then, $\tilde{\chi}^0 \rightarrow q\bar{q}q$ via an off-shell sfermion and the $u^c d^c d^c$ operator
$X \rightarrow jj\ell_\alpha$	MSSM+RPV. LLP is bino LSP ($\tilde{\chi}^0$) that is produced via $pp \rightarrow \tilde{\chi}^+ \tilde{\chi}^-, \tilde{\chi}^+ \rightarrow W^+ \tilde{\chi}^0$. Then, $\tilde{\chi}^0 \rightarrow q\bar{q}\ell_\alpha$ via an off-shell sfermion and the $Qd^c L_\alpha$ operator
$X \rightarrow \ell_\alpha^+ \ell_\alpha^-$ or $\ell_\alpha^+ \ell_\beta^-$	MSSM+RPV. LLP is sneutrino ($\tilde{\nu}$) that is produced via $pp \rightarrow \tilde{\chi}^+ \tilde{\chi}^-, \tilde{\chi}^+ \rightarrow \ell^+ \tilde{\nu}$. Then, $\tilde{\nu}_\alpha \rightarrow \ell_\alpha^+ \ell_\beta^-$ or $\tilde{\nu}_\beta \rightarrow \ell_\alpha^+ \ell_\alpha^-$ via the $L_\alpha L_\beta E_\alpha^c$ operator
$X \rightarrow \ell_\alpha^+ \ell_\alpha^- + \text{inv.}$	MSSM. LLP is second neutralino ($\tilde{\chi}_2^0$) that is produced via $pp \rightarrow \tilde{\chi}^+ \tilde{\chi}^-, \tilde{\chi}^+ \rightarrow W^+ \tilde{\chi}_2^0$. Then, $\tilde{\chi}_2^0 \rightarrow \ell_\alpha^+ \ell_\alpha^- \tilde{\chi}_1^0$ via an off-shell slepton and the $L_\alpha L_\beta E_\alpha^c$ operator
$X \rightarrow \ell_\alpha^+ \ell_\beta^- + \text{inv.}$	MSSM+RPV. LLP is bino ($\tilde{\chi}^0$) that is produced via $pp \rightarrow \tilde{\chi}^+ \tilde{\chi}^-, \tilde{\chi}^+ \rightarrow W^+ \tilde{\chi}^0$. Then, $\tilde{\chi}^0 \rightarrow \ell_\alpha^+ \ell_\alpha^- \nu_\beta$ or $\tilde{\chi}^0 \rightarrow \ell_\alpha^+ \ell_\beta^- \nu_\alpha$ via an off-shell slepton and the $L_\alpha L_\beta E_\alpha^c$ operator

Table A4. Simplified model library process proposals for Higgs (HIG) production mode where the Higgs decays to two LLPs. These modes are particularly important because they can come in association with forward jets (VBF) or leptons and \cancel{E}_T (VH). Note that, in cases of $M_X > M_h/2$, the same production modes could still occur if the Higgs is taken to be off-shell.

Decay mode	Simplified model library process
$X \rightarrow \gamma\gamma$	(N)MSSM. LLP is the lightest pseudoscalar Higgs (a) produced via $pp \rightarrow h, h \rightarrow aa$ Then, $a \rightarrow \gamma\gamma$.
$X \rightarrow \gamma\gamma + \text{inv.}$	MSSM. LLP is the second neutralino ($\tilde{\chi}_2^0$) produced via $pp \rightarrow h, h \rightarrow \tilde{\chi}_2^0 \tilde{\chi}_2^0$. Then, $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 \gamma\gamma$ via an off-shell SM Higgs
$X \rightarrow jj$	(N)MSSM. LLP is the lightest pseudoscalar Higgs (a) produced via $pp \rightarrow h, h \rightarrow aa$. Then, $a \rightarrow jj$.
$X \rightarrow jj + \text{inv.}$	MSSM. LLP is the second neutralino ($\tilde{\chi}_2^0$) that is produced via $pp \rightarrow h, h \rightarrow \tilde{\chi}_2^0 \tilde{\chi}_2^0$. Then, $\tilde{\chi}_2^0 \rightarrow jj \tilde{\chi}_1^0$ via an off-shell squark
$X \rightarrow \ell_\alpha^+ \ell_\alpha^-$	(N)MSSM. LLP is the lightest pseudoscalar Higgs (a) produced via $pp \rightarrow h, h \rightarrow aa$. Then, $a \rightarrow \ell_\alpha^+ \ell_\alpha^-$.
$X \rightarrow \ell_\alpha^+ \ell_\alpha^- + \text{inv.}$	MSSM. LLP is the second neutralino ($\tilde{\chi}_2^0$) produced via $pp \rightarrow h, h \rightarrow \tilde{\chi}_2^0 \tilde{\chi}_2^0$. Then, $\tilde{\chi}_2^0 \rightarrow \ell_\alpha^+ \ell_\alpha^- \tilde{\chi}_1^0$ via an off-shell slepton
$X \rightarrow \ell_\alpha^+ \ell_\beta^- + \text{inv.}$	MSSM+RPV. LLP is the second neutralino ($\tilde{\chi}_2^0$) produced via $pp \rightarrow h, h \rightarrow \tilde{\chi}_2^0 \tilde{\chi}_2^0$. Then, $\tilde{\chi}_2^0 \rightarrow \ell_\alpha^+ \ell_\beta^- \nu$ via an off-shell slepton and RPV couplings
$X \rightarrow \ell_\alpha^+ jj$	MSSM+RPV. LLP is the second neutralino ($\tilde{\chi}_2^0$) produced via $pp \rightarrow h, h \rightarrow \tilde{\chi}_2^0 \tilde{\chi}_2^0$. Then, $\tilde{\chi}_2^0 \rightarrow \ell_\alpha^+ jj$ via an off-shell slepton and RPV couplings

however, it does allow theorists and experimentalists to simulate these particular dark showers. The currently-included models are:

1. *Radiating lepton jets*: This is a model where DM is charged under a dark $U(1)'$ gauge interaction. The dark photon can be produced via final-state radiation from DM produced in a collider. Due to the collinear enhancement radiation of the dark photon, this can lead to a perturbative dark shower [589].
2. *Semi-visible jets*: This is a model in which hidden-sector states are charged under a new confining gauge group, leading to QCD-like showers. The showers produce both stable, invisible DM, as well as unstable states that decay back to the SM. This produces ‘semi-visible’ jets [372, 574].

Table A5. Simplified model library process proposals for Higgs (HIG) production mode where the Higgs decays to two LLPs plus invisible. These modes are particularly important because they can come in association with forward jets (VBF) or leptons and E_T (VH). Note that, in cases of $M_X > M_h/2$, the same production modes could still occur if the Higgs is taken to be off-shell.

Decay mode	Simplified model library process
$X \rightarrow \gamma\gamma$	(N)MSSM. LLP is a pseudoscalar or singlino (a) produced via $pp \rightarrow h, h \rightarrow \tilde{\chi}_2^0 \tilde{\chi}_2^0,$ $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 a$. Finally, $a \rightarrow \gamma\gamma$
$X \rightarrow \gamma\gamma + \text{inv.}$	MSSM. LLP is the second neutralino ($\tilde{\chi}_2^0$) produced via $pp \rightarrow h, h \rightarrow \tilde{\nu}\tilde{\nu}^*$, $\tilde{\nu} \rightarrow \tilde{\chi}_2^0 \nu$. Then, $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 \gamma\gamma$ via an off-shell SM Higgs
$X \rightarrow jj$	MSSM+RPV. LLP is a sneutrino ($\tilde{\nu}$) produced via $pp \rightarrow h$, $h \rightarrow \tilde{\chi}_1^0 \tilde{\chi}_1^0, \tilde{\chi}_1^0 \rightarrow \tilde{\nu}\tilde{\nu}$. Then, $\tilde{\nu} \rightarrow jj$ via the RPV operator LQd^c
$X \rightarrow jj + \text{inv.}$	MSSM. LLP is the second neutralino ($\tilde{\chi}_2^0$) that is produced via $pp \rightarrow h$, $h \rightarrow \tilde{\nu}\tilde{\nu}^*, \tilde{\nu} \rightarrow \nu\tilde{\chi}_2^0$. Then, $\tilde{\chi}_2^0 \rightarrow jj\tilde{\chi}_1^0$ via an off-shell squark
$X \rightarrow \ell_\alpha^+ \ell_\alpha^-$	MSSM+RPV. LLP is a sneutrino ($\tilde{\nu}_\beta$) produced via $pp \rightarrow h$, $h \rightarrow \tilde{\chi}_1^0 \tilde{\chi}_1^0, \tilde{\chi}_1^0 \rightarrow \tilde{\nu}_\beta \tilde{\nu}_\beta$. Then, $\tilde{\nu}_\beta \rightarrow \ell_\alpha^+ \ell_\alpha^-$ via the RPV operator $L_\alpha L_\beta E_\alpha^c$
$X \rightarrow \ell_\alpha^+ \ell_\alpha^- + \text{inv.}$	MSSM. LLP is the second neutralino ($\tilde{\chi}_2^0$) that is produced via $pp \rightarrow h$, $h \rightarrow \tilde{\nu}\tilde{\nu}^*, \tilde{\nu} \rightarrow \nu\tilde{\chi}_2^0$. Then, $\tilde{\chi}_2^0 \rightarrow \ell_\alpha^+ \ell_\alpha^- \tilde{\chi}_1^0$ via an off-shell slepton

Table A6. Simplified model library process proposals for Higgs (HIG) production mode where the Higgs decays to single LLP plus invisible. These modes are particularly important because they can come in association with forward jets (VBF) or leptons and E_T (VH). Note that, in cases of $M_X > M_h/2$, the same production modes could still occur if the Higgs is taken to be off-shell.

Decay mode	Simplified model library process
$X \rightarrow \gamma\gamma + \text{inv.}$	MSSM. LLP is the second neutralino ($\tilde{\chi}_2^0$) produced via $pp \rightarrow h$, $h \rightarrow \tilde{\chi}_2^0 \tilde{\chi}_1^0$. Then, $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 \gamma\gamma$ via an off-shell SM Higgs
$X \rightarrow jj + \text{inv.}$	MSSM. LLP is the second neutralino ($\tilde{\chi}_2^0$) that is produced via $pp \rightarrow h$, $h \rightarrow \tilde{\chi}_2^0 \tilde{\chi}_1^0$. Then, $\tilde{\chi}_2^0 \rightarrow jj\tilde{\chi}_1^0$ via an off-shell squark
$X \rightarrow \ell_\alpha^+ \ell_\alpha^- + \text{inv.}$	MSSM. LLP is the second neutralino ($\tilde{\chi}_2^0$) that is produced via $pp \rightarrow h$, $h \rightarrow \tilde{\chi}_2^0 \tilde{\chi}_1^0$. Then, $\tilde{\chi}_2^0 \rightarrow \ell_\alpha^+ \ell_\alpha^- \tilde{\chi}_1^0$ via an off-shell slepton
$X \rightarrow \ell_\alpha^+ jj$	MSSM+RPV. LLP is the second neutralino ($\tilde{\chi}_2^0$) that is produced via $pp \rightarrow h$, $h \rightarrow \tilde{\chi}_2^0 \tilde{\chi}_1^0$. Then, $\tilde{\chi}_2^0 \ell_\alpha^+ jj$ via the RPV operator QLd^c
$X \rightarrow \ell_\alpha^+ \ell_\beta^- + \text{inv.}$	MSSM+RPV. LLP is the second neutralino ($\tilde{\chi}_2^0$) that is produced via $pp \rightarrow h$, $h \rightarrow \tilde{\chi}_2^0 \tilde{\chi}_1^0$. Then, $\tilde{\chi}_2^0 \rightarrow \ell_\alpha^+ \ell_\beta^- \nu_\sigma$ via the RPV operator $L_\beta L_\sigma E_\alpha^c$

Table A7. Simplified model library process proposals for Z/Z' (ZP) production mode where the Z' decays to two LLPs. For this section, we mostly use a DM simplified model, where fermion x_2 is the LLP for $X \rightarrow SM +$ inv modes and scalar s_2 is the LLP for $X \rightarrow SM$ modes. The same models can also be used for off-shell Z/Z' where $M_X > M_{Z'}/2$. The model file includes all processes in the table; the undesired couplings can be set to zero and energy scales to be very large. The final entry uses a left-right symmetric model; the mass of the intermediate $Z/Z'/h$ can be changed to simulate the desired decay kinematics.

Decay mode	Simplified model library process
$X \rightarrow \gamma\gamma$	DMSM. LLP is scalar s_2 , produced via $pp \rightarrow Z' \rightarrow s_2 s_2$, then $s_2 \rightarrow \gamma\gamma$
$X \rightarrow \gamma\gamma +$ inv.	DMSM. LLP is fermion x_2 , produced via $pp \rightarrow Z' \rightarrow x_2 x_2$, then $x_2 \rightarrow \gamma\gamma x_1$.
$X \rightarrow jj$	DMSM. LLP is scalar s_2 , produced via $pp \rightarrow Z' \rightarrow s_2 s_2$, then $s_2 \rightarrow q\bar{q}$ $s_2 \rightarrow \ell_\alpha^+ \ell_\alpha^-$ couplings are proportional to SM Yukawa couplings.
$X \rightarrow jj +$ inv.	DMSM. LLP is fermion x_2 , produced via $pp \rightarrow Z' \rightarrow x_2 x_2$, then $x_2 \rightarrow q\bar{q} x_1$
$X \rightarrow \ell_\alpha^+ \ell_\alpha^-$	DMSM. LLP is scalar s_2 , produced via $pp \rightarrow Z' \rightarrow s_2 s_2$, then $s_2 \rightarrow \ell_\alpha^+ \ell_\alpha^-$ $s_2 \rightarrow \ell_\alpha^+ \ell_\alpha^-$ couplings are proportional to SM Yukawa couplings
$X \rightarrow \ell_\alpha^+ \ell_\alpha^- +$ inv.	DMSM. LLP is fermion x_2 , produced via $pp \rightarrow Z' \rightarrow x_2 x_2$, then $x_2 \rightarrow \ell_\alpha^+ \ell_\alpha^- x_1$
$X \rightarrow \ell_\alpha^+ jj$.	LRSM. LLP is fermion ν_R , produced via $pp \rightarrow Z/Z'/h \rightarrow \nu_R \nu_R$, then $\nu_R \rightarrow \ell_\alpha^+ jj$ via off-shell W/W'

Table A8. Simplified model library process proposals for Z/Z' (ZP) production mode where the Z/Z' decays to two LLPs plus invisible. For this section, we use a DM simplified model, where the Z' decays into $x_3 x_3$, and x_3 then decays into the LLP (fermion x_2 for SM+inv decay mode or scalar s_2 for SM decay mode), plus invisible (scalar s_1 or fermion x_1 , respectively). The same models can also be used for off-shell Z/Z' where $M_X > M_{Z'}/2$. One model file (DMSM) includes all processes in the table.

Decay mode	Simplified model library process
$X \rightarrow \gamma\gamma$	DMSM. LLP is scalar s_2 , produced via $pp \rightarrow Z' \rightarrow x_3 x_3$, $x_3 \rightarrow s_2 x_1$, then $s_2 \rightarrow \gamma\gamma$
$X \rightarrow \gamma\gamma +$ inv.	DMSM. LLP is fermion x_2 , produced via $pp \rightarrow Z' \rightarrow x_3 x_3$, $x_3 \rightarrow x_2 s_1$, then $x_2 \rightarrow \gamma\gamma x_1$
$X \rightarrow jj$	DMSM. LLP is scalar s_2 , produced via $pp \rightarrow Z' \rightarrow x_3 x_3$, $x_3 \rightarrow s_2 x_1$, then $s_2 \rightarrow q\bar{q}$ $s_2 \rightarrow \ell_\alpha^+ \ell_\alpha^-$ couplings are proportional to SM Yukawa couplings
$X \rightarrow jj +$ inv.	DMSM. LLP is fermion x_2 , produced via $pp \rightarrow Z' \rightarrow x_3 x_3$, $x_3 \rightarrow x_2 s_1$, then $x_2 \rightarrow q\bar{q} x_1$
$X \rightarrow \ell_\alpha^+ \ell_\alpha^-$	DMSM. LLP is scalar s_2 , produced via $pp \rightarrow Z' \rightarrow x_3 x_3$, $x_3 \rightarrow s_2 x_1$, then $s_2 \rightarrow \ell_\alpha^+ \ell_\alpha^-$ $s_2 \rightarrow \ell_\alpha^+ \ell_\alpha^-$ couplings are proportional to SM Yukawa couplings
$X \rightarrow \ell_\alpha^+ \ell_\alpha^- +$ inv.	DMSM. LLP is fermion x_2 , produced via $pp \rightarrow Z' \rightarrow x_3 x_3$, $x_3 \rightarrow x_2 s_1$, then $x_2 \rightarrow \ell_\alpha^+ \ell_\alpha^- x_1$

Table A9. Simplified model library process proposals for Z/Z' (ZP) production mode where the Z/Z' decays to single LLP plus invisible. For this section, we use a DM simplified model, where the Z' couples to an $x_1 x_2$ or $s_1 s_2$ pair. x_1 and s_1 behave as DM, the LLP x_2 decays into $x_1 + \text{SM}$, and the LLP s_2 decays into SM. The DMSM model file again includes all processes in the table except for the last, which is LRSM.

Decay mode	Simplified model library process
$X \rightarrow \gamma\gamma + \text{inv.}$	DMSM. LLP is fermion x_2 , produced via $pp \rightarrow Z' \rightarrow x_1 x_2$, then $x_2 \rightarrow \gamma\gamma x_1$; or a scalar s_2 , produced via $pp \rightarrow Z' \rightarrow s_1 s_2$, then $s_2 \rightarrow \gamma\gamma$
$X \rightarrow jj + \text{inv.}$	DMSM. LLP is fermion x_2 , produced via $pp \rightarrow Z' \rightarrow x_1 x_2$, then $x_2 \rightarrow jj x_1$; or a scalar s_2 , produced via $pp \rightarrow Z' \rightarrow s_1 s_2$, then $s_2 \rightarrow jj$
$X \rightarrow \ell_\alpha^+ \ell_\alpha^- + \text{inv.}$	DMSM. LLP is fermion x_2 , produced via $pp \rightarrow Z' \rightarrow x_1 x_2$, then $x_2 \rightarrow \ell^+ \ell^- x_1$; or a scalar s_2 , produced via $pp \rightarrow Z' \rightarrow s_1 s_2$, then $s_2 \rightarrow \ell^+ \ell^-$
$X \rightarrow \ell_\alpha^+ jj$	LRSM. LLP is fermion ν_R , produced via $pp \rightarrow Z/Z'/h \rightarrow \bar{\nu}_L \nu_R$, then $\nu_R \rightarrow \ell^+ jj$

Table A10. Simplified model library process proposals for charged current (CC) production mode, $W_{\text{SM}}^\pm/W'^\pm \rightarrow X + \ell^\pm$; these can be simulated using left-right symmetric models using either the W or W' (for simplicity, in the table above we only state explicitly W). Right-handed neutrino lifetimes are most naturally long for sub-weak-scale masses.

Decay mode	Simplified model library process
$X \rightarrow jj + \text{inv.}$	LRSM. LLP is the right-handed neutrino (ν_R) produced via $pp \rightarrow W^\pm, W^\pm \rightarrow \ell^\pm \nu_R$. Then, $\nu_R \rightarrow q\bar{q}\nu$ via an off-shell Z . For massive invisible state, it may be possible to use a cascade $\nu_{R2} \rightarrow q\bar{q}\nu_{R1}$ treating the lightest right-handed neutrino as stable
$X \rightarrow jj\ell^\pm$	LRSM. LLP is the right-handed neutrino (ν_R) produced via $pp \rightarrow W^\pm, W^\pm \rightarrow \ell^\pm \nu_R$. Then, $\nu_R \rightarrow q\bar{q}\ell^\pm$ via an off-shell W . Alternately, production and decay can be mediated by W_R .
$X \rightarrow \ell_\alpha^+ \ell_\alpha^- + \text{inv.}$ or $X \rightarrow \ell_\alpha^+ \ell_\beta^- + \text{inv.}$	LRSM. LLP is the right-handed neutrino (ν_R) produced via $pp \rightarrow W^\pm, W^\pm \rightarrow \ell^\pm \nu_R$. Then, $\nu_R \rightarrow \ell_\alpha^+ \ell_\alpha^- \nu_\beta$ or $\nu_R \rightarrow \ell_\alpha^+ \ell_\beta^- \nu_\alpha$ via an off-shell W/Z

ORCID iDs

Juliette Alimena
 <https://orcid.org/0000-0001-6030-3191>

James Beacham
 <https://orcid.org/0000-0003-3623-3335>

Martino Borsato
 <https://orcid.org/0000-0001-5760-2924>

Xabier Cid Vidal
 <https://orcid.org/0000-0002-0468-541X>

Giovanna Cottin
 ⓘ <https://orcid.org/0000-0002-5308-5808>
 David Curtin
 ⓘ <https://orcid.org/0000-0003-0263-6195>
 Albert De Roeck
 ⓘ <https://orcid.org/0000-0002-9228-5271>
 Nishita Desai
 ⓘ <https://orcid.org/0000-0001-7942-1649>
 Jared A Evans
 ⓘ <https://orcid.org/0000-0001-9871-6557>
 Simon Knapen
 ⓘ <https://orcid.org/0000-0002-6733-9231>
 Sabine Kraml
 ⓘ <https://orcid.org/0000-0002-2613-7000>
 Zhen Liu
 ⓘ <https://orcid.org/0000-0002-3143-1976>
 Sascha Mehlhase
 ⓘ <https://orcid.org/0000-0002-1281-2060>
 Michael J Ramsey-Musolf
 ⓘ <https://orcid.org/0000-0001-8110-2479>
 Heather Russell
 ⓘ <https://orcid.org/0000-0003-4181-0678>
 Jessie Shelton
 ⓘ <https://orcid.org/0000-0002-0959-6360>
 Brian Shuve
 ⓘ <https://orcid.org/0000-0002-3524-2021>
 Monica Verducci
 ⓘ <https://orcid.org/0000-0001-8209-4757>
 Jose Zurita
 ⓘ <https://orcid.org/0000-0001-5710-3965>

References

- [1] Shuve B 2017 Theory overview of long-lived particles at the LHC https://indico.cern.ch/event/607314/contributions/2542308/attachments/1447888/2231430/LHC-LLP_Shuve.pdf
- [2] Curtin D *et al* 2019 Long-lived particles at the energy frontier: the MATHUSLA physics case *Rep. Prog. Phys.* **82** 116201
- [3] Russell H 2017 An experimental introduction to long-lived particle searches at the LHC https://indico.cern.ch/event/607314/contributions/2542309/attachments/1447873/2231444/20170424_LLPs.pdf
- [4] CMS Collaboration 2016 Search for heavy stable charged particles with 12.9 fb^{-1} of 2016 data CMS-PAS-EXO-16-036 <https://cds.cern.ch/record/2205281>
- [5] Alexander J *et al* 2016 Dark sectors 2016 workshop: community report arXiv:1608.08632 [hep-ph]
- [6] Battaglieri M *et al* 2017 *US Cosmic Visions: New Ideas in Dark Matter 2017: Community Report. U.S. Cosmic Visions: New Ideas in Dark Matter (23–25,March 2017) (College Park, MD, USA*arXiv:1707.04591 [hep-ph]
- [7] Beacham J *et al* 2020 Physics beyond colliders at CERN: beyond the Standard Model working group report *J. Phys. G: Nucl. Part. Phys.* **47** 010501
- [8] Fayet P 1976 Supersymmetry and weak, electromagnetic and strong interactions *Phys. Lett. B* **64** 159

[9] Fayet P 1977 Spontaneously broken supersymmetric theories of weak, electromagnetic and strong interactions *Phys. Lett. B* **69** 489

[10] Farrar G R and Fayet P 1978 Phenomenology of the production, decay, and detection of new hadronic states associated with supersymmetry *Phys. Lett. B* **76** 575–9

[11] Fayet P 1979 Relations between the masses of the superpartners of leptons and quarks, the goldstino couplings and the neutral currents *Phys. Lett. B* **84** 416

[12] Dimopoulos S and Georgi H 1981 Softly broken supersymmetry and SU(5) *Nucl. Phys. B* **193** 150–62

[13] Chacko Z, Goh H-S and Harnik R 2006 The twin Higgs: natural electroweak breaking from mirror symmetry *Phys. Rev. Lett.* **96** 231802

[14] Burdman G *et al* 2007 Folded supersymmetry and the LEP paradox *J. High Energy Phys.* **JHEP02(2007)009**

[15] Cai H, Cheng H-C and Terning J 2009 A quirky little Higgs model *J. High Energy Phys.* **JHEP05(2009)045**

[16] Boehm C, Ensslin T A and Silk J 2004 Can annihilating dark matter be lighter than a few GeVs? *J. Phys. G: Nucl. Part. Phys.* **30** 279–86

[17] Boehm C, Fayet P and Silk J 2004 Light and heavy dark matter particles *Phys. Rev. D* **69** 101302

[18] Pospelov M, Ritz A and Voloshin M B 2008 Secluded WIMP dark matter *Phys. Lett. B* **662** 53–61

[19] Pospelov M 2009 Secluded U(1) below the weak scale *Phys. Rev. D* **80** 095002

[20] Arkani-Hamed N *et al* 2009 A theory of dark matter *Phys. Rev. D* **79** 015014

[21] Pospelov M and Ritz A 2009 Astrophysical signatures of secluded dark matter *Phys. Lett. B* **671** 391–7

[22] Bouquet A and Salati P R 1987 Parity breaking and cosmological consequences *Nucl. Phys. B* **284** 557

[23] Campbell B A *et al* 1991 Cosmological baryon asymmetry constraints on extensions of the Standard Model *Phys. Lett. B* **256** 484–90

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

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

[26] Cui Y and Shuve B 2015 Probing baryogenesis with displaced vertices at the LHC *J. High Energy Phys.* **JHEP02(2015)049**

[27] Cui Y 2015 A review of WIMP baryogenesis mechanisms *Mod. Phys. Lett. A* **30** 1530028

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

[29] Feng J L and Kumar J 2008 The WIMPless miracle: dark-matter particles without weak-scale masses or weak interactions *Phys. Rev. Lett.* **101** 231301

[30] Baumgart M *et al* 2009 Non-Abelian dark sectors and their collider signatures *J. High Energy Phys.* **JHEP04(2009)014**

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

[32] Chan Y F *et al* 2012 LHC signatures of a minimal supersymmetric hidden valley *J. High Energy Phys.* **JHEP05(2012)155**

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

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

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

[36] Fidalgo J *et al* 2009 Neutrino physics and spontaneous CP violation in the $\mu\nu$ SSM *J. High Energy Phys.* **JHEP08(2009)105**

[37] Ghosh P *et al* 2013 Probing the μ -from- ν supersymmetric Standard Model with displaced multileptons from the decay of a Higgs boson at the LHC *Phys. Rev. D* **88** 015009

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

[39] Helo J C, Hirsch M and Kovalenko S 2016 *Phys. Rev. D* **93** 099902 (erratum)

[39] Antusch S, Cazzato E and Fischer O 2016 Displaced vertex searches for sterile neutrinos at future lepton colliders *J. High Energy Phys.* **JHEP12(2016)007**

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

[41] Graesser M L 2007 Experimental constraints on Higgs Boson decays to TeV-scale right-handed neutrinos arXiv:0705.2190 [hep-ph]

[42] Ghosh P *et al* 2015 Hunting physics beyond the Standard Model with unusual W^\pm and Z decays *Phys. Rev. D* **91** 035020

[43] Izaguirre E and Shuve B 2015 Multilepton and lepton jet probes of sub-weak-scale right-handed neutrinos *Phys. Rev. D* **91** 093010

[44] Maiezza A, Nemevšek M and Nesti F 2015 Lepton number violation in Higgs decay at LHC *Phys. Rev. Lett.* **115** 081802

[45] Batell B, Pospelov M and Shuve B 2016 Shedding light on neutrino masses with dark forces *J. High Energy Phys.* JHEP08(2016)052

[46] Das A, Dev P S B and Kim C S 2017 Constraining sterile neutrinos from precision Higgs data *Phys. Rev. D* **95** 115013

[47] Lara I *et al* 2018 Looking for the left sneutrino LSP with displaced-vertex searches *Phys. Rev. D* **98** 075004

[48] Cottin G, Helo J C and Hirsch M 2018 Searches for light sterile neutrinos with multitrack displaced vertices *Phys. Rev. D* **97** 055025

[49] Nemevšek M, Nesti F and Popara G 2018 Keung–Senjanović process at LHC: from LNV to displaced vertices to invisible decays *Phys. Rev. D* **97** 115018

[50] Curtin D, Dienes K R and Thomas B 2018 Dynamical dark matter, MATHUSLA, and the lifetime frontier *Phys. Rev. D* **98** 115005

[51] Das A 2018 Searching for the minimal Seesaw models at the LHC and beyond *Adv. High Energy Phys.* **2018** 9785318

[52] Cvetic G, Das A and Zamora-Saá J 2019 Probing heavy neutrino oscillations in rare W boson decays *J. Phys. G: Nucl. Part. Phys.* **46** 075002

[53] Giudice G F and Rattazzi R 1999 Theories with gauge mediated supersymmetry breaking *Phys. Rep.* **322** 419–99

[54] Fan J, Reece M and Ruderman J T 2011 Stealth supersymmetry *J. High Energy Phys.* JHEP11(2011)012

[55] Barbier R *et al* 2005 R-parity violating supersymmetry *Phys. Rep.* **420** 1–202

[56] Csaki C, Kuflik E and Volansky T 2014 Dynamical R-parity violation *Phys. Rev. Lett.* **112** 131801

[57] Arvanitaki A *et al* 2013 Mini-split *J. High Energy Phys.* JHEP02(2013)126

[58] Arkani-Hamed N *et al* 2012 Simply unnatural supersymmetry arXiv:1212.6971 [hep-ph].

[59] de Vries J, Dreiner H K and Schmeier D 2016 R-parity violation and light neutralinos at SHiP and the LHC *Phys. Rev. D* **94** 035006

[60] Dercks D *et al* 2019 R-parity violation and light neutralinos at CODEX-b, FASER, and MATHUSLA *Phys. Rev. D* **99** 055039

[61] Dercks D *et al* 2019 Long-lived fermions at AL3X *Phys. Rev. D* **99** 055020

[62] Bond A D *et al* 2017 Directions for model building from asymptotic safety *J. High Energy Phys.* JHEP08(2017)004

[63] Chen C H, Drees M and Gunion J F 1996 Searching for invisible and almost invisible particles at $e + e -$ colliders *Phys. Rev. Lett.* **76** 2002–5

[64] Thomas S D and Wells J D 1998 Phenomenology of massive vectorlike doublet leptons *Phys. Rev. Lett.* **81** 34–7

[65] Feng J L *et al* 1999 Discovering supersymmetry at the Tevatron in wino LSP scenarios *Phys. Rev. Lett.* **83** 1731–4

[66] Strassler M J and Zurek K M 2007 Echoes of a hidden valley at hadron colliders *Phys. Lett. B* **651** 374–9

[67] Strassler M J and Zurek K M 2008 Discovering the Higgs through highly-displaced vertices *Phys. Lett. B* **661** 263–7

[68] Strassler M J 2006 Possible effects of a hidden valley on supersymmetric phenomenology arXiv: hep-ph/0607160 [hep-ph]

[69] Han T *et al* 2008 Phenomenology of hidden valleys at hadron colliders *J. High Energy Phys.* JHEP07(2008)008

[70] Strassler M J 2008 Why unparticle models with mass gaps are examples of hidden valleys arXiv:0801.0629 [hep-ph]

[71] Strassler M J 2008 On the phenomenology of hidden valleys with heavy flavor arXiv:[0806.2385](https://arxiv.org/abs/0806.2385) [hep-ph]

[72] Arkani-Hamed N *et al* 2006 Supersymmetry and the LHC inverse problem *J. High Energy Phys.* **JHEP08(2006)070**

[73] Knuteson B and Mrenna S 2006 BARD: interpreting new frontier energy collider physics arXiv: [hep-ph/0602101](https://arxiv.org/abs/hep-ph/0602101) [hep-ph]

[74] Arkani-Hamed N *et al* 2007 MAMOSET: the path from LHC data to the new Standard Model via on-shell effective theories arXiv:[hep-ph/0703088](https://arxiv.org/abs/hep-ph/0703088) [hep-ph]

[75] Altonen T *et al* 2008 Model-independent and quasi-model-independent search for new physics at CDF *Phys. Rev. D* **78** 012002

[76] Alwall J, Schuster P and Toro N 2009 Simplified models for a first characterization of new physics at the LHC *Phys. Rev. D* **79** 075020

[77] Alwall J *et al* 2009 Model-independent jets plus missing energy searches *Phys. Rev. D* **79** 015005

[78] Alves D 2012 Simplified models for LHC new physics searches *J. Phys. G: Nucl. Part. Phys.* **39** 105005

[79] Petriello F J, Quackenbush S and Zurek K M 2008 The invisible Z' at the CERN LHC *Phys. Rev. D* **77** 115020

[80] Dudas E *et al* 2009 (In)visible Z' -prime and dark matter *J. High Energy Phys.* **JHEP08(2009)014**

[81] Goodman J and Shepherd W 2011 LHC bounds on UV-complete models of dark matter arXiv:[1111.2359](https://arxiv.org/abs/1111.2359) [hep-ph]

[82] An H, Ji X and Wang L-T 2012 Light dark matter and Z' dark force at colliders *J. High Energy Phys.* **JHEP07(2012)182**

[83] Frandsen M T *et al* 2012 LHC and tevatron bounds on the dark matter direct detection cross-section for vector mediators *J. High Energy Phys.* **JHEP07(2012)123**

[84] Dreiner H, Schmeier D and Tattersall J 2013 Contact interactions probe effective dark matter models at the LHC *EPL* **102** 51001

[85] Cotta R C *et al* 2014 Particle physics implications and constraints on dark matter interpretations of the CDMS signal *Phys. Rev. D* **90** 013020

[86] Abdallah J *et al* 2015 Simplified models for dark matter searches at the LHC *Phys. Dark Univ.* **9-10** 8-23

[87] Abercrombie D *et al* 2019 Dark matter benchmark models for early LHC Run-2 searches: report of the ATLAS/CMS dark matter forum *Phys. Dark Univ.* **26** 100371

[88] Busoni G *et al* 2016 Recommendations on presenting LHC searches for missing transverse energy signals using simplified s-channel models of dark matter 2020 *Phys. Dark Univ.* **27** 100365

[89] Arina C *et al* 2016 A comprehensive approach to dark matter studies: exploration of simplified top-philic models *J. High Energy Phys.* **JHEP11(2016)111**

[90] Abe T *et al* (LHC Dark Matter Working Group) 2020 LHC dark matter working group: next-generation spin-0 dark matter models *Phys. Dark Univ.* **27** 100351

[91] Edelhäuser L, Krämer M and Sonneveld J 2015 Simplified models for same-spin new physics scenarios *J. High Energy Phys.* **JHEP04(2015)146**

[92] Edelhäuser L *et al* 2014 Constraining supersymmetry at the LHC with simplified models for squark production *J. High Energy Phys.* **JHEP12(2014)022**

[93] Arina C *et al* 2015 Constraints on sneutrino dark matter from LHC Run 1 *J. High Energy Phys.* **JHEP05(2015)142**

[94] Kraml S *et al* 2016 Scalar versus fermionic top partner interpretations of $t\bar{t} + E_T^{\text{miss}}$ searches at the LHC *J. High Energy Phys.* **JHEP11(2016)107**

[95] Heisig J and Kersten J 2012 Long-lived staus from strong production in a simplified model approach *Phys. Rev. D* **86** 055020

[96] Liu Z and Tweedie B 2015 The fate of long-lived superparticles with hadronic decays after LHC Run 1 *J. High Energy Phys.* **JHEP06(2015)042**

[97] Heisig J, Lessa A and Quertenmont L 2015 Simplified models for exotic BSM searches *J. High Energy Phys.* **JHEP12(2015)087**

[98] Khoze V V, Plascencia A D and Sakurai K 2017 Simplified models of dark matter with a long-lived co-annihilation partner *J. High Energy Phys.* **JHEP06(2017)041**

[99] Mahbubani R, Schwaller P and Zurita J 2017 Closing the window for compressed Dark Sectors with disappearing charged tracks *J. High Energy Phys.* **JHEP06(2017)119**

[100] Buchmueller O *et al* 2017 Simplified models for displaced dark matter signatures *J. High Energy Phys.* **JHEP09(2017)076**

[101] Pinfold J *et al* 2009 Technical design report of the MoEDAL experiment *Technical Report CERN-LHCC-2009-006. MoEDAL-TDR-001* CERN <https://cds.cern.ch/record/1181486>

[102] Haas A *et al* 2015 Looking for milli-charged particles with a new experiment at the LHC *Phys. Lett. B* **746** 117–20

[103] Chou J P, Curtin D and Lubatti H J 2017 New detectors to explore the lifetime frontier *Phys. Lett. B* **767** 29–36

[104] Gligorov V V *et al* 2018 Searching for long-lived particles: a compact detector for exotics at LHCb *Phys. Rev. D* **97** 015023

[105] Feng J L *et al* 2018 ForwArd search ExpeRiment at the LHC *Phys. Rev. D* **97** 035001

[106] Ambrogi F *et al* 2018 On the coverage of the pMSSM by simplified model results *Eur. Phys. J. C* **78** 215

[107] Essig R *et al* 2012 Heavy flavor simplified models at the LHC *J. High Energy Phys.* **JHEP01(2012)074**

[108] Brust C *et al* 2012 SUSY, the third generation and the LHC *J. High Energy Phys.* **JHEP03(2012)103**

[109] Papucci M, Ruderman J T and Weiler A 2012 Natural SUSY endures *J. High Energy Phys.* **JHEP09(2012)035**

[110] Lopez-Fogliani D E and Munoz C 2006 Proposal for a supersymmetric Standard Model *Phys. Rev. Lett.* **97** 041801

[111] Ghosh P *et al* 2018 Searching for left sneutrino LSP at the LHC *Int. J. Mod. Phys. A* **33** 1850110

[112] Dimopoulos S *et al* 1996 Experimental signatures of low-energy gauge mediated supersymmetry breaking *Phys. Rev. Lett.* **76** 3494–7

[113] Arkani-Hamed N and Dimopoulos S 2005 Supersymmetric unification without low energy supersymmetry and signatures for fine-tuning at the LHC *J. High Energy Phys.* **JHEP06(2005)073**

[114] Byrne M 2004 Universal extra dimensions and charged LKPs *Phys. Lett. B* **583** 309–14

[115] Khachatryan V *et al* (CMS Collaboration) 2015 Precise determination of the mass of the Higgs boson and tests of compatibility of its couplings with the Standard Model predictions using proton collisions at 7 and 8 TeV *Eur. Phys. J. C* **75** 212

[116] Aad G *et al* (ATLAS Collaboration) 2015 Constraints on new phenomena via Higgs boson couplings and invisible decays with the ATLAS detector *J. High Energy Phys.* **JHEP11(2015)206**

[117] Curtin D *et al* 2014 Exotic decays of the 125 GeV Higgs boson *Phys. Rev. D* **90** 075004

[118] Craig N *et al* 2015 Naturalness in the dark at the LHC *J. High Energy Phys.* **JHEP07(2015)105**

[119] Silveira V and Zee A 1985 Scalar phantoms *Phys. Lett. B* **161** 136–40

[120] Beauchesne H, Bertuzzo E and Grilli di Cortona G 2017 Constraints on the relaxion mechanism with strongly interacting vector-fermions *J. High Energy Phys.* **JHEP08(2017)093**

[121] Langacker P 2009 The physics of heavy Z' Gauge bosons *Rev. Mod. Phys.* **81** 1199–228

[122] Holdom B 1986 Two U(1)'s and epsilon charge shifts *Phys. Lett. B* **166** 196–8

[123] Feng J L, Kaplinghat M and Yu H-B 2010 Halo shape and relic density exclusions of sommerfeld-enhanced dark matter explanations of cosmic ray excesses *Phys. Rev. Lett.* **104** 151301

[124] Buckley M R and Fox P J 2010 Dark matter self-interactions and light force carriers *Phys. Rev. D* **81** 083522

[125] Tulin S, Yu H-B and Zurek K M 2013 Resonant dark forces and small scale structure *Phys. Rev. Lett.* **110** 111301

[126] Boehm C and Fayet P 2004 Scalar dark matter candidates *Nucl. Phys. B* **683** 219–63

[127] Arkani-Hamed N and Weiner N 2008 LHC signals for a SuperUnified theory of dark matter *J. High Energy Phys.* **JHEP12(2008)104**

[128] Cirelli M, Fornengo N and Strumia A 2006 Minimal dark matter *Nucl. Phys. B* **753** 178–94

[129] Cirelli M and Strumia A 2009 Minimal dark matter: model and results *New J. Phys.* **11** 105005

[130] Fileviez Perez P *et al* 2009 Triplet scalars and dark matter at the LHC *Phys. Rev. D* **79** 055024

[131] Hochberg Y, Kuflik E and Murayama H 2016 SIMP spectroscopy *J. High Energy Phys.* **JHEP05(2016)090**

[132] Tucker-Smith D and Weiner N 2001 Inelastic dark matter *Phys. Rev. D* **64** 043502

[133] Bai Y and Tait T M P 2012 Inelastic dark matter at the LHC *Phys. Lett. B* **710** 335–8

[134] Weiner N and Yavin I 2012 How dark are Majorana WIMPs? Signals from MiDM and Rayleigh dark matter *Phys. Rev. D* **86** 075021

[135] Izaguirre E, Krnjaic G and Shuve B 2016 Discovering inelastic thermal-relic dark matter at colliders *Phys. Rev. D* **93** 063523

[136] Griest K and Seckel D 1991 Three exceptions in the calculation of relic abundances *Phys. Rev. D* **43** 3191–203

[137] Baker M J *et al* 2015 The coannihilation codex *J. High Energy Phys.* **JHEP12(2015)120**

[138] De Simone A, Sanz V and Sato H P 2010 Pseudo-Dirac dark matter leaves a trace *Phys. Rev. Lett.* **105** 121802

[139] Davoli A *et al* 2017 Displaced vertices from pseudo-dirac dark matter *J. High Energy Phys.* **JHEP11(2017)025**

[140] El Hedri S *et al* 2017 Simplified phenomenology for colored dark sectors *J. High Energy Phys.* **JHEP04(2017)118**

[141] Davoli A *et al* 2018 LHC phenomenology of dark matter with a color-octet partner *J. High Energy Phys.* **JHEP07(2018)054**

[142] Garny M *et al* 2017 Coannihilation without chemical equilibrium *Phys. Rev. D* **96** 103521

[143] D’Agnolo R T, Pappadopulo D and Ruderman J T 2017 Fourth exception in the calculation of relic abundances *Phys. Rev. Lett.* **119** 061102

[144] Garny M *et al* 2018 Top-philic dark matter within and beyond the WIMP paradigm *Phys. Rev. D* **97** 075002

[145] Cheng H-C, Li L and Zheng R 2018 Coscattering/coannihilation dark matter in a fraternal twin Higgs model *J. High Energy Phys.* **JHEP09(2018)098**

[146] Hall L J *et al* 2010 Freeze-in production of FIMP dark matter *J. High Energy Phys.* **JHEP03(2010)080**

[147] Co R T *et al* 2015 Freeze-in dark matter with displaced signatures at colliders *J. Cosmol. Astropart. Phys.* **JCAP12(2015)024**

[148] Hessler A G *et al* 2017 Probing the scotogenic FIMP at the LHC *J. High Energy Phys.* **JHEP01(2017)100**

[149] Ghosh A, Mondal T and Mukhopadhyaya B 2017 Heavy stable charged tracks as signatures of non-thermal dark matter at the LHC : a study in some non-supersymmetric scenarios *J. High Energy Phys.* **JHEP12(2017)136**

[150] D’Eramo F, Fernandez N and Profumo S 2018 Dark matter freeze-in production in fast-expanding universes *J. Cosmol. Astropart. Phys.* **JCAP02(2018)046**

[151] Calibbi L *et al* 2018 Singlet-doublet dark matter freeze-in: LHC displaced signatures versus cosmology *J. High Energy Phys.* **JHEP09(2018)037**

[152] Garny M and Heisig J 2018 SuperWIMP meets freeze-in *Phys. Rev. D* **98** 095031

[153] Bélanger G *et al* 2019 LHC-friendly minimal freeze-in models *J. High Energ. Phys.* **JHEP02(2019)186**

[154] Heisig J 2018 Dark matter and long-lived particles at the LHC *53rd Rencontres de Moriond on Electroweak Interactions and Unified Theories (Moriond EW 2018) La Thuile 10–17, March 2018 (Italy, 10–17 March)* arXiv:1805.07361 [hep-ph]

[155] Falkowski A *et al* 2010 Hidden Higgs decaying to lepton jets *J. High Energy Phys.* **JHEP05(2010)077**

[156] Primulando R, Salvioni E and Tsai Y 2015 The dark penguin shines light at colliders *J. High Energy Phys.* **JHEP07(2015)031**

[157] Bai Y, Bourbeau J and Lin T 2015 Dark matter searches with a mono- Z' jet *J. High Energy Phys.* **JHEP06(2015)205**

[158] Minkowski P 1977 $\mu \rightarrow e\gamma$ at a rate of one out of 10^9 Muon decays? *Phys. Lett.* **67B** 421–8

[159] Yanagida T 1979 *Horizontal Symmetry and Masses of Neutrinos Conf. Proc. C7902131* pp 95–9

[160] Mohapatra R N and Senjanovic G 1980 Neutrino mass and spontaneous parity violation *Phys. Rev. Lett.* **44** 912

[161] Glashow S L 1980 The future of elementary particle physics *NATO Sci. Ser. B* **61** 687

[162] Mohapatra R N and Valle J W F 1986 Neutrino mass and baryon number nonconservation in superstring models *Phys. Rev. D* **34** 1642

[163] Keung W-Y and Senjanovic G 1983 Majorana neutrinos and the production of the right-handed charged Gauge boson *Phys. Rev. Lett.* **50** 1427

[164] Ferrari A *et al* 2000 Sensitivity study for new gauge bosons and right-handed Majorana neutrinos in pp collisions at $s = 14$ TeV *Phys. Rev. D* **62** 013001

[165] Basso L *et al* 2009 Phenomenology of the minimal B-L extension of the Standard model: Z' and neutrinos *Phys. Rev. D* **80** 055030

[166] Atre A *et al* 2009 The search for heavy Majorana neutrinos *J. High Energy Phys.* **JHEP05** (2009)030

[167] Fileviez Perez P, Han T and Li T 2009 Testability of type I seesaw at the CERN LHC: revealing the existence of the B-L symmetry *Phys. Rev. D* **80** 073015

[168] Nemevsek M *et al* 2011 First limits on left-right symmetry scale from LHC data *Phys. Rev. D* **83** 115014

[169] Nemevsek M, Nesti F and Vasquez J C 2017 Majorana Higgses at colliders *J. High Energy Phys.* **JHEP04**(2017)114

[170] Accomando E *et al* 2017 Novel SM-like Higgs decay into displaced heavy neutrino pairs in $U(1)'$ models *J. High Energy Phys.* **JHEP04**(2017)081

[171] Accomando E *et al* 2016 Higgses and heavy neutrinos in $U(1)'$ models: from the LHC to the GUT scale *J. High Energy Phys.* **JHEP07**(2016)086

[172] Caputo A *et al* 2017 The seesaw portal in testable models of neutrino masses *J. High Energy Phys.* **JHEP06**(2017)112

[173] Accomando E *et al* 2018 Extra Higgs boson and Z' as portals to signatures of heavy neutrinos at the LHC *J. High Energy Phys.* **JHEP02**(2018)109

[174] Helo J C, Hirsch M and Wang Z S 2018 Heavy neutral fermions at the high-luminosity LHC *J. High Energy Phys.* **JHEP07**(2018)056

[175] Asaka T, Blanchet S and Shaposhnikov M 2005 The nuMSM, dark matter and neutrino masses *Phys. Lett. B* **631** 151–6

[176] Asaka T and Shaposhnikov M 2005 The nuMSM, dark matter and baryon asymmetry of the Universe *Phys. Lett. B* **620** 17–26

[177] Pati J C and Salam A 1974 Lepton number as the fourth color *Phys. Rev. D* **10** 275–89
Pati J C and Salam A 1975 *Phys. Rev. D* **11** 703 (erratum)

[178] Mohapatra R N and Pati J C 1975 A natural left-right symmetry *Phys. Rev. D* **11** 2558

[179] Senjanovic G and Mohapatra R N 1975 Exact left-right symmetry and spontaneous violation of parity *Phys. Rev. D* **12** 1502

[180] Senjanovic G 1979 Spontaneous breakdown of parity in a class of Gauge theories *Nucl. Phys. B* **153** 334–64

[181] Das A, Gao Y and Kamon T 2019 Heavy neutrino search via semileptonic Higgs decay at the LHC *Eur. Phys. J. C* **79** 424

[182] Kang J and Luty M A 2009 Macroscopic strings and ‘Quarks’ at colliders *J. High Energy Phys.* **JHEP11**(2009)065

[183] Chacko Z, Curtin D and Verhaaren C B 2016 A quirky probe of neutral naturalness *Phys. Rev. D* **94** 011504

[184] Autran M *et al* 2015 Searches for dark matter in events with a resonance and missing transverse energy *Phys. Rev. D* **92** 035007

[185] Blinov N, Izaguirre E and Shuve B 2018 Rare Z Boson decays to a hidden sector *Phys. Rev. D* **97** 015009

[186] Craig N *et al* 2016 The Higgs portal above threshold *J. High Energy Phys.* **JHEP02**(2016)127

[187] Khachatryan V *et al* (CMS Collaboration) 2017 Search for narrow resonances in dilepton mass spectra in proton–proton collisions at $\sqrt{s} = 13$ TeV and combination with 8 TeV data *Phys. Lett. B* **768** 57–80

[188] Sirunyan A M *et al* (CMS Collaboration) 2017 Search for dijet resonances in proton–proton collisions at $\sqrt{s} = 13$ TeV and constraints on dark matter and other models *Phys. Lett. B* **769** 520–42
Sirunyan A M *et al* (CMS Collaboration) 2017 *Phys. Lett. B* **772** 882 (erratum)

[189] Aaboud M *et al* (ATLAS Collaboration) 2017 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. D* **96** 052004

[190] Sirunyan A M *et al* (CMS Collaboration) 2017 Search for low mass vector resonances decaying to quark-antiquark pairs in proton–proton collisions at $\sqrt{s} = 13$ TeV *Phys. Rev. Lett.* **119** 111802

[191] Aaboud M *et al* (ATLAS Collaboration) 2017 Search for new high-mass phenomena in the dilepton final state using 36 fb^{-1} of proton–proton collision data at $\sqrt{s} = 13$ TeV with the ATLAS detector *J. High Energy Phys.* **JHEP10**(2017)182

[192] Aaboud M *et al* (ATLAS Collaboration) 2018 Search for $W' \rightarrow tb$ decays in the hadronic final state using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector *Phys. Lett. B* **781** 327–48

[193] Aaboud M *et al* (ATLAS Collaboration) 2019 Search for light resonances decaying to boosted quark pairs and produced in association with a photon or a jet in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector *Phys. Lett. B* **788** 316–35

[194] Ibe M, Moroi T and Yanagida T T 2007 Possible signals of wino LSP at the Large Hadron Collider *Phys. Lett. B* **644** 355–60

[195] Buckley M R, Randall L and Shuve B 2011 LHC searches for non-chiral weakly charged multiplets *J. High Energy Phys.* **JHEP05(2011)097**

[196] Aad G *et al* (ATLAS Collaboration) 2015 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. D* **92** 072004

[197] CMS Collaboration 2016 Search for displaced leptons in the e-mu channel *CMS-PAS-EXO-16-022* <https://cds.cern.ch/record/2205146>

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

[199] Dimopoulos S, Thomas S D and Wells J D 1997 Sparticle spectroscopy and electroweak symmetry breaking with gauge mediated supersymmetry breaking *Nucl. Phys. B* **488** 39–91

[200] del Aguila F and Aguilar-Saavedra J A 2009 Distinguishing seesaw models at LHC with multi-lepton signals *Nucl. Phys. B* **813** 22–90

[201] Ambrosanio S, Kribs G D and Martin S P 1997 Signals for gauge mediated supersymmetry breaking models at the CERN LEP-2 collider *Phys. Rev. D* **56** 1761–77

[202] Delgado A, Giudice G F and Slavich P 2007 Dynamical mu term in gauge mediation *Phys. Lett. B* **653** 424–33

[203] Meade P, Reece M and Shih D 2010 Long-lived neutralino NLSPs *J. High Energy Phys.* **JHEP10(2010)067**

[204] Allanach B *et al* 2015 Light sparticles from a light singlet in Gauge mediation *Phys. Rev. D* **92** 015006

[205] Evans J A and Shelton J 2016 Long-lived staus and displaced leptons at the LHC *J. High Energy Phys.* **JHEP04(2016)056**

[206] Allanach B C *et al* 2016 Prompt signals and displaced vertices in sparticle searches for next-to-minimal Gauge mediated supersymmetric models *Eur. Phys. J. C* **76** 482

[207] Abbiendi G *et al* 2003 Search for stable and longlived massive charged particles in $e + e -$ collisions at $s^{**}(1/2) = 130\text{--}209$ GeV *Phys. Lett. B* **572** 8–20

[208] Jittoh T *et al* 2006 Long life stau in the minimal supersymmetric Standard Model *Phys. Rev. D* **73** 055009
Jittoh T *et al* 2013 *Phys. Rev. D* **87** 019901 (erratum)

[209] Brandenburg A *et al* 2005 Signatures of axinos and gravitinos at colliders *Phys. Lett. B* **617** 99–111

[210] Heisig J *et al* 2014 A survey for low stau yields in the MSSM *J. High Energy Phys.* **JHEP04(2014)053**

[211] Bajc B and Senjanovic G 2007 Seesaw at LHC *J. High Energy Phys.* **JHEP08(2007)014**

[212] Bajc B, Nemevsek M and Senjanovic G 2007 Probing seesaw at LHC *Phys. Rev. D* **76** 055011

[213] Franceschini R, Hambye T and Strumia A 2008 Type-III see-saw at LHC *Phys. Rev. D* **78** 033002

[214] Arhrib A *et al* 2010 Collider signatures for heavy lepton triplet in type I+III seesaw *Phys. Rev. D* **82** 053004

[215] Aad G *et al* (ATLAS Collaboration) 2016 Search for magnetic monopoles and stable particles with high electric charges in 8 TeV pp collisions with the ATLAS detector *Phys. Rev. D* **93** 052009

[216] Khachatryan V *et al* (CMS Collaboration) 2016 Search for long-lived charged particles in proton–proton collisions at $\sqrt{s} = 13$ TeV *Phys. Rev. D* **94** 112004

[217] Khachatryan V *et al* (CMS Collaboration) 2015 Search for disappearing tracks in proton–proton collisions at $\sqrt{s} = 8$ TeV *J. High Energy Phys.* **JHEP01(2015)096**

[218] Aaboud M *et al* (ATLAS Collaboration) 2018 *J. High Energy Phys.* **JHEP06(2018)022**

[219] Sjostrand T, Mrenna S and Skands P Z 2008 A brief introduction to PYTHIA 8.1 *Comput. Phys. Commun.* **178** 852–67

[220] Sjostrand T *et al* 2015 An introduction to PYTHIA 8.2 *Comput. Phys. Commun.* **191** 159–77

[221] Buccella F, Farrar G R and Pugliese A 1985 R Baryon masses *Phys. Lett.* **153B** 311–4

[222] Farrar G R *et al* 2011 Limit on the mass of a long-lived or stable gluino *J. High Energy Phys.* **JHEP02(2011)018**

[223] Aaboud *et al* (ATLAS Collaboration) 2019 Search for heavy charged long-lived particles in the ATLAS detector in 36.1 fb^{-1} of proton–proton collision data at $\sqrt{s} = 13 \text{ TeV}$ *Phys. Rev. D* **99** 092007

[224] Alwall J *et al* 2007 A standard format for Les Houches event files *Comput. Phys. Commun.* **176** 300–4

[225] Degrande C *et al* 2012 UFO—the universal FeynRules output *Comput. Phys. Commun.* **183** 1201–14

[226] Alwall J *et al* 2014 The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations *J. High Energy Phys.* **JHEP07(2014)079**

[227] Artoisenet P *et al* 2013 Automatic spin-entangled decays of heavy resonances in Monte Carlo simulations *J. High Energy Phys.* **JHEP03(2013)015**

[228] Djouadi A *et al* 1998 The Minimal supersymmetric Standard Model: group summary report *GDR (Groupement De Recherche)—Supersymetrie Montpellier* 15–17 April, 1998 (FrancearXiv: [hep-ph/9901246](https://arxiv.org/abs/hep-ph/9901246))

[229] Berger C F *et al* 2009 Supersymmetry without prejudice *J. High Energy Phys.* **JHEP02(2009)023**

[230] Curtin D *et al* 2015 Illuminating dark photons with high-energy colliders *J. High Energy Phys.* **JHEP02(2015)157**

[231] Agostinelli S *et al* 2003 GEANT4: a simulation toolkit *Nucl. Instrum. Methods A* **506** 250–303

[232] Mackeprang R and Rizzi A 2007 Interactions of coloured heavy stable particles in matter *Eur. Phys. J. C* **50** 353–62

[233] Mackeprang R and Milstead D 2010 An updated description of heavy-hadron interactions in GEANT-4 *Eur. Phys. J. C* **66** 493–501

[234] Khachatryan V *et al* (CMS Collaboration) 2015 Search for decays of stopped long-lived particles produced in proton–proton collisions at $\sqrt{s} = 8 \text{ TeV}$ *Eur. Phys. J. C* **75** 151

[235] Sirunyan A M *et al* (CMS Collaboration) 2018 Search for decays of stopped exotic long-lived particles produced in proton–proton collisions at $\sqrt{s} = 13 \text{ TeV}$ *J. High Energy Phys.* **JHEP05** (2018)127

[236] Aad G *et al* (ATLAS Collaboration) 2013 Search for long-lived stopped R-hadrons decaying out-of-time with pp collisions using the ATLAS detector *Phys. Rev. D* **88** 112003

[237] Aaboud M *et al* (ATLAS Collaboration) 2018 Search for long-lived, massive particles in events with displaced vertices and missing transverse momentum in $\sqrt{s} = 13 \text{ TeV}$ pp collisions with the ATLAS detector *Phys. Rev. D* **97** 052012

[238] Alwall J *et al* 2011 MadGraph 5: going beyond *J. High Energy Phys.* **JHEP06(2011)128**

[239] Farina M and Low M 2017 Constraining quirky tracks with conventional searches *Phys. Rev. Lett.* **119** 111801

[240] Knapen S *et al* 2017 Tracking down quirks at the Large Hadron Collider *Phys. Rev. D* **96** 115015

[241] Aad G *et al* (ATLAS Collaboration) 2013 Triggers for displaced decays of long-lived neutral particles in the ATLAS detector *J. Instrum.* **8** P07015

[242] Knapen S *et al* 2017 Triggering soft bombs at the LHC *J. High Energy Phys.* **JHEP08(2017)076**

[243] ATLAS Collaboration 2018 Reinterpretation of searches for supersymmetry in models with variable R-parity-violating coupling strength and long-lived R-hadrons *ATLAS-CONF-2018-003* <http://cds.cern.ch/record/2308391>

[244] Lee L, Ohm C, Soffer A and Yu T T 2018 Collider searches for long-lived particles beyond the Standard Model *Prog. Part. Nucl. Phys.* **106** 210–55

[245] Fanti M 2017 Search for long-lived particles at LEP https://indico.cern.ch/event/517268/contributions/2041301/attachments/1272370/1886064/fanti_160512_LLWorkshop_LEP.pdf

[246] Aaboud M *et al* (ATLAS Collaboration) 2019 Search for long-lived neutral particles in pp collisions at $\sqrt{s} = 13 \text{ TeV}$ that decay into displaced hadronic jets in the ATLAS calorimeter *Eur. Phys. J. C* **79** 481

[247] Aad G *et al* (ATLAS Collaboration) 2015 Search for pair-produced long-lived neutral particles decaying in the ATLAS hadronic calorimeter in pp collisions at $\sqrt{s} = 8 \text{ TeV}$ *Phys. Lett. B* **743** 15–34

[248] Aad G *et al* (ATLAS Collaboration) 2015 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. D* **92** 012010

[249] Sirunyan A M *et al* (CMS Collaboration) 2018 Search for new long-lived particles at $\sqrt{s} = 13$ TeV *Phys. Lett. B* **780** 432–54

[250] Sirunyan A M *et al* (CMS Collaboration) 2019 Search for long-lived particles decaying into displaced jets in proton–proton collisions at $\sqrt{s} = 13$ TeV *Phys. Rev. D* **99** 032011

[251] Khachatryan V *et al* (CMS Collaboration) 2015 Search for long-lived neutral particles decaying to quark-antiquark pairs in proton–proton collisions at $\sqrt{s} = 8$ TeV *Phys. Rev. D* **91** 012007

[252] Sirunyan A M *et al* (CMS Collaboration) 2018 Search for long-lived particles with displaced vertices in multijet events in proton–proton collisions at $\sqrt{s} = 13$ TeV *Phys. Rev. D* **98** 092011

[253] Aaij R *et al* 2017 Updated search for long-lived particles decaying to jet pairs *Eur. Phys. J. C* **77** 812

[254] Aaij R *et al* (LHCb Collaboration) 2016 Search for Higgs-like bosons decaying into long-lived exotic particles *Eur. Phys. J. C* **76** 664

[255] ATLAS Collaboration 2017 Performance of the reconstruction of large impact parameter tracks in the ATLAS inner detector *ATLPHYS-PUB-2017-014* <https://cds.cern.ch/record/2275635>

[256] Aad G *et al* (ATLAS Collaboration) 2013 Triggers for displaced decays of long-lived neutral particles in the ATLAS detector *J. Instrum.* **8** P07015

[257] Aad G *et al* (ATLAS Collaboration) 2014 Standalone vertex finding in the ATLAS muon spectrometer *J. Instrum.* **9** P02001

[258] Aaboud M *et al* (ATLAS Collaboration) 2018 Search for long-lived particles produced in pp collisions at $\sqrt{s} = 13$ TeV that decay into displaced hadronic jets in the ATLAS muon spectrometer *Phys. Rev. D* **99** 052005

[259] Aaboud M *et al* (ATLAS Collaboration) 2019 Search for the production of a long-lived neutral particle decaying within the ATLAS hadronic calorimeter in association with a Z boson from pp collisions at $\sqrt{s} = 13$ TeV *Phys. Rev. Lett.* **122** 151801

[260] Csaki C *et al* 2015 Searching for displaced Higgs boson decays *Phys. Rev. D* **92** 073008

[261] Vasquez Sierra C 2017 Snapshot of current long-lived particle results at LHCb https://indico.cern.ch/event/649760/contributions/2689322/attachments/1542518/2419630/LHCb LLP_CVS.pdf

[262] Coccaro A *et al* 2016 Data-driven model-independent searches for long-lived particles at the LHC *Phys. Rev. D* **94** 113003

[263] Curtin D and Verhaaren C B 2015 Discovering uncolored naturalness in exotic Higgs decays *J. High Energy Phys.* **JHEP12(2015)072**

[264] Cottin G, Helo J C and Hirsch M 2018 Displaced vertices as probes of sterile neutrino mixing at the LHC *Phys. Rev. D* **98** 035012

[265] Khachatryan V *et al* (CMS Collaboration) 2015 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. D* **91** 052012

[266] CMS Collaboration 2015 Search for long-lived particles that decay into final states containing two muons, reconstructed using only the CMS muon chambers *CMS-PAS-EXO-14-012* <https://cds.cern.ch/record/2005761>

[267] Aaij R *et al* (LHCb Collaboration) 2015 Search for hidden-sector bosons in $B^0 \rightarrow K^* 0 \mu^+ \mu^-$ decays *Phys. Rev. Lett.* **115** 161802

[268] Aaij R *et al* (LHCb Collaboration) 2017 Search for long-lived scalar particles in $B^+ \rightarrow K^+ \chi(\mu^+ \mu^-)$ decays *Phys. Rev. D* **95** 071101

[269] Aaij R *et al* (LHCb Collaboration) 2018 Search for dark photons produced in 13 TeV pp collisions *Phys. Rev. Lett.* **120** 061801

[270] Khachatryan V *et al* (CMS Collaboration) 2016 A search for pair production of new light bosons decaying into muons *Phys. Lett. B* **752** 146–68

[271] Sirunyan A M *et al* (CMS Collaboration) 2019 A search for pair production of new light bosons decaying into muons in proton–proton collisions at 13 TeV *Phys. Lett. B* **796** 131–54

[272] Aad G *et al* (ATLAS Collaboration) 2014 Search for long-lived neutral particles decaying into lepton jets in proton–proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector *J. High Energy Phys.* **JHEP11(2014)088**

[273] Aad G *et al* 2019 Search for light long-lived neutral particles produced in pp collisions at $\sqrt{s} = 13$ TeV and decaying into collimated leptons or light hadrons with the ATLAS detector *ATLAS-CONF-2016-042* arXiv:1909.01246 [hep-ex]

[274] Aad G *et al* (ATLAS Collaboration) 2016 A search for prompt lepton-jets in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector *J. High Energy Phys.* **JHEP02(2016)062**

[275] Khachatryan V *et al* (CMS Collaboration) 2015 Search for displaced supersymmetry in events with an electron and a muon with large impact parameters *Phys. Rev. Lett.* **114** 061801

[276] Aaboud M *et al* (ATLAS Collaboration) 2019 Search for long-lived particles in final states with displaced dimuon vertices in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector *Phys. Rev. D* **99** 012001

[277] Aaij R *et al* (LHCb Collaboration) 2016 Tesla: an application for real-time data analysis in high energy physics *Comput. Phys. Commun.* **208** 35–42

[278] Sirunyan A M *et al* (CMS Collaboration) 2018 Search for heavy neutral leptons in events with three charged leptons in proton–proton collisions at $\sqrt{s} = 13$ TeV *Phys. Rev. Lett.* **120** 221801

[279] Aaij R *et al* (LHCb Collaboration) 2017 Search for massive long-lived particles decaying semileptonically in the LHCb detector *Eur. Phys. J. C* **77** 224

[280] Aaij R *et al* (LHCb Collaboration) 2014 Search for Majorana neutrinos in $B^- \rightarrow \pi^+ \mu^- \mu^-$ decays *Phys. Rev. Lett.* **112** 131802

[281] Shuve B and Peskin M E 2016 Revision of the LHCb limit on Majorana neutrinos *Phys. Rev. D* **94** 113007

[282] Gorbunov D and Shaposhnikov M 2007 How to find neutral leptons of the ν MSM? *J. High Energy Phys.* **JHEP10(2007)015**
Gorbunov D and Shaposhnikov M 2013 *J. High Energy Phys.* **JHEP11(2013)101** (erratum)

[283] Milanes D, Quintero N and Vera C E 2016 Sensitivity to Majorana neutrinos in $\Delta L = 2$ decays of B_c meson at LHCb *Phys. Rev. D* **93** 094026

[284] Mermod P 2017 Right-handed neutrinos: the hunt is on! *Proc. of Prospects in Neutrino Physics (NuPhys2016)* (arXiv:1704.08635NuPhys2016-Mermod)

[285] Antusch S, Cazzato E and Fischer O 2017 Sterile neutrino searches via displaced vertices at LHCb *Phys. Lett. B* **774** 114

[286] Drewes M *et al* 2020 A heavy metal path to new physics *Phys. Rev. Lett.* **124** 081801

[287] Gronau M, Leung C N and Rosner J L 1984 Extending limits on neutral heavy leptons *Phys. Rev. D* **29** 2539

[288] Lanfranchi G 2017 Search for hidden sector particles at NA62 *Report PoS EPS-HEP2017 (2017) 301* NA62 Collaboration DOI:10.22323/1.314.0301

[289] SHiP Collaboration 2015 A facility to Search for Hidden Particles (SHiP) at the CERN SPS *CERN-SPSC-2015-016* arXiv:1504.04956 [physics.ins-det]

[290] Bernardi G *et al* 1988 Further limits on heavy neutrino couplings *Phys. Lett. B* **203** 332

[291] CHARM Collaboration 1986 A search for decays of heavy neutrinos in the mass range 0.5–2.8 GeV *Phys. Lett. B* **166** 473

[292] NuTeV Collaboration 1999 Search for neutral heavy leptons in a high-energy neutrino beam *Phys. Rev. Lett.* **83** 4943

[293] DELPHI Collaboration 1997 Search for neutral heavy leptons produced in Z decays *Z. Phys. C* **74** 57

[294] Khachatryan V *et al* (CMS Collaboration) 2015 Search for heavy Majorana neutrinos in $\mu^\pm \mu^\pm$ + jets events in proton–proton collisions at $\sqrt{s} = 8$ TeV *Phys. Lett. B* **748** 144–66

[295] Canetti L *et al* 2013 Dark matter, baryogenesis and neutrino oscillations from right handed neutrinos *Phys. Rev. D* **87** 093006

[296] NA62 Collaboration 2017 The beam and detector of the NA62 experiment at CERN *J. Instrum.* **12** P05025

[297] Kling F and Trojanowski S 2018 Heavy neutral leptons at FASER *Phys. Rev. D* **97** 095016

[298] Asaka T and Shaposhnikov M 2005 The nuMSM, dark matter and baryon asymmetry of the Universe *Phys. Lett. B* **620** 17

[299] Mermod P 2017 Hidden sector searches with SHiP and NA62 *Proc. Int. Workshop on Neutrinos from Accelerators (NuFact2017)* arXiv:1712.01768 [hep-ex]

[300] Drewes M *et al* 2018 ARS leptogenesis *Int. J. Mod. Phys. A* **33** 1842002

[301] Aad G *et al* (ATLAS Collaboration) 2014 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. D* **90** 112005

[302] Aad G *et al* (ATLAS Collaboration) 2013 Search for nonpointing photons in the diphoton and E_T^{miss} final state in $\sqrt{s} = 7$ TeV proton–proton collisions using the ATLAS detector *Phys. Rev. D* **88** 012001

[303] CMS Collaboration 2015 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 *CMS-PAS-EXO-12-035* <https://cds.cern.ch/record/2063495>

[304] CMS Collaboration 2015 Search for displaced photons using conversions at 8 TeV *CMS-PAS-EXO-14-017* <https://cds.cern.ch/record/2019862>

[305] Dine M, Nelson A E and Shirman Y 1995 Low-energy dynamical supersymmetry breaking simplified *Phys. Rev. D* **51** 1362–70

[306] Chatrchyan S *et al* (CMS Collaboration) 2013 Searches for long-lived charged particles in pp collisions at $\sqrt{s} = 7$ and 8 TeV *J. High Energy Phys.* **JHEP07(2013)122**

[307] Aad G *et al* (ATLAS Collaboration) 2015 Searches for heavy long-lived charged particles with the ATLAS detector in proton–proton collisions at $\sqrt{s} = 8$ TeV *J. High Energy Phys.* **JHEP01(2015)068**

[308] Aaboud M *et al* (ATLAS Collaboration) 2019 Search for heavy charged long-lived particles in proton–proton collisions at $\sqrt{s} = 13$ TeV using an ionisation measurement with the ATLAS detector *Phys. Lett. B* **788** 96–116

[309] Aaij R *et al* (LHCb Collaboration) 2015 Search for long-lived heavy charged particles using a ring imaging Cherenkov technique at LHCb *Eur. Phys. J. C* **75** 595

[310] De Roeck A *et al* 2012 Sensitivity of LHC experiments to exotic highly ionising particles *Eur. Phys. J. C* **72** 1985

[311] Acharya B *et al* 2016 Search for magnetic monopoles with the MoEDAL prototype trapping detector in 8 TeV proton–proton collisions at the LHC *J. High Energy Phys.* **JHEP08(2016)067**

[312] Bruce R *et al* 2020 New physics searches with heavy-ion collisions at the LHC *J. Phys. G: Nucl. Part. Phys.* **47** 060501

[313] Hamaguchi K *et al* 2004 A Study of late decaying charged particles at future colliders *Phys. Rev. D* **70** 115007

[314] Asai S *et al* 2011 Stau Kinks at the LHC *J. High Energy Phys.* **JHEP12(2011)041**

[315] Jung S and Lee H-S 2017 Constraining Higgsino Kink tracks from existing LHC searches *Int. J. Mod. Phys. A* **32** 1750070

[316] Sirunyan A M *et al* (CMS Collaboration) 2019 Search for new particles decaying to a jet and an emerging jet *J. High Energy Phys.* **JHEP02(2019)179**

[317] Aad G *et al* (ATLAS Collaboration) 2013 Search for charginos nearly mass degenerate with the lightest neutralino based on a disappearing-track signature in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector *Phys. Rev. D* **88** 112006

[318] Schwaller P and Zurita J 2014 Compressed electroweakino spectra at the LHC *J. High Energy Phys.* **JHEP03(2014)060**

[319] Low M and Wang L-T 2014 Neutralino dark matter at 14 TeV and 100 TeV *J. High Energy Phys.* **JHEP08(2014)161**

[320] Barducci D *et al* 2015 Uncovering natural supersymmetry via the interplay between the LHC and direct dark matter detection *J. High Energy Phys.* **JHEP07(2015)066**

[321] ATLAS Collaboration 2017 Search for direct pair production of Higgsinos by the reinterpretation of the disappearing track analysis with 36.1 fb^{-1} of $\sqrt{s} = 13$ TeV data collected with the ATLAS experiment <https://cds.cern.ch/record/2297480>

[322] Sirunyan A M *et al* (CMS Collaboration) 2018 Search for disappearing tracks as a signature of new long-lived particles in proton–proton collisions at $\sqrt{s} = 13$ TeV *J. High Energy Phys.* **JHEP08(2018)016**

[323] ATLAS Collaboration 2019 July 2019 Summary Plot Update Tech. Rep. ATL-PHYS-PUB-2019-022 CERN Geneva <https://cds.cern.ch/record/2682063>

[324] Gershtein Y 2017 CMS hardware track triggers and LLPs at the HL-LHC <https://indico.cern.ch/event/649760/contributions/2748968/subcontributions/239747/attachments/1544494/2423617/Trieste-2017.pdf>

[325] Ray Holmes T 2017 ATLAS fast tracker (FTK)-Info, status, and prospects <https://indico.cern.ch/event/649760/contributions/2748920/attachments/1542791/2420193/17-10-13-LLPWorkshop.pdf>

[326] Horyn L 2017 ATLAS Fast TracKer (FTK)-Overview of triggering constraints https://indico.cern.ch/event/649760/contributions/2748921/attachments/1542873/2420596/Horyn_LLP_FTK.pdf

[327] Bullock J S 2010 Notes on the missing satellites problem arXiv:1009.4505 [astro-ph.CO]

[328] Boylan-Kolchin M, Bullock J S and Kaplinghat M 2011 Too big to fail? The puzzling darkness of massive Milky Way subhaloes *Mon. Not. R. Astron. Soc.* **415** L40–4

[329] Weinberg D H *et al* 2014 Cold dark matter: controversies on small scales *Proc. Natl Acad. Sci.* **112** 12249–55
Weinberg D H *et al* 2015 Cold dark matter: controversies on small scales *Proc. Natl Acad. Sci.* **112** 2249

[330] Famaey B and McGaugh S 2013 Challenges for Λ CDM and MOND *J. Phys.: Conf. Ser.* **437** 012001

[331] Schwaller P, Stolarski D and Weiler A 2015 Emerging jets *J. High Energy Phys.* **JHEP05** (2015)059

[332] Daci N *et al* 2015 Simplified SIMPs and the LHC *J. High Energy Phys.* **JHEP11(2015)108**

[333] Abazov V M *et al* 2010 Search for new fermions ('Quirks') at the fermilab tevatron collider *Phys. Rev. Lett.* **105** 211803

[334] Evans J A and Luty M A 2019 Stopping quirks at the LHC *J. High Energy Phys.* **JHEP06** (2019)090

[335] CMS Collaboration 2015 Muon reconstruction and identification improvements for run-2 and first results with 2015 run data *CMS-DP-2015-015* <http://cds.cern.ch/record/2037372>

[336] Seez C *et al* (CMS Collaboration) 2017 The phase-2 upgrade of the CMS endcap calorimeter *CERNLHCC- 2017-023. CMS-TDR-019* <https://cds.cern.ch/record/2293646>

[337] Aaboud M *et al* (ATLAS Collaboration) 2016 A measurement of material in the ATLAS tracker using secondary hadronic interactions in 7 TeV pp collisions *J. Instrum.* **11** P11020

[338] CMS Collaboration 2010 Studies of tracker material *CMS-PASTRK-10-003* <https://cds.cern.ch/record/1279138>

[339] Sirunyan A M *et al* 2018 Precision measurement of the structure of the CMS inner tracking system using nuclear interactions *J. Instrum.* **13** P10034

[340] Alexander M *et al* 2018 Mapping the material in the LHCb vertex locator using secondary hadronic interactions *J. Instrum.* **13** P06008

[341] Ilten P *et al* 2016 Proposed inclusive dark photon search at LHCb *Phys. Rev. Lett.* **116** 251803

[342] de Favereau J *et al* 2014 DELPHES 3, A modular framework for fast simulation of a generic collider experiment *J. High Energy Phys.* **JHEP02(2014)057**

[343] Chatrchyan S *et al* (CMS Collaboration) 2012 Search for stopped long-lived particles produced in pp collisions at $\sqrt{s} = 7$ TeV *J. High Energy Phys.* **JHEP08(2012)026**

[344] Khachatryan V *et al* (CMS Collaboration) 2011 Search for stopped gluinos in pp collisions at $\sqrt{s} = 7$ TeV *Phys. Rev. Lett.* **106** 011801

[345] Aad G *et al* (ATLAS Collaboration) 2012 Search for decays of stopped, long-lived particles from 7 TeV pp collisions with the ATLAS detector *Eur. Phys. J. C* **72** 1965

[346] Schmidt B 2016 The high-luminosity upgrade of the LHC: physics and technology challenges for the accelerator and the experiments *J. Phys.: Conf. Ser.* **706** 022002

[347] Apollinari G *et al* 2015 High-luminosity Large Hadron Collider (HL-LHC) : preliminary design report *CERN-2015-005* <https://cds.cern.ch/record/2116337>

[348] CMS Collaboration 2017 The phase-2 upgrade of the CMS tracker *CMS-TDR-014* <https://cds.cern.ch/record/2272264>

[349] CMS Collaboration 2017 The phase-2 upgrade of the CMS barrel calorimeters *CERN-LHCC-2017-011. CMS-TDR-015* <https://cds.cern.ch/record/2283187>

[350] CMS Collaboration 2017 The phase-2 upgrade of the CMS muon detectors *CERN-LHCC-2017-012. CMS-TDR-016* <https://cds.cern.ch/record/2283189>

[351] CMS Collaboration 2017 Technical proposal for a MIP timing detector in the CMS experiment phase-2 upgrade *CERN-LHCC- 2017-027. LHCC-P-009* <https://cds.cern.ch/record/2296612>

[352] CMS Collaboration 2017 The phase-2 upgrade of the CMS L1 trigger interim technical design report *CERN-LHCC-2017-013. CMS-TDR-017* <https://cds.cern.ch/record/2283192>

[353] ATLAS Collaboration 2017 Technical design report for the ATLAS inner tracker pixel detector *CERN-LHCC-2017-021. ATLASTDR- 030* <http://cds.cern.ch/record/2285585>

[354] ATLAS Collaboration 2017 Technical design report for the Phase-II upgrade of the ATLAS TDAQ system *CERN-LHCC-2017-020. ATLAS-TDR-029* <https://cds.cern.ch/record/2285584>

[355] ATLAS Collaboration 2017 Technical design report for the Phase-II Upgrade of the ATLAS tile calorimeter *CERN-LHCC-2017- 019. ATLAS-TDR-028* <https://cds.cern.ch/record/2285583>

[356] ATLAS Collaboration 2017 Technical design report for the Phase-II upgrade of the ATLAS LAr calorimeter *CERN-LHCC-2017- 018. ATLAS-TDR-027* <https://cds.cern.ch/record/2285582>

[357] ATLAS Collaboration 2017 Technical design report for the Phase-II upgrade of the ATLAS muon spectrometer *CERN-LHCC-2017- 017. ATLAS-TDR-026* <https://cds.cern.ch/record/2285580>

[358] ATLAS Collaboration 2017 Technical design report for the ATLAS inner tracker strip detector *CERN-LHCC-2017-005. ATLASTDR-025* <https://cds.cern.ch/record/2257755>

[359] ATLAS Collaboration 2019 Search prospects for dark-photons decaying to displaced collimated jets of muons at HL-LHC *ATLPHYS-PUB-2019-002* <https://cds.cern.ch/record/2257755>

[360] Ruderman J T and Shih D 2010 Slepton co-NLSPs at the tevatron *J. High Energy Phys.* **JHEP11(2010)046**

[361] CMS Collaboration 2017 Estimated sensitivity for new particle searches at the HL-LHC *CMS-PAS-FTR-16-005* <http://cds.cern.ch/record/2274436>

[362] CMS Collaboration 2018 Search sensitivity for dark photons decaying to displaced muons with CMS at the high-luminosity LHC *CMS-PAS-FTR-18-002* <http://cds.cern.ch/record/2644533>

[363] Liu J, Liu Z and Wang L-T 2019 Enhancing long-lived particles searches at the LHC with precision timing information *Phys. Rev. Lett.* **122** 131801

[364] Gershtein Y 2017 CMS hardware track trigger: new opportunities for long-lived particle searches at the HL-LHC *Phys. Rev. D* **96** 035027

[365] CMS Collaboration 2018 First level track jet trigger for displaced jets at high luminosity LHC *CMS-PAS-FTR-18-018* <http://cds.cern.ch/record/2647987>

[366] Cacciari M, Salam G P and Soyez G 2008 The anti- k_t jet clustering algorithm *J. High Energy Phys.* **JHEP04(2008)063**

[367] Cacciari M, Salam G P and Soyez G 2012 Fastjet user manual *Eur. Phys. J. C* **72** 1896

[368] Khachatryan V *et al* (CMS Collaboration) 2016 Search for narrow resonances in dijet final states at $\sqrt{s} = 8$ TeV with the novel CMS technique of data scouting *Phys. Rev. Lett.* **117** 031802

[369] Sirunyan A M *et al* (CMS Collaboration) 2018 Search for narrow and broad dijet resonances in proton–proton collisions at $\sqrt{s} = 13$ TeV and constraints on dark matter mediators and other new particles *J. High Energy Phys.* **JHEP08(2018)130**

[370] Aaboud M *et al* (ATLAS Collaboration) 2018 Search for low-mass dijet resonances using trigger-level jets with the ATLAS detector in pp collisions at $\sqrt{s} = 13$ TeV *Phys. Rev. Lett.* **121** 081801

[371] Hochberg Y *et al* 2014 Mechanism for thermal relic dark matter of strongly interacting massive particles *Phys. Rev. Lett.* **113** 171301

[372] Cohen T, Lisanti M and Lou H K 2015 Semivisible jets: dark matter undercover at the LHC *Phys. Rev. Lett.* **115** 171804

[373] Bediaga I *et al* (LHCb Collaboration) 2012 Framework TDR for the LHCb upgrade: technical design report *CERN-LHCC- 2012-007. LHCb-TDR-12* <https://cds.cern.ch/record/1443882>

[374] Bediaga I *et al* (LHCb Collaboration) 2018 Physics case for an LHCb upgrade II—Opportunities in flavour physics, and beyond, in the HL-LHC era *Technical Report CERN-LHCC-2018-027 LHCb-PUB-2018-009 CERN arXiv:1808.08865*

[375] LHCb Collaboration 2013 LHCb VELO upgrade technical design report *CERN-LHCC-2013-021* <https://cds.cern.ch/record/1624070>

[376] LHCb Collaboration 2014 LHCb tracker upgrade technical design report *CERN-LHCC-2014-001* <https://cds.cern.ch/record/1647400>

[377] LHCb Collaboration 2014 LHCb trigger and online upgrade technical design report *CERN-LHCC-2014-016* <https://cds.cern.ch/record/1701361>

[378] LHCb Collaboration 2017 Upgrade trigger: biannual performance update *LHCb-PUB-2017-005. CERN-LHCb-PUB-2017-005* <https://cds.cern.ch/record/2244312>

[379] LHCb Collaboration 2020 Trigger schemes <http://lhcb.web.cern.ch/lhcb/speakersbureau/html/TriggerScheme.html>

[380] Albrecht J *et al* 2014 The upgrade of the LHCb trigger system *J. Instrum.* **9** C10026

[381] Abba A *et al* 2015 Simulation and performance of an artificial retina for 40 MHz track reconstruction *J. Instrum.* **10** C03008

[382] Ilten P *et al* 2015 Dark photons from charm mesons at LHCb *Phys. Rev. D* **92** 115017

[383] Pierce A *et al* 2018 Searching for confining hidden valleys at LHCb, ATLAS, and CMS *Phys. Rev. D* **97** 095033

[384] Aaij R *et al* LHCb Collaboration 2017 Expression of interest for a Phase-II LHCb upgrade: opportunities in flavour physics, and beyond, in the HL-LHC era *CERN-LHCC-2017-003* <https://cds.cern.ch/record/2244311>

[385] LHCb Collaboration 2018 Prospects for searches for long-lived particles after the LHCb detector upgrades *LHCb-CONF-2018-006*. *CERN-LHCb-CONF-2018-006* <http://cds.cern.ch/record/2649825>

[386] Aaboud M *et al* (ATLAS Collaboration) 2016 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. D* **93** 112015

[387] Aaij R *et al* (LHCb Collaboration) 2015 Search for long-lived particles decaying to jet pairs *Eur. Phys. J. C* **75** 152

[388] NA62 Collaboration 2017 The Beam and detector of the NA62 experiment at CERN *J. Instrum.* **12** P05025

[389] Alekhin S *et al* 2016 A facility to search for hidden particles at the CERN SPS: the SHiP physics case *Rep. Prog. Phys.* **79** 124201

[390] Fairbairn M *et al* 2007 Stable massive particles at colliders *Phys. Rep.* **438** 1–63

[391] Burdin S *et al* 2015 Non-collider searches for stable massive particles *Phys. Rep.* **582** 1–52

[392] Acharya B *et al* 2014 The physics programme of the MoEDAL experiment at the LHC *Int. J. Mod. Phys. A* **29** 1430050

[393] Joergensen M D *et al* 2012 Searching for magnetic monopoles trapped in accelerator material at the Large Hadron Collider arXiv:1206.6793 [physics.ins-det]

[394] De Roeck A *et al* 2012 Development of a magnetometer-based search strategy for stopped monopoles at the Large Hadron Collider *Eur. Phys. J. C* **72** 2212

[395] Dirac P A M 1931 Quantized singularities in the electromagnetic field *Proc. R. Soc. A* **133** 60–72

[396] Dirac P A M 1948 The theory of magnetic poles *Phys. Rev.* **74** 817–30

[397] Hooft G 't 1974 Magnetic monopoles in unified gauge theories *Nucl. Phys. B* **79** 276–84

[398] Polyakov A M 1974 Particle spectrum in the quantum field theory *JETP Lett.* **20** 194–5
Polyakov A M 1974 Particle spectrum in the quantum field theory *Pisma Zh. Eksp. Teor. Fiz.* **20** 430

[399] Vento V and Sari Mantovani V 2013 On the magnetic monopole mass arXiv:1306.4213 [hep-ph]

[400] Cho Y M and Maison D 1997 Monopoles in Weinberg–Salam model *Phys. Lett. B* **391** 360–5

[401] Bae W S and Cho Y M 2005 Finite energy electroweak dyon *J. Korean Phys. Soc.* **46** 791–804 arXiv:hep-th/0210299

[402] Cho Y M, Kimm K and Yoon J H 2016 Mass of the electroweak monopole *Mod. Phys. Lett. A* **31** 1650053

[403] Cho Y M, Kimm K and Yoon J H 2016 Gravitationally coupled electroweak monopole *Phys. Lett. B* **761** 203–6

[404] Ellis J, Mavromatos N E and You T 2016 The price of an electroweak monopole *Phys. Lett. B* **756** 29–35

[405] Barriola M and Vilenkin A 1989 Gravitational field of a global monopole *Phys. Rev. Lett.* **63** 341

[406] Drukier A K and Nussinov S 1982 Monopole pair creation in energetic collisions: is it possible? *Phys. Rev. Lett.* **49** 102

[407] Mazur P O and Papavassiliou J 1991 Gravitational scattering on a global monopole *Phys. Rev. D* **44** 1317–20

[408] Mavromatos N E and Papavassiliou J 2018 Singular lensing from the scattering on special space–time defects *Eur. Phys. J. C* **78** 68

[409] Mavromatos N E and Sarkar S 2017 Magnetic monopoles from global monopoles in the presence of a Kalb–Ramond field *Phys. Rev. D* **95** 104025

[410] Mavromatos N E and Sarkar S 2018 Regularized Kalb–Ramond magnetic monopole with finite energy *Phys. Rev. D* **97** 125010

[411] Zeldovich Y B and Khlopov M Y 1978 On the concentration of relic magnetic monopoles in the Universe *Phys. Lett.* **79B** 239–41

[412] Hill C T 1983 Monopolonium *Nucl. Phys. B* **224** 469–90

[413] Dubrovich V K 2002 Association of magnetic monopoles and antimonopoles in the early universe *5th Int. Conf. COSMION 2001 (Moscow, Russia, May 21–27, 2001)* **8N1** 122–5 <https://inspirehep.net/literature/606767>

[414] Rajantie A 2012 Introduction to magnetic monopoles *Contemp. Phys.* **53** 195–211

[415] Rajantie A 2016 The search for magnetic monopoles *Phys. Today* **69** 40–6

[416] Kephart T W, Leontaris G K and Shafi Q 2017 Magnetic monopoles and free fractionally charged states at accelerators and in cosmic rays *J. High Energy Phys.* **JHEP10(2017)176**

[417] Epele L N *et al* 2008 Monopolium: the key to monopoles *Eur. Phys. J. C* **56** 87–95

[418] Epele L N *et al* 2009 Monopolium production from photon fusion at the Large Hadron Collider *Eur. Phys. J. C* **62** 587–92

[419] Epele L N *et al* 2016 Can the 750 GeV enhancement be a signal of light magnetic monopoles? arXiv:1607.05592 [hep-ph]

[420] Kalliokoski M *et al* 2016 Turning the LHC ring into a new physics search machine arXiv:1604.05778 [hep-ex]

[421] Felea D, Mamuzic J, Maselek R *et al* 2020 Prospects for discovering supersymmetric long-lived particles with MoEDAL *Eur. Phys. J. C* **80** 431

[422] Acharya B *et al* 2018 Search for magnetic monopoles with the MoEDAL forward trapping detector in 2.11 fb^{-1} of 13 TeV proton–proton collisions at the LHC *Phys. Lett. B* **782** 510–6

[423] Acharya B *et al* 2017 Search for magnetic monopoles with the MoEDAL forward trapping detector in 13 TeV proton–proton collisions at the LHC *Phys. Rev. Lett.* **118** 061801

[424] Acharya B *et al* (MoEDAL) 2019 Magnetic monopole search with the full MoEDAL trapping detector in 13 TeV pp collisions interpreted in Photon-Fusion and Drell-Yan production *Phys. Rev. Lett.* **123** 021802

[425] Baines S *et al* 2018 Monopole production via photon fusion and Drell–Yan processes: MadGraph implementation and perturbativity via velocity-dependent coupling and magnetic moment as novel features *Eur. Phys. J. C* **78** 966

[426] Gould O and Rajantie A 2017 Magnetic monopole mass bounds from heavy ion collisions and neutron stars *Phys. Rev. Lett.* **119** 241601

[427] Aad G *et al* (ATLAS Collaboration) 2012 Search for magnetic monopoles in $\sqrt{s} = 7 \text{ TeV}$ pp collisions with the ATLAS detector *Phys. Rev. Lett.* **109** 261803

[428] Patrizii L and Spurio M 2015 Status of searches for magnetic monopoles *Annu. Rev. Nucl. Part. Sci.* **65** 279–302

[429] Abulencia A *et al* 2006 Direct search for Dirac magnetic monopoles in $p\bar{p}$ collisions at $\sqrt{s} = 1.96 \text{ TeV}$ *Phys. Rev. Lett.* **96** 201801

[430] Pinfold J 2017 The MoEDAL experiment at the LHC *EPJ Web Conf.* **145** 12002

[431] Pinfold J L 2018 Contribution to 4th lhc-llp workshop 2018 https://indico.cern.ch/event/744951/contributions/3171317/attachments/1739901/2816344/4th_LLP_Workshop_2018.pdf

[432] Prinz A A *et al* 1998 Search for millicharged particles at SLAC *Phys. Rev. Lett.* **81** 1175–8

[433] De Roeck A *et al* 2017 Searching for trapped magnetic monopoles in LHC accelerator material https://indico.cern.ch/event/623746/attachments/1427507/2190853/expression_of_interest_beam_pipe.pdf

[434] Ball A *et al* 2016 A letter of intent to install a milli-charged particle detector at LHC P5 arXiv:1607.04669 [physics.ins-det]

[435] Izaguirre E and Yavin I 2015 New window to millicharged particles at the LHC *Phys. Rev. D* **92** 035014

[436] Sher M and Stevens J 2018 Detecting a heavy neutrino electric dipole moment at the LHC *Phys. Lett. B* **777** 246–9

[437] Aielli G *et al* 2006 Layout and performance of RPCs used in the Argo-YBJ experiment *Nucl. Instrum. Methods A* **562** 92–6

[438] Iuppa R 2015 Potential of RPCs in cosmic ray experiments for the next decade *J. Instrum.* **10** C04044

[439] Curtin D and Peskin M E 2018 Analysis of long lived particle decays with the MATHUSLA detector *Phys. Rev. D* **97** 015006

[440] Fradette A and Pospelov M 2017 BBN for the LHC: constraints on lifetimes of the Higgs portal scalars *Phys. Rev. D* **96** 075033

[441] Co R T, D’Eramo F and Hall L J 2017 Gravitino or axino dark matter with reheat temperature as high as 10^{16} GeV *J. High Energy Phys.* **JHEP03(2017)005**

[442] Dev P S B, Mohapatra R N and Zhang Y 2017 Long lived light scalars as probe of low scale seesaw models *Nucl. Phys. B* **923** 179–221

[443] Evans J A 2018 Detecting hidden particles with MATHUSLA *Phys. Rev. D* **97** 055046

[444] Evans J A, Gori S and Shelton J 2018 Looking for the WIMP next door *J. High Energy Phys.* **JHEP02(2018)100**

[445] D’Agnolo R T *et al* 2018 Exponentially light dark matter from coannihilation *J. High Energy Phys.* **JHEP08(2018)079**

[446] Deppisch F F, Liu W and Mitra M 2018 Long-lived heavy neutrinos from Higgs decays *J. High Energy Phys.* **JHEP08(2018)181**

[447] Jana S, Okada N and Raut D 2018 Displaced vertex signature of type-I seesaw model *Phys. Rev. D* **98** 035023

[448] Lindner R 2017 private communication

[449] Willey R S and Yu H L 1982 Decays $k \rightarrow \pi \ell \ell \ell$ and limits on the mass of the neutral Higgs boson *Phys. Rev. D* **26** 3287–9

[450] Chivukula R and Manohar A V 1988 Limits on a light Higgs boson *Phys. Lett. B* **207** 86–90

[451] Grinstein B, Hall L J and Randall L 1988 Do B meson decays exclude a light Higgs? *Phys. Lett. B* **211** 363–9

[452] Lanfranchi G 2017 Sensitivity of the SHiP experiment to a light scalar particle mixing with the Higgs <https://cds.cern.ch/record/2243034>

[453] Sierra C 2017 Searches for Long-lived particles at LHCb *Talk at Workshop on the physics of HL-LHC, and perspectives at HE-LHC* https://indico.cern.ch/event/647676/contributions/2757734/attachments/1549579/2435644/HLLHC_LL_P_CVS.pdf

[454] Patrignani C *et al* 2016 Review of particle physics *Chin. Phys. C* **40** 100001

[455] Altmannshofer W *et al* 2018 *The Belle II Physics Book* ed E Kou and P Urquijo arXiv:1808.10567 [hep-ex].

[456] Schabinger R M and Wells J D 2005 A minimal spontaneously broken hidden sector and its impact on Higgs boson physics at the Large Hadron Collider *Phys. Rev. D* **72** 093007

[457] Gopalakrishna S, Jung S and Wells J D 2008 Higgs boson decays to four fermions through an abelian hidden sector *Phys. Rev. D* **78** 055002

[458] Meade P, Papucci M and Volansky T 2009 Dark matter sees the light *J. High Energy Phys.* **12** 052

[459] Gligorov V V *et al* 2019 Leveraging the ALICE/L3 cavern for long-lived particle searches *Phys. Rev. D* **99** 015023

[460] Aaboud M *et al* (ATLAS Collaboration) 2016 Measurement of the Inelastic proton–proton cross section at $\sqrt{s} = 13$ TeV with the ATLAS detector at the LHC *Phys. Rev. Lett.* **117** 182002

[461] Van Haevermaet H 2016 Measurement of the inelastic proton–proton cross section at $\sqrt{s} = 13$ TeV *PoS DIS2016* 198 arXiv:1607.02033 [hep-ex]

[462] Ariga A *et al* 2018 Letter of intent for FASER: forward search experiment at the LHC arXiv:1811.10243

[463] Ariga A *et al* 2019 FASER’s physics reach for long-lived particles *Phys. Rev. D* **99** 095011

[464] Ariga A *et al* 2018 Technical proposal for FASER: forward search experiment at the LHC arXiv:1812.09139

[465] Ariga A *et al* 2019 FASER: forward search experiment at the LHC arXiv:1901.04468

[466] Ferrari A *et al* 2005 FLUKA: a multi-particle transport code (program version 2005) *CERN Yellow Reports: Monographs*, CERN-2005-010; INFN-TC-2005-11; SLAC-R-773; CERN-2005-10 CERN, Geneva doi:10.5170/CERN-2005-010

[467] Böhnen T T *et al* 2014 The FLUKA code: developments and challenges for high energy and medical applications *Nucl. Data Sheets* **120** 211–4

[468] CERN Sources, Targets, and Interactions Group, *et al* 2018 Characterization of the radiation field for the FASER experiment

[469] Feng J L *et al* 2018 Dark Higgs bosons at the forward search experiment *Phys. Rev. D* **97** 055034

[470] Feng J L *et al* 2018 Axionlike particles at FASER: The LHC as a photon beam dump *Phys. Rev. D* **98** 055021

[471] Berlin A and Kling F 2019 Inelastic dark matter at the lhc lifetime frontier: ATLAS, CMS, LHCb, CODEX-b, FASER, and MATHUSLA *Phys. Rev. D* **99** 015021

[472] Batell B *et al* 2018 Flavor-specific scalar mediators *Phys. Rev. D* **98** 055026

[473] Bauer M, Foldenauer P and Jaeckel J 2018 Hunting all the hidden photons *J. High Energy Phys.* **07** 094

[474] Hochberg Y *et al* 2018 Strongly interacting massive particles through the axion portal *Phys. Rev. D* **98** 115031

[475] Giudice G F and Romanino A 2004 Split supersymmetry *Nucl. Phys. B* **699** 65–89
Giudice G F and Romanino A 2005 Split supersymmetry *Nucl. Phys. B* **706** 487

[476] Pagels H and Primack J R 1982 Supersymmetry, cosmology and new TeV physics *Phys. Rev. Lett.* **48** 223

[477] Covi L, Kim J E and Roszkowski L 1999 Axinos as cold dark matter *Phys. Rev. Lett.* **82** 4180–3

[478] Zurek K M 2014 Asymmetric dark matter: theories, signatures, and constraints *Phys. Rep.* **537** 91–121

[479] CERN 2020 Forum on the Interpretation of the LHC Results for BSM studies, <https://twiki.cern.ch/twiki/bin/view/LHCPhysics/InterpretingLHCResults>

[480] Okawa H 2011 Interpretations of SUSY searches in ATLAS with simplified models. in particles and fields, particles and fields *Proc., Meeting of the Division of the American Physical Society, DPF 2011* (9–13, August 2011) (Providence, USA arXiv:1110.0282)

[481] Chatrchyan S and (CMS Collaboration) 2013 Interpretation of searches for supersymmetry with simplified models *Phys. Rev. D* **88** 052017

[482] Kraml S *et al* 2014 SModelS: a tool for interpreting simplified-model results from the LHC and its application to supersymmetry *Eur. Phys. J. C* **74** 2868

[483] Ambrogi F *et al* 2018 SModelS v1.1 user manual: Improving simplified model constraints with efficiency maps *Comput. Phys. Commun.* **227** 72–98

[484] Ambrogi F *et al* 2020 SModelS v1.2: long-lived particles, combination of signal regions, and other novelties *Comput. Phys. Commun.* **251** 106848

[485] Papucci M *et al* 2014 Fastlim: a fast LHC limit calculator *Eur. Phys. J. C* **74** 3163

[486] Drees M *et al* 2015 CheckMATE: confronting your favourite new physics model with LHC data *Comput. Phys. Commun.* **187** 227–65

[487] Derckx D *et al* 2017 Check MATE 2: from the model to the limit *Comput. Phys. Commun.* **221** 383–418

[488] Conte E *et al* 2014 Designing and recasting LHC analyses with madanalysis 5 *Eur. Phys. J. C* **74** 3103

[489] Dumont B *et al* 2015 Toward a public analysis database for LHC new physics searches using MADANALYSIS 5 *Eur. Phys. J. C* **75** 56

[490] Buckley A *et al* 2013 Rivet user manual *Comput. Phys. Commun.* **184** 2803–19

[491] Balázs C *et al* 2017 ColliderBit: a GAMBIT module for the calculation of high-energy collider observables and likelihoods *Eur. Phys. J. C* **77** 795

[492] CMS Collaboration 2017 Simplified likelihood for the re-interpretation of public CMS results *CMS-NOTE-2017-001* <https://cds.cern.ch/record/2242860>

[493] Heisig J, Kraml S and Lessa A 2019 Constraining new physics with searches for long-lived particles: Implementation into SModelS *Phys. Lett. B* **788** 87–95

[494] Chang J 2018 Madanalysis5 implementation of CMS-EXO-16-022 doi:10.7484/INSPIREHEP. DATA.UFU4.99E3

[495] Khachatryan V and (CMS Collaboration) 2015 Constraints on the pMSSM, AMSB model and on other models from the search for long-lived charged particles in proton–proton collisions at $\sqrt{s} = 8$ TeV. *Eur. Phys. J. C* **75** 325

[496] Cranmer K and Yavin I 2011 RECAST: extending the impact of existing analyses *J. High Energy Phys.* **04** 038

[497] Belanger G *et al* 2015 Probing U(1) extensions of the MSSM at the LHC Run I and in dark matter searches *J. High Energy Phys.* **09** 151

[498] Barducci D *et al* 2016 Status and prospects of the nMSSM after LHC Run-1 *J. High Energy Phys.* **01** 050

[499] Spite F and Spite M 1982 Abundance of lithium in unevolved halo stars and old disk stars: interpretation and consequences *Astron. Astrophys.* **115** 357–66 <http://adsabs.harvard.edu/full/1982A&A...115..357S>

[500] Cyburt R H, Fields B D and Olive K A 2008 An Update on the big bang nucleosynthesis prediction for Li-7: The problem worsens *J. Cosmol. Astropart. Phys.* **JCAP11(2008)012**

[501] Sjöstrand T, Mrenna S and Skands P 2006 PYTHIA 6.4 physics and manual *J. High Energy Phys.* **JHEP05(2006)026**

[502] Bagnaschi E *et al* 2017 Likelihood analysis of supersymmetric SU(5) GUTs *Eur. Phys. J. C* **77** 104

[503] Heisig J *et al* 2017 Trilinear-Augmented gaugino mediation *J. High Energy Phys.* **JHEP05** (2017)003

[504] Brooijmans G *et al* 2017 Les Houches 2017: physics at TeV colliders new physics working group report arXiv:1803.10379

[505] CMS Collaboration 2014 Search for displaced SUSY in dilepton final states *Report CMS-PAS-B2G-12-024 CERN* <http://cds.cern.ch/record/1706154>

[506] Displaced SUSY Parametrisation Study For User (2014). <https://twiki.cern.ch/twiki/bin/view/CMSPublic/DisplacedSusyParametrisationStudyForUser>

[507] Graham P W *et al* 2012 Displaced supersymmetry *J. High Energy Phys.* **JHEP07(2012)149**

[508] CMS Collaboration 2013 Search for long-lived neutral particles decaying to dijets *CMS-PAS-EXO-12-038* <https://cds.cern.ch/record/1563591>

[509] Cui Y 2013 Natural baryogenesis from unnatural supersymmetry *J. High Energy Phys.* **JHEP12(2013)067**

[510] ATLAS Collaboration 2016 Search for long-lived neutral particles decaying into displaced lepton jets in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector *ATLAS-CONF-2016-042* **94** <http://cds.cern.ch/record/2206083>

[511] Barello G, Chang S and Newby C A 2016 Correlated signals at the energy and intensity frontiers from non-Abelian kinetic mixing *Phys. Rev. D* **94** 055018

[512] Arguelles C A *et al* 2017 Dark gauge bosons: LHC signatures of non-abelian kinetic mixing *Phys. Lett. B* **770** 101–7

[513] Dine M and Fischler W 1982 A phenomenological model of particle physics based on supersymmetry *Phys. Lett.* **110B** 227–31

[514] Freytsis M, Robinson D J and Tsai Y 2015 Galactic center gamma-ray excess through a dark shower *Phys. Rev. D* **91** 035028

[515] Freytsis M *et al* 2016 Gamma-rays from dark showers with twin Higgs models *J. High Energy Phys.* **JHEP05(2016)018**

[516] Cottin G 2018 Reconstructing particle masses in events with displaced vertices *J. High Energy Phys.* **JHEP03(2018)137**

[517] Park M and Zhao Y 2011 Recovering particle masses from missing energy signatures with displaced tracks arXiv:1110.1403

[518] ATLAS Collaboration 2013 Search for long-lived, heavy particles in final states with a muon and a multi-track displaced vertex in proton–proton collisions at $\sqrt{s} = 8$ TeV *ATLAS-CONF-2013-092* <http://cds.cern.ch/record/1595755>

[519] ATLAS Collaboration 2017 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, Auxiliary Information Report *SUSY-2016-08* https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PAPERS/SUSY-2016-08/hepdata_info.pdf (2017)

[520] Cacciari M and Salam G P 2008 Pileup subtraction using jet areas *Phys. Lett. B* **659** 119–26

[521] Bertolini D *et al* 2014 Pileup per particle identification *J. High Energy Phys.* **JHEP10(2014)059**

[522] Cacciari M, Salam G P and Soyez G 2015 SoftKiller a particle-level pileup removal method *Eur. Phys. J. C* **75** 59

[523] Conte E, Fuks B and Serret G 2013 MadAnalysis 5, A user-friendly framework for collider phenomenology *Comput. Phys. Commun.* **184** 222–56

[524] Cranmer K, Heinrich L and (ATLAS Collaboration) 2018 Analysis preservation and systematic reinterpretation within the ATLAS experiment *J. Phys. Conf. Ser.* **1085** 042011

[525] Chen X *et al* 2016 *CERN Analysis Preservation: A Novel Digital Library Service to Enable Reusable and Reproducible Research and Advanced Technology for Digital Libraries: 20th International Conf. on Theory and Practice of Digital Libraries, TPDL 2016, (5–9 September, 2016, Hannover, Germany)* ed N Fuhr *et al* (Cham: Springer International Publishing) pp 347–56 DOI:10.1007/978-3-319-43997-6_27

[526] Cranmer K and Heinrich L 2017 Yadage and Packtivity—analysis preservation using parametrized workflows *J. Phys.: Conf. Ser.* **898** 102019

[527] Amstutz P *et al* 2016 Common workflow language v1.0, https://figshare.com/articles/Common_Workflow_Language_draft_3/3115156/2

[528] Aad G and (ATLAS Collaboration) 2015 Summary of the ATLAS experiments sensitivity to supersymmetry after LHC Run 1-interpreted in the phenomenological MSSM *J. High Energy Phys.* **JHEP10(2015)134**

[529] Aad G and (ATLAS Collaboration) 2014 Search for direct production of charginos, neutralinos and sleptons in final states with two leptons and missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector *J. High Energy Phys.* **JHEP05(2014)071**

[530] Aboud M and (ATLAS Collaboration) 2016 Dark matter interpretations of ATLAS searches for the electroweak production of supersymmetric particles in $\sqrt{s} = 8$ TeV proton–proton collisions *J. High Energy Phys.* **JHEP09(2016)175**

[531] ATLAS Collaboration 2016 A re-interpretation of $\sqrt{s} = 8$ TeV ATLAS results on electroweak supersymmetry production to explore general gauge mediated models *ATLAS-CONF-2016-033* <https://cds.cern.ch/record/2198316>

[532] Dobbs M and Hansen J B 2001 The HepMC C++ Monte Carlo event record for high energy physics *Comput. Phys. Commun.* **134** 41–6

[533] Marchesini G *et al* 1992 HERWIG: A Monte Carlo event generator for simulating hadron emission reactions with interfering gluons. Version 5.1-April 1991 *Comput. Phys. Commun.* **67** 465–508

[534] Gleisberg T *et al* 2009 Event generation with SHERPA 1.1 *J. High Energy Phys.* **JHEP02(2009)007**

[535] CMS Collaboration 2016 Search for pair-production of second-generation scalar leptoquarks in pp collisions at $\sqrt{s} = 13$ TeV with the CMS detector *CMS-PAS-EXO-16-007* <http://cds.cern.ch/record/2139349>

[536] Sirunyan A M *et al* 2019 Search for pair production of second-generation leptoquarks at $\sqrt{s} = 13$ TeV *Phys. Rev. D* **99** 032014

[537] ATLAS Collaboration 2014 Limits on metastable gluinos from ATLAS SUSY searches at 8 TeV *ATLAS-CONF-2014-037* <http://cds.cern.ch/record/1735199>

[538] Sirunyan A M and (CMS Collaboration) 2018 Search for natural and split supersymmetry in proton–proton collisions at $\sqrt{s} = 13$ TeV in final states with jets and missing transverse momentum *J. High Energy Phys.* **JHEP05(2018)025**

[539] Aboud M and (ATLAS Collaboration) 2018 Search for the Higgs boson produced in association with a vector boson and decaying into two spin-zero particles in the $H \rightarrow aa \rightarrow 4b$ channel in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector *J. High Energy Phys.* **JHEP10(2018)031**

[540] Juknevich J E, Melnikov D and Strassler M J 2009 A pure-glue hidden valley I states and decays *J. High Energy Phys.* **JHEP07(2009)055**

[541] Craig N *et al* 2016 The vector-like twin Higgs *J. High Energy Phys.* **JHEP07(2016)002**

[542] Renner S and Schwaller P 2018 A flavoured dark sector *J. High Energy Phys.* **JHEP08(2018)052**

[543] Izaguirre E and Stolarski D 2018 The platinum channel: Higgs decays to as many as 8 leptons *Phys. Rev. Lett.* **121** 221803

[544] Sterman G F and Weinberg S 1977 Jets from quantum chromodynamics *Phys. Rev. Lett.* **39** 1436

[545] Banks T and Zaks A 1982 On the phase structure of vector-like gauge theories with massless fermions *Nucl. Phys. B* **196** 189–204

[546] Larkoski A J 2017 An unorthodox introduction to QCD arXiv:[1709.06195](https://arxiv.org/abs/1709.06195)

[547] Polchinski J and Strassler M J 2000 The string dual of a confining four-dimensional gauge theory arXiv:[hep-th/0003136](https://arxiv.org/abs/hep-th/0003136)

[548] Klebanov I R and Strassler M J 2000 Supergravity and a confining gauge theory: duality cascades and χ SB-resolution of naked singularities *J. High Energy Phys.* **JHEP08(2000)052**

[549] Hofman D M and Maldacena J 2008 Conformal collider physics: energy and charge correlations *J. High Energy Phys.* **JHEP05(2008)012**

[550] Harnik R and Wizansky T 2009 Signals of new physics in the underlying event *Phys. Rev. D* **80** 075015

[551] Ellis R K, Stirling W J and Webber B R 1996 QCD and collider physics *Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology* vol 8 (Cambridge: Cambridge University Press) <https://doi.org/10.1017/CBO9780511628788>

[552] Giele W T, Kosower D A and Skands P Z 2008 A simple shower and matching algorithm *Phys. Rev. D* **78** 014026

[553] Randall L and Sundrum R 1999 A large mass hierarchy from a small extra dimension *Phys. Lett.* **83** 3370–3

[554] Csaki C, Reece M and Terning J 2009 The AdS/QCD correspondence: still undelivered *J. High Energy Phys.* **JHEP05(2009)067**

[555] Bjorken J D and Brodsky S J 1970 Statistical model for electron–positron annihilation into hadrons *Phys. Rev. D* **1** 1416–20

[556] Farhi E A 1977 QCD test for jets *Phys. Rev. Lett.* **39** 1587–8

[557] Fischer N *et al* 2016 Vincia for hadron colliders *Eur. Phys. J. C* **76** 589

[558] Morningstar C J and Peardon M J 1999 The Glueball spectrum from an anisotropic lattice study *Phys. Rev. D* **60** 034509

[559] Juknevich J E 2010 Pure-glue hidden valleys through the Higgs portal *J. High Energy Phys.* **JHEP08(2010)121**

[560] Clarke J D, Foot R and Volkas R R 2014 Phenomenology of a very light scalar (100 MeV $< m_h < 10$ GeV) mixing with the SM Higgs *J. High Energy Phys.* **JHEP02(2014)123**

[561] Ilten P *et al* 2018 Serendipity in dark photon searches *J. High Energy Phys.* **JHEP06(2018)004**

[562] Bauer M, Neubert M and Thamm A 2017 Collider probes of axion-like particles *J. High Energy Phys.* **JHEP12(2017)044**

[563] Deppisch F F, Bhupal Dev P S and Pilaftsis A 2015 Neutrinos and collider physics *New J. Phys.* **17** 075019

[564] Sirunyan A M and (CMS Collaboration) 2017 Search for heavy resonances that decay into a vector boson and a Higgs boson in hadronic final states at $\sqrt{s} = 13$ TeV *Eur. Phys. J. C* **77** 636

[565] Sirunyan A M and (CMS Collaboration) 2018 Search for massive resonances decaying into WW, WZ, ZZ, qW, and qZ with dijet final states at $\sqrt{s} = 13$ TeV *Phys. Rev. D* **97** 072006

[566] Thaler J and Van Tilburg K 2011 Identifying boosted objects with N-subjettiness *J. High Energy Phys.* **JHEP03(2011)015**

[567] Ellis S D, Roy T S and Scholtz J 2013 Jets and photons *Phys. Rev. Lett.* **110** 122003

[568] Ellis S D, Roy T S and Scholtz J 2013 Phenomenology of photon-jets *Phys. Rev. D* **87** 014015

[569] Aaij R *et al* (LHCb Collaboration) 2013 Search for the rare decay $K_S \rightarrow \mu^+ \mu^-$ *J. High Energy Phys.* **JHEP01(2013)090**

[570] Aaij R and (LHCb Collaboration) 2016 Measurement of CP violation in $B^0 \rightarrow D^+ D^-$ decays *Phys. Rev. Lett.* **117** 261801

[571] Nachman B and Rubbo F 2018 Search strategy using LHC pileup interactions as a zero bias sample *Phys. Rev. D* **97** 092002

[572] Falkowski A *et al* 2010 Discovering Higgs decays to lepton jets at hadron colliders *Phys. Rev. Lett.* **105** 241801

[573] Chatrchyan S and (CMS Collaboration) 2011 Search for light resonances decaying into pairs of muons as a signal of new physics *J. High Energy Phys.* **JHEP07(2011)098**

[574] Cohen T *et al* 2017 LHC searches for dark sector showers *J. High Energy Phys.* **JHEP11(2017)196**

[575] Park M and Zhang M 2017 Tagging a jet from a dark sector with jet-substructures at colliders arXiv:1712.09279

[576] Gallicchio J *et al* 2011 Multivariate discrimination and the Higgs +W/Z search *J. High Energy Phys.* **JHEP04(2011)069**

[577] Larkoski A J, Salam G P and Thaler J 2013 Energy correlation functions for jet substructure *J. High Energy Phys.* **JHEP06(2013)108**

[578] Larkoski A J, Moult I and Neill D 2014 Power counting to better jet observables *J. High Energy Phys.* **JHEP12(2014)009**

[579] Gras P *et al* 2017 Systematics of quark/gluon tagging *J. High Energy Phys.* **JHEP07(2017)091**

[580] Frye C *et al* 2017 Casimir meets poisson: improved quark/gluon discrimination with counting observables *J. High Energy Phys.* **JHEP09(2017)083**

[581] Dery L M *et al* 2017 Weakly supervised classification in high energy physics *J. High Energy Phys.* **JHEP05(2017)145**

[582] Cohen T, Freytsis M and Ostdiek B 2018 (Machine) learning to do more with less *J. High Energy Phys.* **JHEP02(2018)034**

[583] Metodiev E M, Nachman B and Thaler J 2017 Classification without labels: learning from mixed samples in high energy physics *J. High Energy Phys.* **JHEP10(2017)174**

[584] Aad G and (ATLAS Collaboration) 2016 Charged-particle distributions in $\sqrt{s} = 13$ TeV pp interactions measured with the ATLAS detector at the LHC *Phys. Lett. B* **758** 67–88

[585] Sirunyan A M and (CMS Collaboration) 2017 Particle-flow reconstruction and global event description with the CMS detector *J. Instrum.* **12** P10003

[586] Aad G and (ATLAS Collaboration) 2016 Measurement of the charged-particle multiplicity inside jets from $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector *Eur. Phys. J. C* **76** 322

[587] Rovere M 2015 CMS reconstruction improvements for the tracking in large pileup events *J. Phys.: Conf. Ser.* **664** 8

[588] Chatrchyan S and (CMS Collaboration) 2013 Identification of b-quark jets with the CMS experiment *J. Instrum.* **8** P04013

[589] Buschmann M *et al* 2015 Lepton jets from radiating dark matter *J. High Energy Phys.* **JHEP07** (2015)045

[590] Cheung C *et al* 2010 Lepton jets in (supersymmetric) electroweak processes *J. High Energy Phys.* **JHEP04(2010)116**

[591] Katz A and Sundrum R 2009 Breaking the dark force *J. High Energy Phys.* **JHEP06(2009)003**

[592] Bai Y and Han Z 2009 Measuring the dark force at the LHC *Phys. Rev. Lett.* **103** 051801

[593] Gupta A, Primulando R and Saraswat P 2015 A new probe of dark sector dynamics at the LHC *J. High Energy Phys.* **JHEP09(2015)079**

[594] Aad G and (ATLAS Collaboration) 2013 A search for prompt lepton-jets in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector *Phys. Lett. B* **719** 299–317

[595] Carloni L and Sjostrand T 2010 Visible effects of invisible hidden valley radiation *J. High Energy Phys.* **JHEP09(2010)105**

[596] Carloni L, Rathsman J and Sjostrand T 2011 Discerning secluded sector gauge structures *J. High Energy Phys.* **JHEP04(2011)091**

[597] Davoudiasl H, Lee H-S and Marciano W J 2012 Dark side of Higgs diphoton decays and muon g-2 *Phys. Rev. D* **86** 095009

[598] Endo M, Hamaguchi K and Mishima G 2012 Constraints on hidden photon models from electron g-2 and hydrogen spectroscopy *Phys. Rev. D* **86** 095029

[599] Agakishiev G *et al* 2014 Searching a dark photon with HADES *Phys. Lett. B* **731** 265–71

[600] Babusci D *et al* 2013 Limit on the production of a light vector gauge boson in phi meson decays with the KLOE detector *Phys. Lett. B* **720** 111–5

[601] Babusci D *et al* 2014 Search for light vector boson production in $e^+e^- \rightarrow \mu^+\mu^-\gamma$ interactions with the KLOE experiment *Phys. Lett. B* **736** 459–64

[602] Abrahamyan S *et al* 2011 Search for a new gauge boson in electron-nucleus fixed-target scattering by the APEX experiment *Phys. Rev. Lett.* **107** 191804

[603] Aubert B *et al* 2009 Search for dimuon decays of a light scalar boson in radiative transitions $\text{upsilon} \rightarrow \text{gamma A0}$ *Phys. Rev. Lett.* **103** 081803

[604] Lees J *et al* 2014 Search for a dark photon in e^+e^- collisions at BaBar *Phys. Rev. Lett.* **113** 201801

[605] Blümlein J and Brunner J 2011 New exclusion limits for dark gauge forces from beam-dump data *Phys. Lett. B* **701** 155–9

[606] Bjorken J D *et al* 2009 New fixed-target experiments to search for dark gauge forces *Phys. Rev. D* **80** 075018

[607] Bross A *et al* 1991 A search for shortlived particles produced in an electron beam dump *Phys. Rev. Lett.* **67** 2942–5

[608] Merkel H *et al* 2011 Search for light gauge bosons of the dark sector at the mainz microtron *Phys. Rev. Lett.* **106** 251802

[609] Davier M and Nguyen Ngoc H 1989 An unambiguous search for a light Higgs boson *Phys. Lett. B* **229** 150

[610] Blümlein J and Brunner J 2014 New exclusion limits on dark gauge forces from proton bremsstrahlung in beam-dump data *Phys. Lett. B* **731** 320–6

[611] Gnenienko S 2012 Constraints on sub-GeV hidden sector gauge bosons from a search for heavy neutrino decays *Phys. Lett. B* **713** 244–8

[612] Essig R *et al* 2010 Discovering new light states at neutrino experiments *Phys. Rev. D* **82** 113008

[613] Dent J B, Ferrer F and Krauss L M 2012 Constraints on light hidden sector gauge bosons from supernova cooling arXiv:1201.2683

[614] Dreiner H K *et al* 2014 Supernova constraints on MeV dark sectors from e^+e^- annihilations *Phys. Rev. D* **89** 105015

[615] Batley J R *et al* 2015 Search for the dark photon in π^0 decays *Phys. Lett. B* **746** 178–85

[616] Alloul A *et al* 2014 FeynRules 2.0-A complete toolbox for tree-level phenomenology *Comput. Phys. Commun.* **185** 2250–300

[617] Beauchesne H *et al* 2018 Collider phenomenology of hidden valley mediators of spin 0 or 1/2 with semivisible jets *J. High Energy Phys.* **JHEP08(2018)030**

[618] Hatta Y and Matsuo T 2009 Thermal hadron spectrum in $e + e^-$ annihilation from gauge/string duality *Phys. Rev. Lett.* **102** 062001