DISSERTATION

Exploring Long-Lived Particles decaying into Displaced Dimuons with the CMS experiment

Ausgeführt zum Zwecke der Erlangung des akademischen Grades

Doktor der technischen Wissenschaften

unter der Leitung von

Dipl. Ing. Dr. Robert Schöfbeck, Privatdozent

und

Dr. Alberto Escalante del Valle

als verantwortlich mitwirkender Assistent

am Institut für Hochenergiephysik der Östereichischen Akademie der Wissenschaften

unterstützt durch

den Österreichischen Wissenschaftsfonds (FWF) im Rahmen des Einzelprojekts P 33693

und des Doktoratskollegs 'Particles and Interactions' (DK-PI)

Eingereicht an der

Technischen Universität Wien

Fakultät für Physik

Von

Mangesh Sonawane, MSc

Matrikelnummer : 12046014

Wien, 18th September 2024

Mangesh Sonawane   Robert Schöfbeck

Declaration

I declare in lieu of oath, that I wrote this thesis and performed the associated research myself, using only literature cited in this volume. If text passages from sources are used literally, they are marked as such.
I confirm that this work is original and has not been submitted elsewhere for any examination, nor is it currently under consideration for a thesis elsewhere.
I acknowledge that the submitted work will be checked electronically/technically using suitable and state-of-the-art means (plagiarism detection software). On the one hand, this ensures that the submitted work was prepared according to the high-quality standards within the applicable rules to ensure good scientific practice 'Code of Conduct' at the TU Wien. On the other hand, a comparison with other student theses avoids violations of my personal copyright.

Mangesh Sonawane, MSc
September 2024

Abstract

Many well-motivated theories beyond the Standard Model of particle physics propose the existence of exotic-long lived particles. These long-lived particles could be produced in high-energy proton-proton collisions and go on to decay into known particles, which could be detected a macroscopic distance away from the proton-proton collision point. This thesis presents an inclusive search for long-lived exotic particles decaying to a pair of muons. The experimental signature consists of two oppositely charged muons originating from a common secondary vertex, displaced from the proton-proton interaction point by distances ranging between several hundred micrometers to several meters. The search uses data collected by the CMS experiment at the CERN LHC in proton-proton collisions at $\sqrt{s} = 13.6$ TeV in 2022, corresponding to an integrated luminosity of 36.6 fb^{-1}, using new trigger strategies designed to maximize sensitivity to events bearing such a signature. These trigger strategies employ new and existing reconstruction algorithms innovatively, benefiting from the higher resolution of the tracker as well as the extensive coverage of the muon system, to achieve lower trigger thresholds while simultaneously suppressing prompt backgrounds. The results of the search are interpreted in the framework of the hidden Abelian Higgs model, in which the Higgs boson decays to a pair of long-lived dark photons, which in turn decay into a pair of displaced muons, and of an R-parity violating supersymmetry model, in which long-lived neutralinos decay to a pair of muons and a neutrino. No statistically significant excess was observed over the predicted Standard Model background hypothesis, and limits were set on the model parameters. The results show substantial improvements with respect to a previous analysis performed using data taken at $\sqrt{s} = 13$ TeV, especially at low masses and long lifetimes, which are mainly driven by the improved triggers and a refined analysis strategy.

Kurzfassung

Viele plausible Theorien jenseits des Standardmodells der Elementarteilchenphysik sagen die Existenz exotischer, langlebiger Teilchen voraus. Diese langlebigen Teilchen könnten in hochenergetischen Proton-Proton-Kollisionen erzeugt werden und in bekannte Teilchen zerfallen, die in makroskopischer Entfernung zum Kollisionspunkt detektiert werden können. Diese Arbeit präsentiert eine inklusive Suche nach exotischen, langlebigen Teilchen, die in ein Paar Myonen zerfallen. Die experimentelle Signatur besteht aus zwei gegensätzlich geladenen Myonen, die aus einem Sekundärvertex stammen, welcher zwischen einigen Hundert Mikrometern und einigen Metern vom Proton-Proton-Kollisionspunkt entfernt liegt. Die Suche basiert auf Daten, die vom CMS-Experiment am CERN LHC in Proton-Proton-Kollisionen bei $\sqrt{s} = 13.6$ TeV in 2022 aufgenommen wurden und einer integrierten Luminosität von 36.6 fb^{-1} entsprechen. Dabei werden neue Trigger-Strategien eingesetzt, die darauf ausgelegt sind, die Sensitivität auf Ereignisse mit einer solchen Signatur zu maximieren. Diese Trigger-Strategien kombinieren neue und bestehende Rekonstruktionsalgorithmen auf innovative Weise, wodurch sie von der höheren Auflösung des Trackers sowie der großräumigen geometrischen Abdeckung des Myonsystems profitieren um Trigger-Schwellenwerte zu reduzieren und gleichzeitig Untergrundereignisse zu unterdrücken. Die Resultate der Suche werden zum einen im Kontext des "Hidden Abelian Higgs Model" interpretiert, in dem das Higgs-Boson in ein Paar dunkle Photonen zerfällt, welche an einem sekundären Vertex in ein Paar Myonen zerfallen. Zum anderen wird ein R-Parität-verletzendes Supersymmetrie-Modell getestet, in dem langlebige Neutralinos in ein Paar Myonen und ein Neutrino zerfallen. Da kein statistisch signifikanter Überschuss im Vergleich zur Standardmodellhypothese beobachtet wurde, wurden Grenzen der entsprechenden Modellparameter bestimmt. Durch die verbesserten Trigger und die verfeinerte Analysestrategie zeigt diese Arbeit im Vergleich zu einer vorherigen Analyse bei $\sqrt{s} = 13$ TeV wesentliche Verbesserungen, besonders bei kleinen Massen und großen Lebenszeiten.

I dedicate this thesis to my Gurus, chief among them my mother, my father, and my mentor and friend, Alberto.

Acknowledgements

I would like to express my deepest gratitude to everyone who were part of this long journey towards a PhD.

First and foremost, I would like to thank my advisor, Dr. Alberto Escalante del Valle. This thesis simply would not have been possible without your vision, steady guidance and unwavering confidence in me. You are one of the finest physicists I know, and more so, you are among the best human beings I have ever had the pleasure of knowing. I have learned not just physics, but also essential life skills such as team collaboration, resource and time management, presentation skills, and perhaps most importantly, communication. Your ability to appreciate my efforts and achievements while also conveying constructive criticism for further improvement is inspiring, and something I will strive to emulate in my own future career and life. You will always be a mentor to me, and now, I am proud to say, a close friend for a lifetime. I could honestly write a few hundred pages on the many, many things that I have learned from you. And perhaps, some day, I will.

Reaching this milestone would not have been possible without the guidance and advice of my mentors, Robert Schöfbeck, Ivan Mikulec, Wolfgang Adam and Manfred Jeitler. You have always been a steady lighthouse of incredible knowledge and experience in the misty oceans of physics. You have also been a tremendous source of gentle encouragement to me in times of self-doubt. I am privileged to have had this opportunity to learn from you.

I am also incredibly grateful to the members of my advisory committee, Prof. Florian Bernlochner and Dr. Sabine Kraml, for giving your time to evaluate my thesis and provide insightful comments and constructive feedback. Your suggestions have been invaluable.

This PhD was made possible through the Austrian Science Fund (FWF) project P33693[1]. I would also like to express my gratitude to Prof. Claudia Wulz for supervising me under the Doktoratskollegs ‘Particles and Interactions’ (DK-PI) programme, which allowed me to travel freely and often to attend conferences and schools, allowing me to learn much of what I know. I would like to extend my gratitude to the Institute for High Energy Physics (HEPHY) and our director Prof. Jochen Schieck for giving me the opportunity to pursue my research at the institute, and for being an excellent instructor, and to the admins Simone, Natalie and Barbara for their guidance, help and support at every step of the way. I would also like to thank Prof. Christoph Schwanda and Dr.

[1]https://www.doi.org/10.55776/

Claudius Krause for being such superb teachers and for the lovely conversations we've had along the way.

This search was a collaborative effort between our team at HEPHY[2] Vienna, and teams from UCLA[3], University of Oviedo, IFCA[4] Santander, Boston University, CIEMAT[5] Madrid. Therefore, I would like to express my heartfelt gratitude to the Displaced Dimuons analysis team : our experts Slava, Ansar, Bob, Santi, Efe, Jay and Pablo, as well as my peers, Celia, Alejandro and Ruben. I owe much of my understanding of this project to them. I have had an incredibly positive experience working with such a cohesive and well functioning team of brilliant physicists. If competitive physics research were a category in the Olympics a la synchronized swimming, I believe this team would have a fair shot at the gold medal! I would like to express my heartfelt gratitude to my predecessors Sebastian, Will and Riju for paving the path I walked on. Your theses helped me understand the complexities of the project, and went on to inspire much of this very thesis. I would also like to thank Elisa, Jaana and Caterina from the L1 trigger team, who taught me everything I needed to know about the trigger development, and thus played a significant role in achieving these results. I would like to extend my gratitude to the CMS collaboration and every member who is and was a part of this incredible enterprise. None of this would be possible without their contributions. To paraphrase Sir Isaac Newton, I am merely standing on the shoulders of giants.

I would like to express my gratitude towards my colleagues and peers in the CMS analysis group at HEPHY. Our postdocs : Suman, who helped me reach Vienna and get settled during the initial pandemic days; Dennis, who provided the German translation of the abstract; Ang, Lisa and Dietrich for being a veritable font of knowledge and expertise. It has been an incredible pleasure getting to know you and learn from you. And to my fellow students : Sahana, Cristina, Maryam, Kaan and Priya, you are wonderful, incredibly talented, and brilliant scientists. I truly admire your passion and focus, and there is so much I learned from watching you work. I am proud to count you among my closest friends. Also, thanks for laughing at my jokes occasionally, though I enjoy the groans as well. Heehee.

I would also like to thank my peers Geraldine and Huw for helping me navigate the final stages of the PhD, and advising me about the university requirements. My heartfelt thanks to Angira and Kari as well, your award-winning theses provided the inspiration and an excellent reference to base my own thesis on. My thanks also to Shubhangi for your support and encouragement, for being an inspiration, and helping me finish writing the thesis.

On a personal note, I am so very grateful to my friends here in Vienna for making this city so beautiful, and for providing the unwavering support and love to keep me going

[2]Institute of High Energy Physics, Austrian Academy of Sciences
[3]University of California, Los Angeles
[4]Instituto de Física de Cantabria
[5]Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas

through the difficult days. Arianna, Christian, Erich, Flint, Cristian, Yerke, Ion, Klea, Sabrina, Fionnuala, Laura, Rahul, Uroosa, and many many others. I cannot express how essential you are to me. I am so incredibly lucky to have met you, and to have your friendship. You are my home away from home. I love you all to bits.

I want to convey my love and gratitude to my friends Manish, Tanishq, Yash, Srujan, Satyaprasad and Aniruddh. You've been with me literally my whole life, and I am overjoyed that you are still an essential and irreplaceable part of my life. Thank you for over 25 years worth of encouragement and companionship. Anish, Prajeet, I met you just a few years ago, but it feels like I've known you forever. You are the real ones.

I would like to extend my gratitude also to my friends from my school days, Shubham, Alka, Aishvarya and Donear. You have achieved incredible heights in such varied professions, and I am privileged to have been a small part of that journey. My fond thanks also to Kinnari, for always being a source of calm wisdom and refuge. Thank you all for your friendship, advice and support.

To my friends and colleagues from IISER, I offer my salutations along with rib-crushing hugs. Rugved, Shreyas, Vignesh, Bhavin, Naven, Siddharth, Lokahith, Khilav, Lakshya, Karthik, Rohit, Bharat, Pratyush, Nidhi, Nandini, Snehal and Tanushree. I have learned so much from you, and I can honestly say that I have only reached this point because of your constant love and support, especially at the time when I truly needed it. You are my strength.

I am extremely grateful to my family for always being in my corner and cheering me on. My siblings : Sharvari, Sayali, Soham, Hrishikesh, Hema, Nandita and Asmita, every day you inspire me. My uncles and aunts : Aditi, Maya, Sanju, Ravi, and the Mahajans. You have raised me to be the person I am today, and I am proud to be your nephew.

I would like to thank everyone who were part of this journey but could not complete it with me. You have touched my life, however briefly, and made it richer. You helped me find a hidden strength within myself that I did not know I possessed, and taught me the value of real friendship.

गुरूर्ब्रह्मा गुरूर्विष्णुः गुरूर्देवो महेश्वरः ।
गुरूर्साक्षात परब्रह्म तस्मै श्री गुरवे नमः ।।

Guru is Brahma, Vishnu and Maheshwar. Guru is the One Above All, and to you I offer my salutations.

I would like to thank my gurus for imparting their wisdom, teaching me everything I know today, and making me who I am today. I thank Prof. Seema Sharma, Prof. Sourabh Dube and Prof. Sunil Mukhi for introducing me to the field of particle physics, and mentoring me at every step of the way. I am here because of you. I would like to convey my gratitude to Prof. Suneeta Vardarajan for being one of the finest professors I have ever had the privilege of learning from, and equipping me with the mathematical foundations necessary to do well in this field. I would also like to thank Dr. Nishita Desai for introducing me to the subject of long-lived particles, and setting me on the path to this PhD. My deepest thanks to Mr. Ravindra Devadiga, who introduced me to the subject I love so much, in his very first physics lecture when I was in the 9th grade. I remember every word you spoke vividly, even 15 years later, and I am here today because you inspired me.

Music is an essential part of my life, and helped me find the peace and focus necessary to complete this work. Therefore, I owe so much of my mental and spiritual well-being to my music gurus. I thank you, Farhan, for introducing me to the field, and teaching me the very basics despite your busy schedule, and Achintya for being an inspiration and teaching me the nuances of so many raagas. I would like to express my heartfelt gratitude to Mr. Deepak Bhanuse, Pt. Keshav Ginde, Ustad Chote Rahimat Khan, Pt. Umakant Gundecha, Pt. Ramakant Gundecha and Pt. Anant Gundecha. The knowledge and wisdom you have imparted to me have helped me find peace and strength within myself.

I have the deepest gratitude and love for my late grandmothers, Sushanta Nesarikar and Suneeta Sonawane. They were among my first gurus. Both teachers, they kindling within me a deep passion for knowledge and learning, a great love for my culture, and a strong moral compass. I miss them every day.

And finally, I would like to thank my very first gurus in life, my parents, Ranjana and Avinash Sonawane. You have taught me everything I know, You have been my pillars of strength and support through all of my decisions and my struggles and my triumphs. I owe you my very existence, everything I have ever been since then, and everything I will ever be. Mere words are inadequate to express the depth of my gratitude to you, but I hope these few will suffice for now. I hope I have made you proud. I love you so so very much.

I am because of you. Always.

Table of contents

List of figures

List of tables

Nomenclature

Acronyms / Abbreviations

BMTF Barrel Muon Track Finder

BSM Beyond the Standard Model

CERN *Conseil Européen pour la Recherche Nucléaire* (European Organization for Nuclear Research)

CMS Compact Muon Solenoid

CSC Cathode Strip Chambers

DAQ Data Acquisition

DSA Displaced Standalone

DT Drift Tube

ECAL Electromagnetic Calorimeter

EMTF Endcap Muon Track Finder

EW Electroweak

HAHM Hidden Abelian Higgs Model

HCAL Hadronic Calorimeter

HL-LHC High Luminosity LHC

HLT High Level Trigger

L1T Level 1 Trigger

LEP Large Electron-Positron Collider

LHC Large Hadron Collider

LINAC Linear Accelerator

LLP Long Lived Particles

LS Long Shutdown

MET Missing E_T, or missing transverse momentum

MIP Minimum Ionizing Particles

OMTF Overlap Muon Track Finder

PF Particle Flow

PSB Proton Synchrotron Booster

PS Proton Synchrotron

QCD Quantum Chromodynamics

QED Quantum Electrodynamics

QFT Quantum Field Theory

RF Radio Frequency

RMS Root Mean Square

RPC Resistive Plate Chambers

RPV R-parity violating

RSA Refitted Standalone

SA Standalone

SLC Stanford Linear Collider

SM Standard Model

SPS Super Proton Synchrotron

SSB Spontaneous Symmetry Breaking

SUSY Supersymmetry

WLCG Worldwide LHC Computing Grid

Chapter 1

Introduction

"What is the universe made of?"

This question has occurred to all of us at one point or another during our lives. Indeed, it encapsulates so much about human nature in a mere six words: the curiosity and the drive to understand ourselves and the world around us, to push the limits of our knowledge of the universe that we are a part of. While the answer is not a simple one, a part of the effort to arrive at it is to approach it in terms of building blocks. Is there something more fundamental that makes up everything in the universe, including us?

The concept of atoms dates back to ancient civilizations such as Greece and India. In fact, the word 'atom' is derived from the ancient Greek *atomos*, meaning 'indivisible'. For a long time, atoms were considered to be the building blocks of matter. However, it wasn't until the late $19^{\text{th}}$and early 20^{th}centuries that our current understanding of atoms took shape, thanks to the collective efforts of many great minds. As it turns out, atoms are not indivisible either but are composed of electrons, protons, and neutrons. Are these, then, the most elementary objects that make up the universe?

The short answer is no. The long answer is fascinating.

Four fundamental forces govern all interactions in the universe: the strong nuclear, the weak nuclear, the electromagnetic, and the gravitational forces. These forces govern the interactions between matter at various distances and energy scales. We are most familiar with two of these: the electromagnetic force – responsible for most physical phenomena that we can observe in our daily lives, such as the light that enables us to see, the electricity that powers our homes, and even our ability to touch objects; and the gravitational force – responsible for pulling objects towards the Earth and keeping planets in orbit around the Sun. However, the strong and weak nuclear forces are equally important. The strong nuclear force holds protons and neutrons together in the atomic nucleus, and is responsible for forming these nucleons from even more fundamental particles called 'quarks' and 'gluons'. The weak nuclear force is responsible for the radioactive decay of atoms and the thermonuclear fission reaction that powers the Sun. These particles, along with visible matter, constitute less than 5% of the total matter-energy content of the universe [1, 2].

The Standard Model (SM) of particle physics is a mathematical framework that paints a seemingly consistent and comprehensive picture of all known elementary particles and their interactions [3]. It is an effort to harmonize three of the fundamental forces (excluding gravity) in a common theoretical framework, to explain physical observations, make reliable predictions, and understand the evolution of the universe from almost the moment of the Big Bang till today. It is capable of describing a wide range of physical phenomena with an astonishing degree of accuracy (for example, the electron magnetic dipole moment was successfully described up to a precision of 11 significant digits). Its predictions have been repeatedly confirmed under the most stringent scientific scrutiny, including the existence of new particles that were discovered later. The discovery of the Higgs boson is one of the most recent confirmations of the SM predictions. It was discovered at the Large Hadron Collider (LHC) [4, 5] in 2012 by the ATLAS [6] and CMS [7] collaborations, decades after being hypothesized to explain the mechanism by which particles gain masses. In recent years, its couplings to gauge bosons were measured [8–10], and its couplings to third-generation fermions were directly observed [11–13].

Despite its triumphs, however, the SM is not a complete theory. There are several physical phenomena which it does not adequately explain. Several measurements are in tension with SM predictions, such as the recent observations of the anomalous magnetic moment of the muon [14], or the fact that neutrinos are observed to oscillate and have mass [15]. The SM is still unable to explain why the Higgs boson has a mass of only 125 GeV, indicating some kind of fine-tuning that cancels out the large expected radiative corrections to its mass. The SM also does not describe observed cosmological such as Dark Matter (DM), which is believed to make up 27% of the universe, or the fact that the remaining 68% of the universe is thought to be composed of Dark Energy – a mysterious force that might explain why the universe is expanding at an accelerated rate.

Several solutions have been proposed that address the above gaps in the SM, such as new particles that cancel out the large quantum corrections to the Higgs mass and 'correct' it to its observed value, such as Supersymmetry (SUSY) [16], or a new understanding of gravity (with extra dimensions [17]). Such models are collectively called Beyond the Standard Model (BSM) theories. However, as of today, there is no convincing experimental evidence for any of them, at least up until the energies probed so far, of the order of a few TeV.

Most initial efforts at the LHC approached the search for DM in the form of significant missing energy in the detectors, targeting mono-particle signatures, but did not find any statistically significant deviations from SM-only predictions. However, many DM scenarios postulate the existence of particles that mediate the interaction between the SM and the dark sector (e.g., dark photons [18, 19]). While invisible to our detectors, these mediators can have suppressed decays back into SM particles, resulting in them being long-lived particles (LLPs). These types of models have received growing attention in recent years, but the experimental phase space is still relatively unexplored [20, 21].

The work presented in this thesis aims to fill in the gap with a generic and inclusive search for pairs of displaced dimuons covering a large phase space in terms of LLP masses and lifetimes, and improve on the existing results by employing an innovative trigger strategy, motivated by lessons learned from previous analyses, that significantly increases sensitivity to a broad range of BSM physics models.

This thesis focuses on my significant contributions to the search – the design and development of the new triggers, and the analysis of the 'STA-STA' dimuon category. It is structured as follows:

- Chapter 2 introduces the theoretical formulation of the Standard Model, the motivation for long-lived particles, and the theoretical framework of the signal models used in this thesis.
- Chapter 3 introduces the LHC, the CMS detector, and the technical aspects relevant to this search.
- Chapter 4 discusses the data and simulation samples used in this search.
- Chapter 5 discusses the search, describing the trigger strategy and design and their performance in signal and background samples, followed by the analysis selection, background estimation, and the results.
- Chapter 6 provides the results of the search, and concludes with a summary as well as an outlook of the goals of the next phase of the project.

The work performed during this thesis resulted in contributions to the following publications and conferences:

Publications

[22] CMS Collaboration, "Search for long-lived particles decaying to a pair of muons in proton-proton collisions at $\sqrt{s} = 13$ TeV", *Journal of High Energy Physics* **2023**, 228 (2023), https://doi.org/10.1007/JHEP05(2023)228.

[23] CMS Collaboration, "Search for long-lived particles decaying to final states with a pair of muons in proton-proton collisions at $\sqrt{s} = 13.6$ TeV", *Journal of High Energy Physics* **2024**, 47 (2024), https://doi.org/10.1007/JHEP05(2024)047.

[24] CMS Collaboration, "Dark sector searches with the CMS experiment", CMS-EXO-23-005, CERN-EP-2024-106, *Submitted to Physics Reports*, `arXiv:`2405.13778.

CMS Technical Reports

- "Search for long-lived particles decaying to a pair of muons in proton-proton collisions at $\sqrt{s} = 13.6$ TeV with 2022 data", AN-23-050, November 2023.

Conferences

[25] **(Proceedings)** E. Augé, J. Dumarchez and J. Trân Thanh Vân, "Search for long-lived particles decaying to a pair of muons in proton-proton collisions at $\sqrt{s} = 13$ TeV", Rencontres de Moriond 2022: Proceedings of the ElectroWeak Session, La Thuile, ARISF, 2022, https://doi.org/10.58027/9qzt-x624.

[26] Mangesh Sonawane, on behalf of the CMS Collaboration, "Exploring long-lived particles decaying into Displaced Dimuons at $\sqrt{s} = 13.6$ TeV", FAKT Workshop 2024: Particle Physics Retreat, https://indico.cern.ch/event/1336414/contributions/5783068/

[27] Mangesh Sonawane, on behalf of the CMS Collaboration, "Exploring long-lived particles decaying into Displaced Dimuons at $\sqrt{s} = 13.6$ TeV", La Thuile 2024 - Les Rencontres de Physique de la Vallée d'Aoste, https://agenda.infn.it/event/38205/timetable/?view=standard#73-exploring-long-lived-partic

[28] Mangesh Sonawane, on behalf of the CMS Collaboration, "Exploring long-lived particles decaying into Displaced Dimuons at $\sqrt{s} = 13.6$ TeV", LLP2024 Tokyo, https://indico.cern.ch/event/1381368/contributions/5963316/.

Posters

- Mangesh Sonawane, “Search for Long Lived Particles in Higgs Decays with the CMS experiment”, Clues to a mysterious Universe - exploring the interfact of particle, gravity and quantum physics, Kitzbühel Austria, 2022.
- Mangesh Sonawane, on behalf of the CMS Collaboration, “Search for Displaced Dimuons at $\sqrt{s} = 13.6$ TeV in proton-proton collisions with the CMS experiment, DKPI, 2024.

Chapter 2

The Standard Model and Beyond

Science, distilled into its essence, is the quest to know the universe around us, not merely to understand it but also to use that understanding to predict what happens. The Standard Model of particle physics is our most refined understanding of the universe's fundamental particles and forces. It is capable of predicting, with astounding precision and detail, a breadth of observable physical phenomena on the most fundamental of scales. For all its triumphs, however, the SM does not paint a complete picture. There are conspicuous blanks – mysteries like dark matter, or the asymmetry between the matter-antimatter content of the universe, or why gravity behaves so differently when it comes to the quantum scale.

The following sections give an overview of the SM, introducing the elementary particles and the fundamental forces that govern their interactions. The mathematical formulation of the quantum field theory is briefly summarized in Appendix A, introducing the various parts of the SM and fitting them together. It does not provide a complete and detailed description of the theory, but rather introduces some key notions that are vital for the context of this work.

This chapter is based on impressions and explanations from several sources [29–33]. Further, this chapter briefly touches upon the inadequacies of the SM, and the efforts made to explain the as-yet unexplained observations of physical phenomena using extensions to the SM, called Beyond Standard Model theories. Long-lived particles naturally arise in several of these theories, and two such models – featuring LLPs decaying into muons – are introduced.

2.1 An Overview of Particles and Interactions

Based on their behavior in statistical ensembles, all known elementary particles can be classified into two kinds – **fermions** and **bosons**. These fundamental particles have an intrinsic quantum number **(spin)** which governs their behavior. According to the spin-statistics theorem, particles with integer spin (bosons) obey Bose-Einstein statistics and can simultaneously occupy the same quantum state. Meanwhile, particles with half-integer spins (fermions) follow Fermi-Dirac statistics and are governed by the Pauli exclusion principle, which prevents them from occupying the same quantum state. The SM fermions all carry a spin of 1/2. Among the bosons, the W and Z bosons and the photon have a spin of 1 and are known as 'vector' bosons, while the Higgs boson carries a spin of 0 and is a 'scalar' boson. This property leads to entirely different behavior – fermions make up the matter of the universe, and bosons mediate the interactions.

There are four known fundamental forces in nature :

- The **electromagnetic** force – mediated by the massless photon.
- The **weak** force – mediated by the massive W^+, W^- and Z^0 bosons.
- The **strong** force – mediated by the massless gluons.
- The **gravitational** force – mediated by hypothesized[1] massless gravitons.

All forces, except for gravity, are described in the Standard Model of elementary particles. Attempts to formulate a complete quantum theory of gravity have yet to provide testable predictions. However, the relative strength of gravitational interactions is vanishingly small compared to the others at collider energy scales. Therefore, it is possible to ignore gravity entirely and still obtain highly accurate predictions from the SM.

Fermions can be further classified into **quarks**, which carry a color charge and therefore participate in strong interactions, and **leptons**, which do not. Both quarks and leptons occur in three generations, each consisting of two types. Quarks occur as 'up' type, consisting of the **up, charm, and top** quarks and carrying a charge of $+{}^{2}\!/_{3}\ e$, or as 'down' type, consisting of the **down, strange, and bottom** quarks and carrying a charge of $-{}^{1}\!/_{3}\ e$, where e is in units of the charge of an electron. Similarly, leptons occur as the charged leptons (**electrons, muons, and taus**) carrying a charge of -1 e, and as neutrinos of the corresponding lepton flavor, which are neutral. The first generation of fermions make the basic building blocks of most of the matter we see around us: the up and down quarks make up the protons and neutrons that make up the nuclei of atoms, while electrons constitute the 'shell' around the nucleus.

Particles can also occur as antiparticles, identical to their regular counterparts, except that they carry the opposite physical quantum numbers. The exceptions are particles that are their own antiparticles (Majorana particles).

Mass is a property that seems quite intuitive and commonplace in our daily, macroscopic, and non-relativistic lives. However, in the world of fundamental particles and interactions, it carries great depth and far-reaching implications. Whether the particles have mass or not, and how much, has significant consequences for their interactions. Each interaction arises from a fundamental symmetry of the universe, and the experimental observation that the mediators of the weak interaction have a non-zero mass indicates

[1]Currently, we do not have any experiment sensitive enough to provide direct or indirect evidence for the existence of gravitons

that the related symmetry is broken. The mechanism for this symmetry breaking, called **spontaneous symmetry breaking (SSB)**, gives rise to the mass of all particles – called the Higgs-Englert-Brout mechanism – and predicts the existence of yet another fundamental particle, the Higgs boson. The existence of this new particle was the last piece of the puzzle at the turn of the century, and in fact, it was the search for the Higgs boson that was the driving force behind the existence of the LHC. Its subsequent discovery in 2012 was a resounding triumph and confirmed the Standard Model and its predictions.

Figure 2.1 illustrates the particles in the SM, strategically grouping them and displaying some of their important properties.

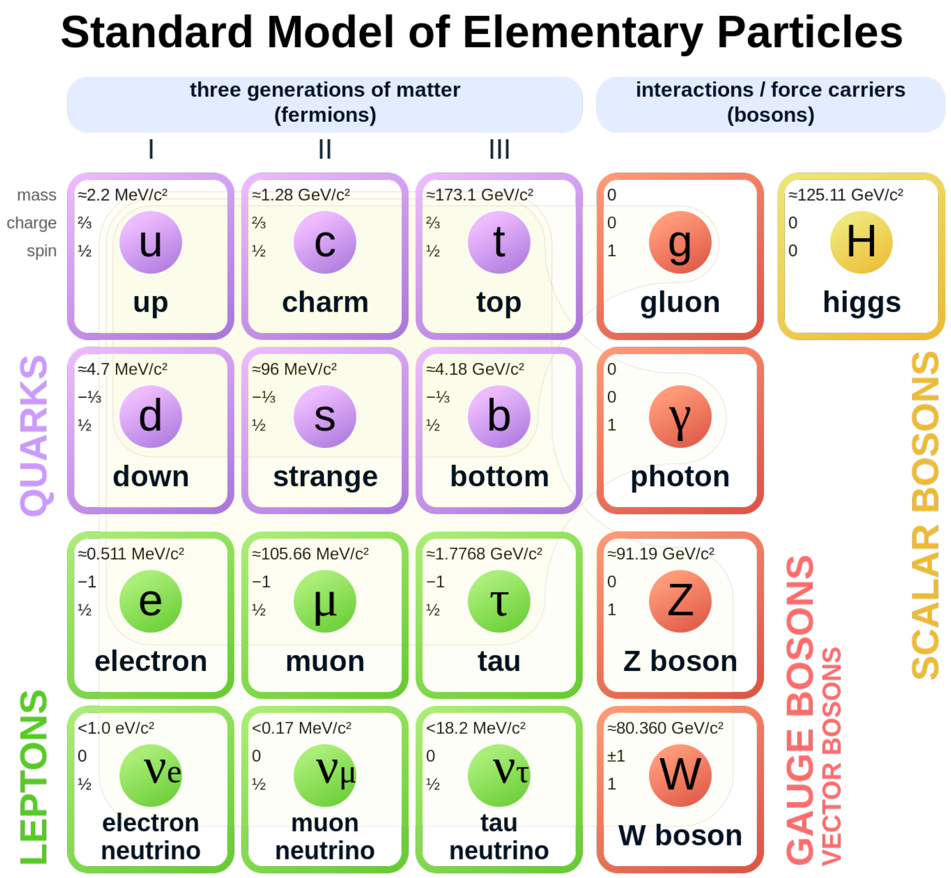

Figure 2.1: The Standard Model of elementary particles. Image source [34].

The mathematical formulation of the SM is based on quantum field theories (QFTs) that describe the fundamental particles and their interactions. In a QFT, particles are treated as excited states (quanta) of their underlying quantum fields. A topical introduction of this mathematical formulation is presented in Appendix A.

2.2 Beyond the Standard Model

The SM provides an elegant and coherent description of three of the four fundamental interactions in the universe under just one framework to an astounding degree of precision. However, despite these triumphs, it is not a complete theory in that it is unable to explain several observed phenomena. Some of these are briefly expanded upon below:

- Gravity does not fit into the SM, and all attempts thus far to include a complete quantum theory of gravity in the SM have failed. For the most part, the SM can get away with ignoring it and still accurately describe particle physics because the gravitational force is negligibly weak compared to the other forces in play (the electromagnetic, the strong, and the weak nuclear forces).

- This relative difference in interaction strength itself presents a problem, called the 'hierarchy problem': the weak interaction is 10^{24} times stronger than the gravitation force, a scale difference so large it is regarded as 'unnatural'. The weak interaction is mediated by the W and Z bosons, which gain their masses through the Higgs mechanism via their couplings to the Higgs field. Theoretically, quantum corrections to the Higgs mass are expected to be very large and around the Planck scale, $m_p = (\hbar c/G)^2 \equiv 1.2 \times 10^{19}$ GeV, the scale at which gravitational effects become significant. However, the Higgs boson mass is measured to be about 125 GeV, suggesting some sort of incredible and unnatural fine-tuning.

- Dark Matter (DM): cosmological observations indicate an abundance of DM in the universe. Compelling indirect evidence, such as the anisotropies in the Cosmic Microwave Background Radiation, or the rotational behavior of galaxies, suggests that about 27% of matter-energy content in the universe is dark matter [35]. Observations of the Bullet Cluster [36] or the MACS J0025.4-1222 galaxy cluster [37] even provide an expectation of the particle character of DM. However, the Standard Model does not provide a viable candidate for dark matter.

One can extend the Standard Model in a way that explains these observations while also predicting an observable signature in particle physics experiments. Such extensions are known as Beyond Standard Model theories, such as SUSY [38]. A large part of the research effort at the LHC is focused on investigating these BSM theories and searching for the new and exotic particles they predict. Most of these searches focus on 'prompt' phenomena featuring particles produced at or very close to the pp collision point, or on signatures with new hypothetical particles that could escape detection due to their weak interaction with the detectors. We expect energy and momentum to be conserved. Therefore, such particles escaping detection result in an imbalance in the momentum and energy recorded in the detector. Currently, no convincing evidence has been found to support BSM physics, at the LHC or anywhere else. However, another region of phase space contains a rich phenomenology, that of long-lived particles, where there remains much to explore.

The following sections introduce the concept of LLPs in BSM theories and motivate searches from such processes. I also discuss the powerful discovery potential of **displaced muons**, especially with the CMS experiment, and describe two well-established benchmark BSM signal models that are used in the context of the work presented in this thesis.

2.2.1 Long Lived Particles

The proper lifetime of a particle, τ, is inversely proportional to its total decay rate (also known as the 'decay width'), Γ

$$\tau = \frac{1}{\Gamma_{\text{total}}} \tag{2.1}$$

This width represents the probability that the particle will decay, per unit time. Particles often decay via several different channels, with the probability of the particle decaying via a particle channel i given by Γ_i. Thus

$$\Gamma_{\text{total}} = \sum_i \Gamma_i \tag{2.2}$$

The partial decay width can be computed using Fermi's Golden Rule [39].

$$\Gamma_i = \frac{1}{2m_X} \int d\Pi_f |\mathcal{M}(m_X \to \{p_f\})|^2 \tag{2.3}$$

where m_X gives the mass of the particle, $\mathcal{M}$ is the 'matrix element' for the particle's decay into the decay products $\{p_f\}$, and $d\Pi_f$ is the Lorentz-invariant phase space for the decay, assuming that $\hbar = c = 1$. The matrix element can be computed for a given process using Feynman rules.

Then, the probability of a particle surviving up to a time t is given by an exponential probability distribution with the mean parameterized by the mean proper lifetime of the particle, τ.

$$\mathcal{P}(t|\tau) = \frac{e^{-t/\tau}}{\tau} \tag{2.4}$$

The SM is composed of particles that have lifetimes ranging over many orders of magnitude – from the short-lived Z boson ($\tau \approx 2 \times 10^{-25}$ s), to the neutron ($\tau \approx 900$ s), to the proton ($\tau \gtrapprox 10^{34}$ years – more than 10^{24} times the age of the universe!), to the electron (which is stable). Figure 2.2 illustrates the lifetimes and masses of known particles (in units of mean proper decay length $c\tau$ [m]). Many of the SM particles decay within micrometers or much less, especially the heavy particles such as the Z or the Higgs bosons, or the top quark. In the context of collider experiments such as the CMS experiment, such particles are called **promptly** decaying.

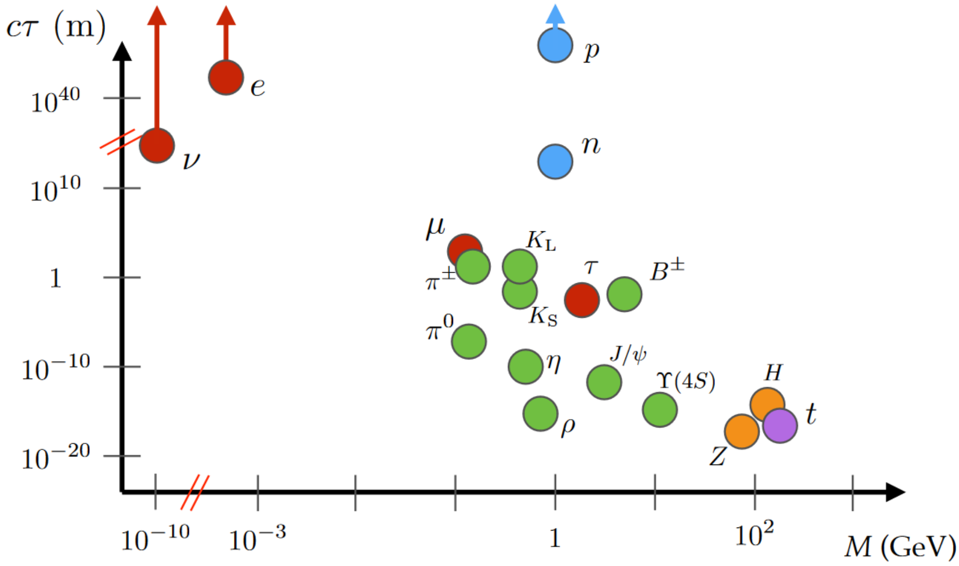

Figure 2.2: Lifetimes and masses of SM particles. Image source : [40]

Particles that don't decay immediately and leave identifiable signatures in the various components of the detector are considered as **stable** final-state particles, such as electrons, muons, or protons. Neutrinos are also stable particles, however they pass through the entire detector without interacting with any of the components, and thus, upon applying conservation of momentum in the transverse plane, show up as an imbalance in the sum of the transverse momenta of all recorded particles in the event. This imbalance is called **missing transverse momentum**[2], and is denoted by E_T^{miss}.

Particles in the intermediate lifetime range, that can traverse a measurable, macroscopic distance within the detector before decaying, are called **long-lived particles**, and have a mean proper decay length of between a few millimeters to a few kilometers. This is not an exact definition of LLPs, but is motivated by the capabilities and dimensions of existing experiments.

Long-lived particles arise naturally in BSM models, with new and exotic particles having a wide range of predicted lifetimes. While there are a wide range of BSM models, long lifetimes are typically generated by the following mechanisms [40]:

- A combination of scale hierarchies in interactions that mediate decays, for example, decays via off-shell mediators.

- small couplings between the long-lived particles and the SM particles, resulting in a small decay width into SM particles and longer lifetimes.

[2] E_T^{miss} is often abbreviated as 'missing E_T', or MET.

- Phase space considerations, such as a small mass splitting between particles and their decay products, or a large number (called multiplicity) of particles in the final state.

To illustrate an LLP generation mechanism, consider the following scenario: assume the existence of a heavy mediator X with a mass m_X, and two BSM particles Y and Z, with respective weak-scale masses m_Y and m_Z. The couplings between the BSM and SM particles are defined as g_X, g_Y, and g_Z. Then, the following configurations will lead to BSM models featuring particles with macroscopic lifetimes [41]. For a collection of references in each scenario, please see Section 2.2 of [40].

Decay via a heavy off-shell mediator

Consider the decay $Y \to X^* \to$ SM. If the mediator X is heavy (with a mass of up to or around the TeV scale), then $m_X/m_Y \gg 1$, leading to the lifetime of Y being determined by

$$c\tau_Y \approx \frac{1.2\text{ fm}}{g_X^4}\left(\frac{m_X}{m_Y}\right)^4\left(\frac{1\text{TeV}}{m_Y}\right) \tag{2.5}$$

This leads to scenarios such as

- $m_X \approx 100$ GeV, $m_Y \approx 10$ GeV, $g_X^4 \approx 10^{-7} \implies c\tau \approx 1$ cm (e.g heavy neutral lepton decays, etc)
- $m_X \approx 1$ TeV, $m_Y \approx 10$ GeV, $g_X^4 \approx 10^{-3} \implies c\tau \approx 1$ cm (e.g. Hidden Valley scenarios, split SUSY, etc.)

Small mass splitting between initial and final states

If the masses of the BSM particles Y and Z are nearly degenerate in the process $Y \to X^* \to Z+$SM, then $m_X/\Delta m \gg 1$ (with $\Delta m = m_Y - m_Z$). The lifetime of Y is then determined by

$$c\tau_Y \approx \frac{1.2\text{ fm}}{g_X^4}\left(\frac{m_X}{\Delta m}\right)^4\left(\frac{1\text{ TeV}}{\Delta m}\right) \tag{2.6}$$

This leads to scenarios such as

- $m_X \approx 100$ GeV, $\Delta m \approx 1$ GeV, $g_X^4 \approx 10^{-2} \implies c\tau \approx 1$ cm (e.g Compressed SUSY)

Small couplings

If the couplings between the BSM and SM particles in a process such as $Y \rightarrow \text{SM}$ is small, $g_Y \ll 1$, then the lifetime of Y is determined by

$$c\tau_Y \approx \frac{1.2 \text{ fm}}{g_Y^2}\left(\frac{1 \text{ TeV}}{m_Y}\right) \tag{2.7}$$

leading to scenarios such as

- $m_Y \approx 1$ GeV, $g_Y^2 \approx 10^{-12} \implies c\tau \approx 1$ cm (e.g dark photons)

In the above scenarios, one can obtain a wide range of Y lifetimes with plausible mass values and couplings. The exact configurations are, however, clearly model-dependent, and consequently, the lifetime of a new particle is commonly treated as a free parameter in experimental LLP searches.

Unlike SM particles with macroscopic lifetimes, which typically have masses roughly below 5 GeV and well-understood signatures, BSM LLPs can have unusual and exotic signatures in experiments. Some of these signatures, which can be probed with the CMS experiment, are illustrated in Fig 2.3. In principle, this makes any BSM processes involving LLPs potentially stand out from SM decays – some searches are even designed to be background-free for a wide range of search parameters.

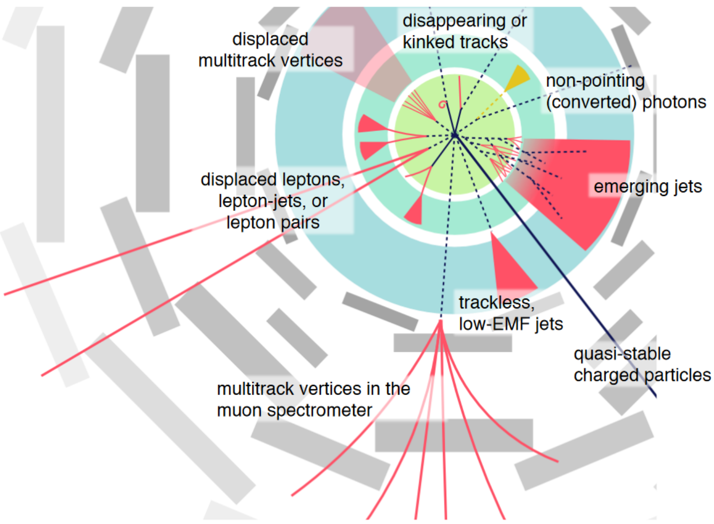

Figure 2.3: Illustration of a variety of plausible detector signatures involving BSM LLPs in a collider experiment like CMS (in the $\phi - r$ plane). Delayed signatures, such as stopped particles or delayed jets, are not shown. The figure is reproduced from Ref. [40].

However, the circumstances that make searches for unknown LLPs appealing also make the same searches experimentally challenging. The very fact that BSM LLP signatures are usually exotic makes it difficult or even impossible to use established and commonly used experimental methods, such as strategies for triggering, particle reconstruction, and physics simulation, especially in experiments originally designed for the study of prompt processes, like CMS or ATLAS. New methods must be developed at most stages of the data acquisition and processing pipeline to improve sensitivity to LLPs. Given the vastly different characteristics of LLP decays, each signature usually requires a dedicated analysis with dedicated methods.

2.3 Probing New Physics with Displaced Muons

The discussion so far suggests that there are a large number of BSM scenarios featuring LLPs that decay in a variety of ways into measurable final-state particles. The current experimental approach is **signature driven** rather than model-dependent to maximize the sensitivity to all possible models. This allows individual searches to cover much more LLP phase space when the experimental results are eventually reinterpreted for a different scenario, resulting in the same, or similar, final state. Moreover, relatively simple models are often sensitive to a large portion of the phase space. Therefore, the following sections focus on single decay modes – BSM decays that lead to one or more pairs of displaced muons ($X \to \mu^+\mu^-$) for a wide range of $c\tau(X)$.

We design the search using two common LLP benchmark models featuring displaced dimuons in the final state, resulting from BSM models using only a minimal set of assumptions. The results of the search presented in this thesis are interpreted in the context of these models, and used to set limits on the model parameters.

2.3.1 Hidden Abelian Higgs Model (HAHM); $h \to Z_D Z_D$

The **Hidden Abelian Higgs Model (HAHM)** [42, 43] is a widely used benchmark with a minimal set of BSM assumptions. The HAHM is part of a class of models featuring a 'hidden' or 'dark' sector of matter that does not interact directly with SM particles, but can be manifest through mixing effects. In this model, the SM is extended with a 'dark' $\mathrm{U}(1)_D$ gauge symmetry mediated by a new vector boson – a 'dark photon' Z_D (called the 'vector portal'). The SM fields are singlets under the new gauge symmetry (i.e., they do not participate in the interactions), and the only interactions between the Z_D and the SM hypercharge gauge bosons (the photon and the Z boson) occur through kinetic mixing with a coupling strength ϵ.

This new symmetry is spontaneously broken via a 'dark Higgs mechanism', analogous to the electroweak-symmetry-breaking Higgs mechanism in the SM, and gives rise to an associated 'dark Higgs boson' S that mixes with the 125 GeV SM-like Higgs boson H (the 'Higgs portal) and gives mass (m_{Z_D}) to the dark photon. In the absence of hidden-sector states with masses smaller than m_{Z_D}, the 'vector portal' mixing with the SM hypercharge bosons causes the dark photon to decay into SM particles exclusively, with a sizeable branching fraction to leptons.

Models featuring hidden sectors near the weak scale, such as the one realized in this model, are motivated by naturalness [44–48], thermal dark matter [19, 49, 50] and electroweak baryogenesis [51], and provide a possible answer to multiple open problems at the same time. Therefore, they represent a general expectation of BSM physics [52].

Using the notation defined in [42], the relevant gauge terms in the Lagrangian are denoted in Equation 2.8.

$$\mathcal{L} \subset -\frac{1}{4}\hat{B}_{\mu\nu}\hat{B}^{\mu\nu} - \frac{1}{4}\hat{Z}_{D\mu\nu}\hat{Z}_D^{\mu\nu} + \frac{1}{2}\frac{\epsilon}{\cos\theta}\hat{Z}_{D\mu\nu}\hat{B}^{\mu\nu} + \frac{1}{2}m_{D,0}^2\hat{Z}_D^{\mu}\hat{Z}_{D\mu} \tag{2.8}$$

where $\hat{B}_{\mu\nu} \equiv \partial_\mu\hat{B}_\nu - \partial_\nu\hat{B}_\mu$, $\hat{Z}_{D\mu\nu} \equiv \partial_\mu\hat{Z}_{D\nu} - \partial_\nu\hat{Z}_{D\mu}$, θ is the Weinberg mixing angle, ϵ is the kinetic mixing parameter between the dark photon Z_D and the SM hypercharge bosons, and $m_{D,0}$ is the dark photon mass before mixing with the SM fields.

The combined potential for the SM and the dark Higgs fields is

$$V_0(H,S) = -\mu^2|H|^2 + \lambda|H|^4 - \mu_S^2|S|^2 + \lambda_S|S|^4 + \kappa|S|^2|H|^2 \tag{2.9}$$

where H is the SM Higgs doublet, S is the dark Higgs boson, and κ is the Higgs portal coupling linking the dark and SM Higgs fields.

The phenomenology of the HAHM is thus defined by four free model parameters: the kinetic mixing ϵ between a broken dark Abelian gauge symmetry $\mathrm{U}(1)_D$ and the SM hypercharge $\mathrm{U}(1)_Y$, the dark photon mass m_D, after mixing with the SM fields, the Higgs portal coupling κ, the dark Higgs boson mass m_s.

The dark sector and the SM interact through two 'interaction portals', facilitated through the couplings ϵ and κ.

- **The 'Higgs portal'** kinetically mixes the SM Higgs boson h with the dark Higgs boson s, which subsequently decays into dark photons, which in turn eventually decays to final state leptons when $m_{Z_D} < m_h/2$, as shown in Figure 2.4 (right). The process $h \to Z_D Z_D$ is facilitated by the h-s mixing, which generates a non-zero hZ_DZ_D coupling.

  We assume dark Higgs to be much heavier than the SM Higgs ($m_s >> m_h$), to forbid the $h \to ss$ process, thereby avoiding strong constraints from direct searches [53]. We then obtain the partial width at lowest order in κ by

$$\Gamma(h \to Z_D Z_D) = (\kappa')^2\frac{1}{32\pi}\frac{v^2}{m_h}\sqrt{1 - \frac{4m_{Z_D}^2}{m_h^2}}\frac{\left(m_h^2 + 2m_{Z_D}^2\right)^2 - 8\left(m_h^2 - m_{Z_D}^2\right)m_{Z_D}^2}{m_h^4}, \tag{2.10}$$

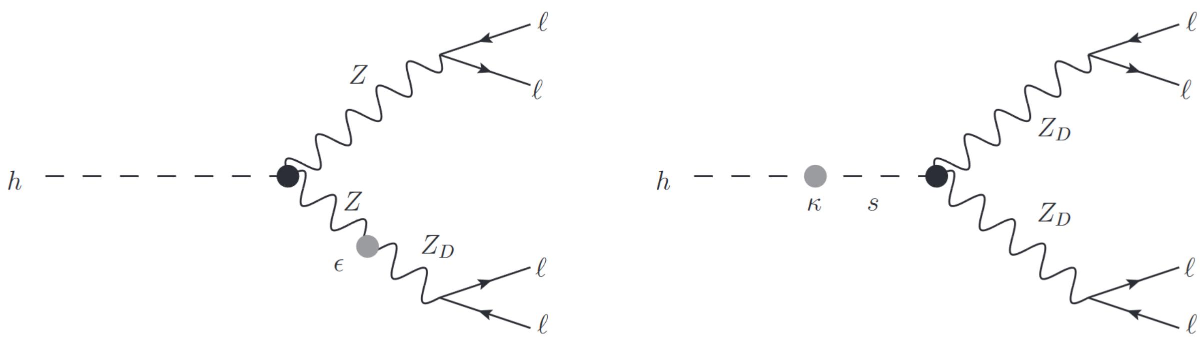

Figure 2.4: Feynman diagrams for the exotic Higgs boson s decays to four leptons in the final state, induced by intermediate dark photons in the HAHM model, for the $h \to Z_D Z^{(*)} \to 4l$ process in the vector-portal (left) and the $h \to Z_D Z_D \to 4l$ process via the Higgs portal (right).

where $v \approx 246$ GeV is the vacuum expectation value of the Higgs field, and κ' is defined as

$$\kappa' \equiv \kappa \frac{m_h^2}{|m_h^2 - m_s^2|}. \tag{2.11}$$

The dimensionless parameter κ' is introduced for convenience, as it encapsulates the dependence on m_s and κ in a single parameter that controls the rate of the exotic decay of the SM Higgs boson. Consequently, the branching fraction $\mathcal{B}(h \to Z_D Z_D)$ is proportional to $(\kappa')^2$ for a given m_{Z_D}. The value of $\mathcal{B}(h \to Z_D Z_D) = 1\%$ (0.01%) is compatible with the existing measurements of the Higgs boson properties, and corresponds to a value of $\kappa' = 10^{-2}$ (10^{-3}) for $m_{Z_D} = 10$ GeV.

- **The 'vector-portal'** (or 'hypercharge portal') kinetically mixes Z_D with the SM photon and Z bosons. If m_{Z_D} is indeed the lowest mass state in the dark sector (i.e., there exist no other lighter dark sector particles), then the dark photon Z_D will decay exclusively into SM particle final states, with a sizeable branching fraction to leptons. These decay and production mechanisms can give rise to dilepton signatures, pp $\to Z_D \to l^+l^-$. The branching fraction of the dark photon into leptons, $\mathcal{B}(Z_D \to l^+l^-)$, is generally determined by the dark photon mass.

  The Z_D lifetime, τ_{Z_D} is proportional to ϵ^{-2}. Figure 2.5 shows the relation between the dark photon mass, kinetic mixing and the dark photon lifetime, while Figure 2.6 shows the branching ratios, $\mathcal{B}(h \to Z_D Z_D \to 4l)$ and $\mathcal{B}(h \to Z_D^* \to 4l)$ for various values of ϵ and κ'. For sufficiently small values of ϵ, the Z_D can have macroscopically large lifetimes – such as if the dark photon has a mass 10 GeV $\lesssim m_{Z_D} < m_h/2$ with $\epsilon < \mathcal{O}(10^{-6})$, the dark photon would have a lifetime $c\tau_{Z_D} \gtrsim \mathcal{O}(100\ \mu\text{m})$. However, the phase space at smaller ϵ is inaccessible via vector-portal interactions, as the direct Z_D production cross-section scales with ϵ^{-2}. This phase space can be explored via the 'Higgs-portal'.

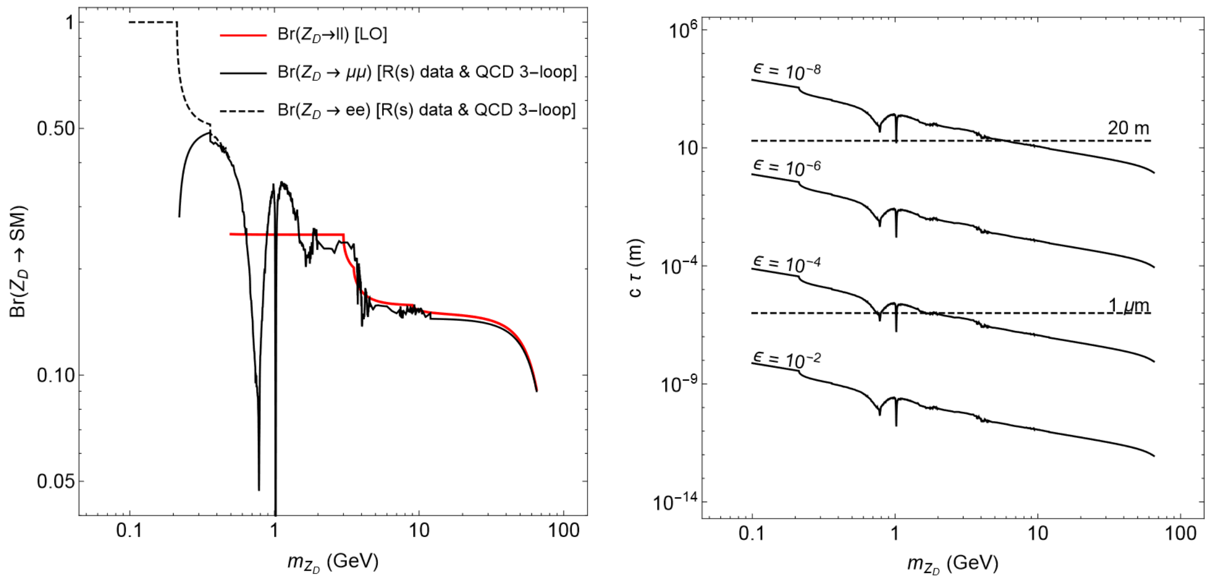

Figure 2.5: Distributions of the leptonic branching fraction of Z_D as a function of m_{Z_D} (left), and the mean proper decay length of Z_D as a function of m_{Z_D} for different values of ϵ (right), for the HAHM model. Figures sourced from [42].

It is important to note that $\mathcal{B}(h \to Z_D Z_D)$ does not depend on ϵ, which only controls the dark photon lifetime. Therefore, a search for the dark photon production via the Higgs portal is by far the most sensitive search of the $\mathrm{U}(1)_D$ hidden sector at $m_{Z_D} > \mathcal{O}(1\ \mathrm{GeV})$ in the long-lived regime with $\epsilon < \mathcal{O}(10^{-6})$.

The most stringent constraints in this ϵ range are set by CMS in two prior displaced muon searches, performed at $\sqrt{s} = 13$ TeV using data from Run 2 of the LHC. The predecessor to this analysis [22] set the (previously) world-leading constraints for 20 GeV $< m_{Z_D} < m_h/2$ for all $c\tau_{Z_D}$ values except those between ≈ 0.5 and 500 mm (depending on m_{Z_D}). In this lifetime region, the search is complemented by a CMS search using data collected with a dedicated high-rate data stream [54].

To summarize, the dark photon model extends the SM with a dark sector $U(1)_D$ gauge field and a scalar dark Higgs field, together with spontaneous symmetry breaking. This relatively minimal extension can give rise to distinctive collider experiment signatures involving muon pairs originating from vertices displaced macroscopically from the primary interaction point.

2.3.2 R-Parity Violating Supersymmetry; $\tilde{q} \to q\tilde{\chi}_1^0$

We also consider a simplified benchmark model inspired by **R-parity violating Supersymmetry (RPV SUSY)** [55, 56], featuring displaced $\mu\mu\nu$ vertices from the decay of long-lived neutralinos, $\tilde{\chi}_1^0$. This decay is mediated by the lepton-number-violating super-potential term $\frac{1}{2}\lambda_{ijk} L_i L_j \bar{E}_k$, resulting in two charged leptons and a neutrino of

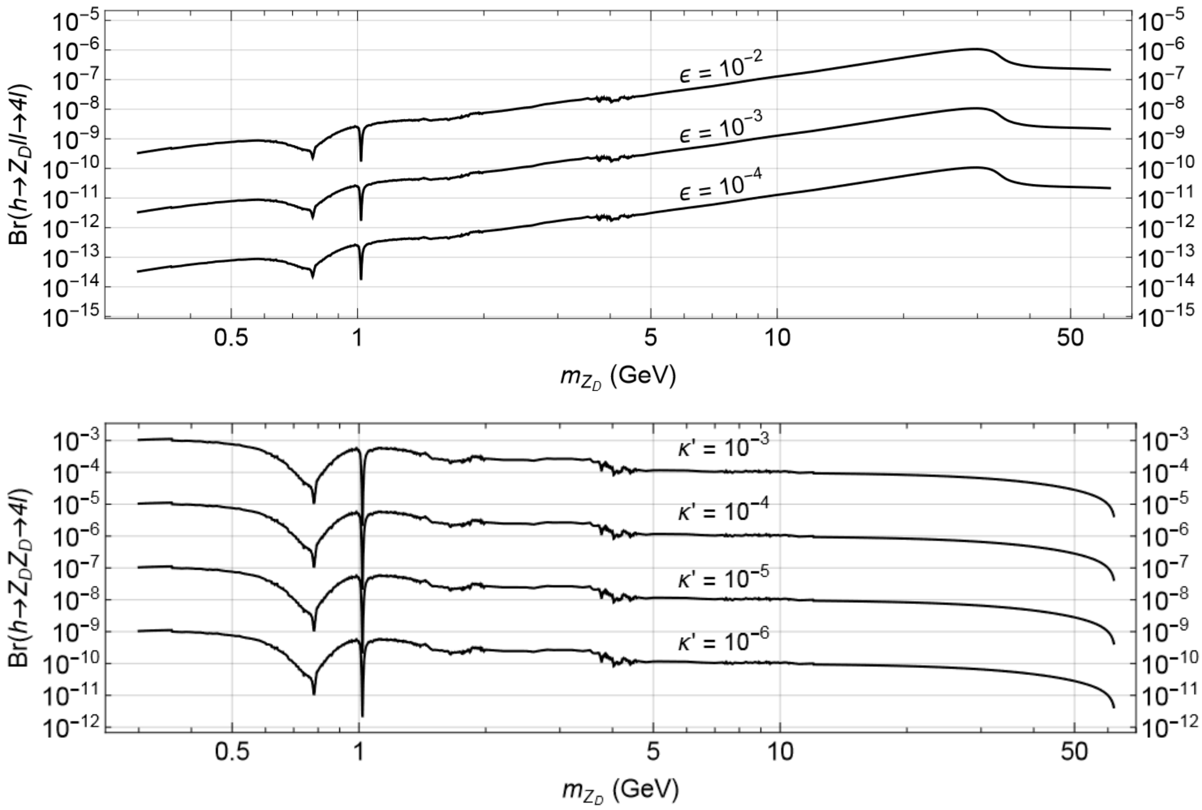

Figure 2.6: $\mathcal{B}(h \to Z_D Z^* \to 4l)$ (top) and $\mathcal{B}(h \to Z_D Z_D \to 4l)$ (bottom) for different values of ϵ and κ'. Figures sourced from [42].

generation i, j and k respectively,

$$\tilde{\chi}_1^0 \to l_k^{\pm} l_{i/j}^{\mp} \nu_{j/i}, \tag{2.12}$$

where L_i and $\bar{E}$ denote lepton SU(2)-doublet and singlet superfields respectively. The λ_{ijk} RPV couplings allow $\tilde{\chi}_1^0$ decays to every possible combination of charged-lepton pairs, where the branching ratio for each combination differs for each λ_{ijk}. The RPV couplings satisfy the relations $\lambda_{ijk} = -\lambda_{jik}$ and $i \neq j$.

Non-zero values of RPV couplings λ_{122} and λ_{232} enable displaced $\tilde{\chi}_1^0$ decays into a pair of muons and a neutrino, $\tilde{\chi}_1^0 \to \mu\mu\nu$. In contrast to the HAHM model, in which two-body LLP decays result in displaced $\mu\mu$ vertices that are reconstructed with a dimuon invariant mass $m_{\mu\mu}$ corresponding to the LLP mass, the neutrino from the $\tilde{\chi}_1^0$ decay leads to a non-peaking $m_{\mu\mu}$ distribution with a broad spectrum below the generated LLP mass. This model was used by CMS in searches for displaced dimuons in Run 1 [57, 58], and by ATLAS in the search for displaced charged leptons in Run 2 [59]. All possible charged lepton final states are listed in Table 3 of Ref [55].

The long-lived particle is assumed to be an RPV neutralino $\tilde{\chi}_1^0$ resulting from the decays of squarks, $\tilde{q} \to q\tilde{\chi}_1^0$, pair produced in pp collisions. Figure 2.7 shows the Feynman diagram for this process. The production cross section, $\sigma(\tilde{q}\bar{\tilde{q}})$ depends on the squark mass

$m(\tilde{q})$, and at $\sqrt{s} = 13.6$ TeV, this ranges from 836 pb for $m(\tilde{q}) = 200$ GeV, to 2.15 fb for $m(\tilde{q}) = 1.6$ TeV [60].

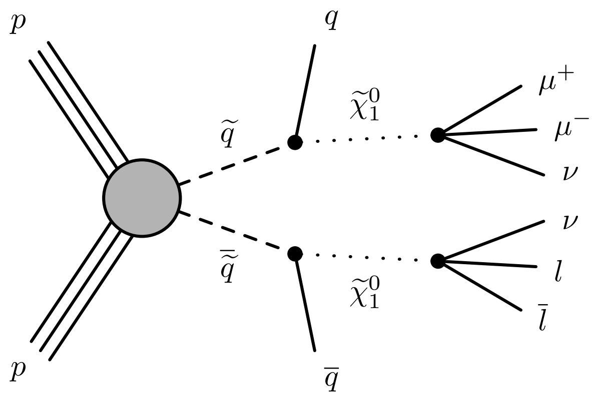

Figure 2.7: Feynman diagram for the pair production of squarks followed by their decay into neutralinos and SM quarks, $\tilde{q} \to q\tilde{\chi}_1^0$. The RPV Neutralino can be long-lived, and subsequently decays into two charged leptons and a neutrino $\tilde{\chi}_1^0 \to \mu\mu\nu$.

2.3.3 Heavy Scalar model; $H \to XX$

The Higgs portal scenario of the HAHM model (described in Section 2.3.1) postulates that dark photons are produced in the decay of the SM Higgs boson. However, the production of LLPs in the decays of SM Higgs bosons is not the only process accessible at the LHC; heavier mediators that decay into LLPs can also be produced. To explore a broader range of signal kinematics and topologies than those offered by HAHM, and to make the search less model-dependent, we consider an additional benchmark model with heavier mediators when optimizing the analysis. In this model [61], which was also used by the Run 1 [57, 58] and Run 2 [22] searches for displaced muons, the long-lived particle is an exotic spin-0 boson, denoted by X. The scalar X decays into dileptons with a non-zero branching fraction, and is pair-produced in the decay of a new heavy scalar boson H, which in turn is produced in gluon-gluon fusion, $gg \to H \to XX; X \to l^+l^-$. This model was instrumental in the design and optimization of the triggers and analysis selection.

2.3.4 Summary of signal models

For the displaced dimuon search presented here, we use the Higgs-portal production of dark photons in the HAHM scenario, $h \to Z_D Z_D$ with the subsequent resonant decay of the dark photon into muons $Z_D \to \mu\mu$. We use an R-parity violating Supersymmetric model, featuring a Neutralino decaying into two muons and a neutrino $\tilde{\chi}_1^0 \to \mu\mu\nu$, resulting in the interesting topology of a non-resonant dimuon invariant mass distribution. The results of the search are interpreted in terms of the above two models, and are used to set limits on the model parameters.

We also use a simplified benchmark model featuring a heavy scalar LLP X that can decay into muons, $H \to XX; X \to \mu\mu$, to design and optimize the analysis.

The signal samples chosen are designed to cover a wide range of LLP masses and lifetimes (detailed in Section 4.2) to maximize discovery potential across the entire available phase space.

Chapter 3

The LHC and the CMS Experiment

This thesis is based on experimental work conducted at the Compact Muon Solenoid (CMS) [62, 63] experiment, which operates at the CERN Large Hadron Collider [4, 5, 64, 65].

3.1 The Large Hadron Collider

The Large Hadron Collider (LHC) is the largest and most powerful particle accelerator, collider and storage ring in the world. It is operated by the European Organization for Nuclear Research (CERN) and is located at the France-Switzerland border near Geneva.

It was commissioned in 2008, after over a decade of planning and construction, and the first particle collisions at the LHC were reported in 2009. This section provides an overview of the CERN accelerator complex, focusing on the LHC, condensing and highlighting some interesting and relevant topics among the design principles, machine specifications, and operating parameters described in these comprehensive references.

The LHC is contained within a circular tunnel 26.7 km in circumference, at a depth of between 45m and 170m underground. This 3.8 m wide tunnel houses two adjacent beam pipes that contain beams that travel in opposite directions around the ring. To facilitate this, the interior of the beam pipes is operated at an ultrahigh vacuum. These beams of charged particles – protons or heavy ion nuclei such as lead – are accelerated, stored, and made to collide at four points around the ring where the beam pipes intersect. Four major particle detectors are situated at these points: ATLAS [66], LHCb [67], ALICE [68] and CMS[62, 63]. Among these, CMS and ATLAS are general-purpose detectors designed to study various phenomena. ALICE is dedicated to heavy ion collisions, and LHCb uses the LHC as a B-meson factory and is designed to study the forward decays of heavy bottom and charm hadrons.

The LHC operates in data-taking periods, called **runs**, separated by long shutdown (LS) periods in which the accelerators and detectors are maintained and upgraded. The LHC is currently in the third data-taking period, called Run 3. The first period, Run 1, operated between 2010 and 2013, with proton-proton collisions occurring at a collision energy of $\sqrt{s} = 7$ TeV in 2010-2011, and 8 TeV in 2012. Run 2 operated between 2015-2018, at $\sqrt{s} = 13$ TeV, while Run 3 data-taking began in 2022 and will continue till the end of 2025 at $\sqrt{s} = 13.6$ TeV. Following the end of Run 3, the LHC will undergo upgrades for the High Luminosity LHC (HL-LHC) runs, which are scheduled to begin operation in 2029. During the runs, the accelerators and detectors also periodically undergo shorter technical stops (TS) for maintenance and upgrades.

The LHC is designed to accelerate protons to energies of up to 7 TeV, or 2.76 TeV per nucleon in the case of Pb ions. This results in a **center-of-mass energy** $\sqrt{s}$ of the collision of two particles,

$$\sqrt{s} = 2 \times E_{beam}, \tag{3.1}$$

which results in a maximum collision energy of $\sqrt{s} = 14$ TeV in proton-proton collisions in the LHC.

3.1.1 The Pre-accelerator Complex

The LHC alone cannot accelerate particles over the full energy range, from $0 \rightarrow 7$ TeV; it requires them to have an energy of at least 450 GeV at the time of injection. The CERN complex uses a chain of pre-accelerators[65] to get particles up to that minimum energy. The full injector sequence is shown in Figure 3.1 shows the full injector sequence.

1. In the first stage, a linear accelerator, **LINAC 4**[1] accelerates H^- ions to an energy of 160 MeV before injecting them into the next stage.

2. The **Proton Synchrotron Booster** (PSB) is the first circular accelerator ring. It is made up of four superimposed synchrotron rings that receive a beam of H^- ions from LINAC4, strip them of the two electrons, and accelerate the resulting protons to an energy of 2 GeV, further injecting them into the next stage.

3. The **Proton Synchrotron** (PS) is a synchrotron ring with a circumference of 628 m. The PS accelerates the proton beams to an energy of 26 GeV and accumulates them to form a train of bunches separated by 25 ns. These are then injected into the next stage.

4. The **Super Proton Synchrotron** (SPS) is a synchrotron 6.9 km in circumference, and further increases the energy of the protons to 450 GeV over a period of a few minutes. The beams are then injected into the main LHC ring[2].

[1]**LINAC 2**, was used in Run 2. It stripped Hydrogen atoms of the electron and accelerated the resulting H^+ ions (which are just protons) up to an energy of 50 MeV. By using H^- instead, **LINAC 4** is able to achieve a much higher beam intensity and energy.

[2]Before the LHC began operations, the SPS was used as a proton-antiproton collider, and witnessed the discovery of the W and Z bosons.

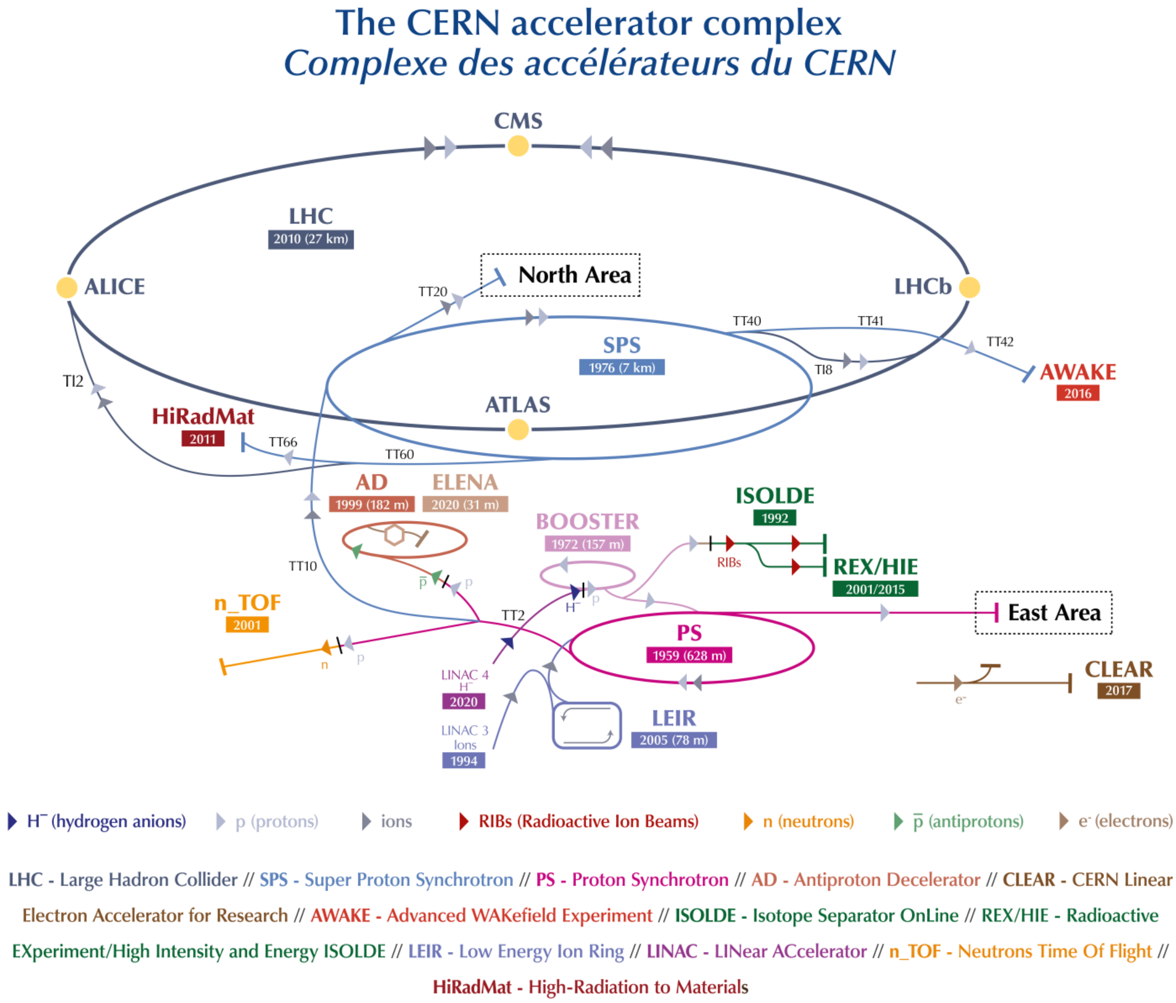

Figure 3.1: The CERN accelerator complex [69].

5. Finally, the **Large Hadron Collider** (LHC) further accelerates the proton bunches over about 20 minutes, after which they reach the peak energy of 6.8 TeV and are ready for colliding.

Rather than a continuous beam, the particles in the beam are organized into longitudinally spaced intervals called **buckets**. A bucket can contain a fill of particles called a particle **bunch**; however, a bucket can also be empty. The exact bunch configuration depends upon the injection cycles of the LHC pre-accelerators. It is optimized to control beam interaction rates and to remain within technical limitations, such as the heat-absorbing capacity of the cryogenic system.

The particle beams are kept in their circular orbits in the LHC by means of 1232 superconducting dipole magnets, which provide a powerful magnetic field of 8T to bend the particle trajectories, with an additional 392 quadrupole magnets used to keep the beams collimated. Such strong fields are only possible at extremely high currents (11500 Amperes) passing through the magnets. This is only possible when the magnets exist in a state of superconductivity, which is achieved at extremely low temperatures using the cryogenic system. Thus, the magnets are operated at temperatures below 2 K using superfluid Helium as a coolant. Arrays of magnets with higher multipole orders are used to focus the beams and counteract the electromagnetic repulsion within the beam, and prevent interactions between the beam particles and the confines of the beam pipe. Upon reaching the desired collision energy, dedicated arrays of focusing quadrupole magnets at the interaction points are used to squeeze the beams to a diameter of about 10 μm and maximize the chance of interaction between the beam particles, whereupon a fraction of all the particles in the bunch collide in each revolution. After the collision, the beams are again redirected into their respective beam pipes for another revolution. This cycle continues for about 10 hours, until the bunches are sufficiently depleted of protons, following which they are dumped.

It is important to have a figure of merit that reflects the rate of collisions. This quantity, used for colliders such as the LHC, is called the **instantaneous luminosity**.

3.1.2 Luminosity

The instantaneous luminosity is a measure of the number of collisions that can be produced in a detector per cm^2 and per second:

$$\mathcal{L} = \frac{1}{\sigma}\frac{dN}{dt} \tag{3.2}$$

where σ is the cross-section – a measure of the probability of a particular interaction occurring, measured in *barns* (b), a unit of area ($1\text{ b} = 10^{-28}\text{ m}^2$). The larger the luminosity, the larger the rate of collisions.

There are several possible ways that protons can interact when colliding: they could glance off of each other, or undergo hard scattering and produce new resonances that further decay into other particles. Each of these processes has its own cross-section that quantifies the probability of that interaction. Thus, for a rare process, the only way to increase the chances of seeing it is to increase the number of collisions, i.e., by increasing the luminosity, or the operation time.

The total size of the recorded dataset is characterized by the total **integrated luminosity**,

$$\mathcal{L}_{int} = \int \mathcal{L}(t)dt \tag{3.3}$$

which is simply the instantaneous luminosity integrated over the data collection period of the LHC. This, along with Eq. 3.2, gives us the total number of expected events for a given physical process with cross-section σ,

$$N = \int \frac{dN}{dt}dt = \sigma \int \mathcal{L}(t)dt = \sigma \cdot \mathcal{L}_{int} \tag{3.4}$$

The luminosity cannot be made arbitrarily high due to technical limitations. It is influenced by several parameters, as shown in Equation 3.5

$$\mathcal{L} = \frac{k_b N_p^2 f_{rev} \gamma}{4\pi \epsilon_n \beta^*} F, \tag{3.5}$$

with Table 3.1 summarizing the terms in Equation 3.5.

Term	Definition	Value
γ	Lorentz factor	6500
f_{rev}	Revolution Frequency	11,246 Hz
k_b	Number of proton bunches	2750
N_p	Number of proton in a bunch	1.4×10^{11}
ϵ_n	Normalized RMS transverse beam emittance	1.3 μm
β^*	Beta function at collision point	0.3 m
F	Reduction factor due to crossing angle	~ 1

Table 3.1: The parameters determining the instantaneous luminosity, as well as their typical values in Run 3, mostly reproduced from Refs. [70, 71].

The beam emittance ϵ_n describes the spread of the beam particles in position-momentum space. The smaller the beam emittance, the more squeezed the beam is and the higher the probability of a collision. The repulsion between the protons in the beam causes the beam to spread out, while the magnetic field squeezes it together. The amplitude function β^* is determined by the quadrupole magnet configuration and power, and refers to the distance over which the beam is squeezed by half to its narrowest at the collision point. As the beams cross in the collision region, the beam-beam interactions cause non-linear perturbations of the charge distributions in the beams, and result in a decrease in particle density in the bunch. The number of protons and the filling scheme is limited by the heat load and temperature margins inside the superconducting magnets and cryogenic systems. The beam dumping system must also be able to safely and quickly absorb the particle beams during normal running operations, as well as in case of an emergency. Further details are provided in [5, 65].

In 2022, the first year of the LHC Run 3, the accelerator provided collisions at center-of-mass energies of $\sqrt{s} = 13.6$ TeV. The typical filling scheme was 2750 bunches of 1.4×10^{11} protons each, spaced out by 25 ns, resulting in a total integrated luminosity of 42 fb^{-1} delivered to the CMS experiment, out of which 38.5 fb^{-1} was recorded. Figure 3.2 shows the integrated luminosity delivered to CMS in Run 2 and Run 3, from 2015 up to 2023.

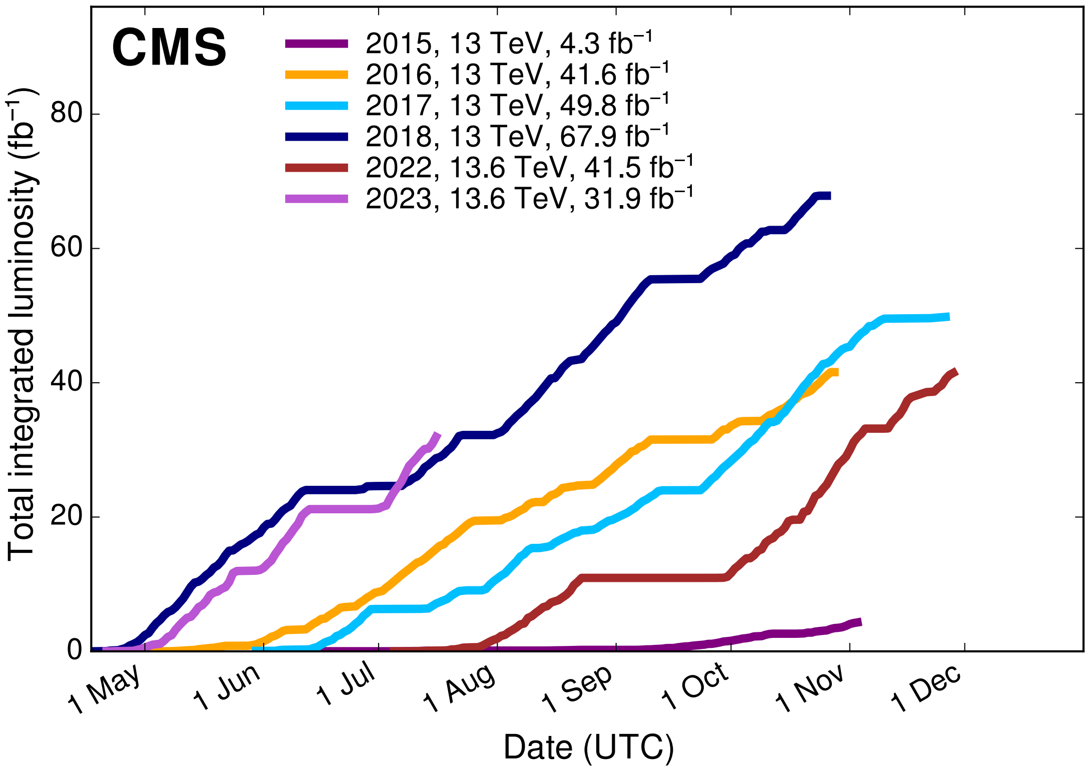

Figure 3.2: Cumulative integrated luminosity per year delivered to CMS during Run 2 and Run 3, during stable beams for proton-proton collisions at nominal center-of-mass energy, from 2015 to 2023, measured by the CMS collaboration [72].

3.2 The Worldwide LHC Computing Grid

The Worldwide LHC Computing Grid (WLCG) is a global collaboration of computer centers and storage systems used to store, distribute, and analyze the 15 petabytes of data generated by the LHC every year, as well as to perform simulations for the LHC and the various experiments operating at it. This is made possible by combining computer facilities at CERN, with the national or regional resources brought in by member institutes and laboratories across the world. The WLCG was instituted as part of the LHC design to handle the significant volume of data generated.

The WLCG is composed of three 'Tiers', each providing a specific set of services.

- **Tier 0** accounts for about 20% of the total computing resources. It stores the digital information directly from the subdetectors (called 'RAW' data), performs a first calibration pass of the reconstruction, and passes the output through high-speed optical-fiber links (10 GB/s) to the next tiers, Tier 1 and 2. When the LHC is not running, it also reprocesses the data.

- **Tier 1** consists of about 15 centers situated globally and is responsible for round-the-clock support for the grid, large-scale reprocessing, storing the corresponding output, and distributing the data to Tier 2. It also stores a copy of the raw data.

- **Tier 2** centers are located mainly in universities or scientific institutions in most regions of the world. They are interconnected with each other and with Tier 1 centers by the general-purpose research network. They are used by physicists to perform their analysis, as well as for simulation. The simulated data are sent back to Tier 1 and stored. Analysis programs are either directly run on Tier 2 infrastructure or by producing data formats on local clusters that are sometimes classified as **Tier 3**.

Users may access the grid with the proper credentials. This analysis was facilitated by the CERN computing cluster called 'LXPLUS', as well as the Tier 2 facility and computing cluster at the Vienna Bio Center called 'CLIP' [73].

3.3 The Compact Muon Solenoid

The Compact Muon Solenoid experiment (CMS) [62, 63] is a general-purpose detector designed to record proton-proton and heavy-ion collisions at the LHC at high instantaneous luminosities and energies up to $\sqrt{s} = 14$ TeV. It is located 100 m underground at the LHC Interaction Point 5 (IP5), near the French town of Cessy, between Lake Geneva and the Jura mountains.

The CMS detector is designed to trigger on [74, 75], identify and measure the properties of particles that traverse the detector and interact with the various detector modules, namely electrons [76], muons [77], photons [78], and (charged and neutral) hadrons [79]. The particles are generally stable and pass through multiple layers, allowing us to reconstruct their trajectories as well as properties like energy and momentum. However, thanks to the precision of these measurements, as well as accurate spatial resolution, we are able to also reconstruct unstable and intermediate particles indirectly based on information from detected particles, allowing us to probe a wide range of physical phenomena.

The CMS detector is cylindrical in shape, with the detector modules arranged in layers around the beam pipe. These modules are arranged in a central 'barrel' section, and two 'endcap' sections on either end of detector, providing an almost complete 4π coverage in solid angle from the proton-proton interaction point. The name of the CMS detector is inspired by its design. The detector is *compact* in its construction, packing an incredible number and variety of sub-detectors in a surprisingly small space: it has a length of 22 m (28.7 m also considering the forward calorimeters), a diameter of 15 m, and weighs about 14 thousand tons. The central feature of CMS is a superconducting *solenoid* 6 m in diameter that suffuses the inner volume of the detector with a powerful 3.8 Tesla magnetic field. This field was chosen to obtain precise measurements of the momenta of charged particles, especially for muons. The third distinguishing feature of the CMS detector is its *muon* system, which can function as an independent detector to accurately reconstruct muons. The CMS detector can combine measurements from the tracker modules with those from the muon system to obtain the momenta of muons with much higher precision [77].

In addition to momentum resolution, CMS also provides good energy resolution to discriminate between other particles and account for their energy. To achieve this, CMS consists of a silicon pixel and strip **tracker**, a lead tungstate crystal **electromagnetic calorimeter (ECAL)**, and a brass and plastic scintillator **hadron calorimeter (HCAL)** within the solenoid volume. Forward calorimeters extend the coverage provided by the barrel and endcap detectors. Muon systems comprising of gas-ionization chambers are arranged outside the solenoid.

Figure 3.3 shows a cutaway rendering of the CMS detector and its sub-detectors. The main components are described in the following sections.

3.3.1 The CMS Coordinate System

The origin of the CMS coordinate system (x, y, z) is defined as the center of the detector, at the collision point of the two LHC beams. The x-axis is defined to point towards the center of the LHC ring, and the y-axis to point upwards. The plane spanned by the x and y axes is called the **transverse plane**. The azimuthal angle, ϕ, in this plane is defined to be zero in the x-direction. The z-axis points along the beam, with the positive z-direction defined by the *right-hand rule*, and is in the direction of the magnetic field.

The polar angle θ is the angle between the particle three-momentum vector and the positive direction of the beam axis. In hadron collision processes, the longitudinal momenta of colliding partons generally cannot be precisely measured. Therefore, the rapidity y is preferred over the polar angle θ, as the difference in rapidity of two particles is Lorentz invariant under boosts along the longitudinal direction. Additionally, in most high-energy hadronic collisions, the distribution of final-state hadrons is nearly uniform in rapidity in the central regions.

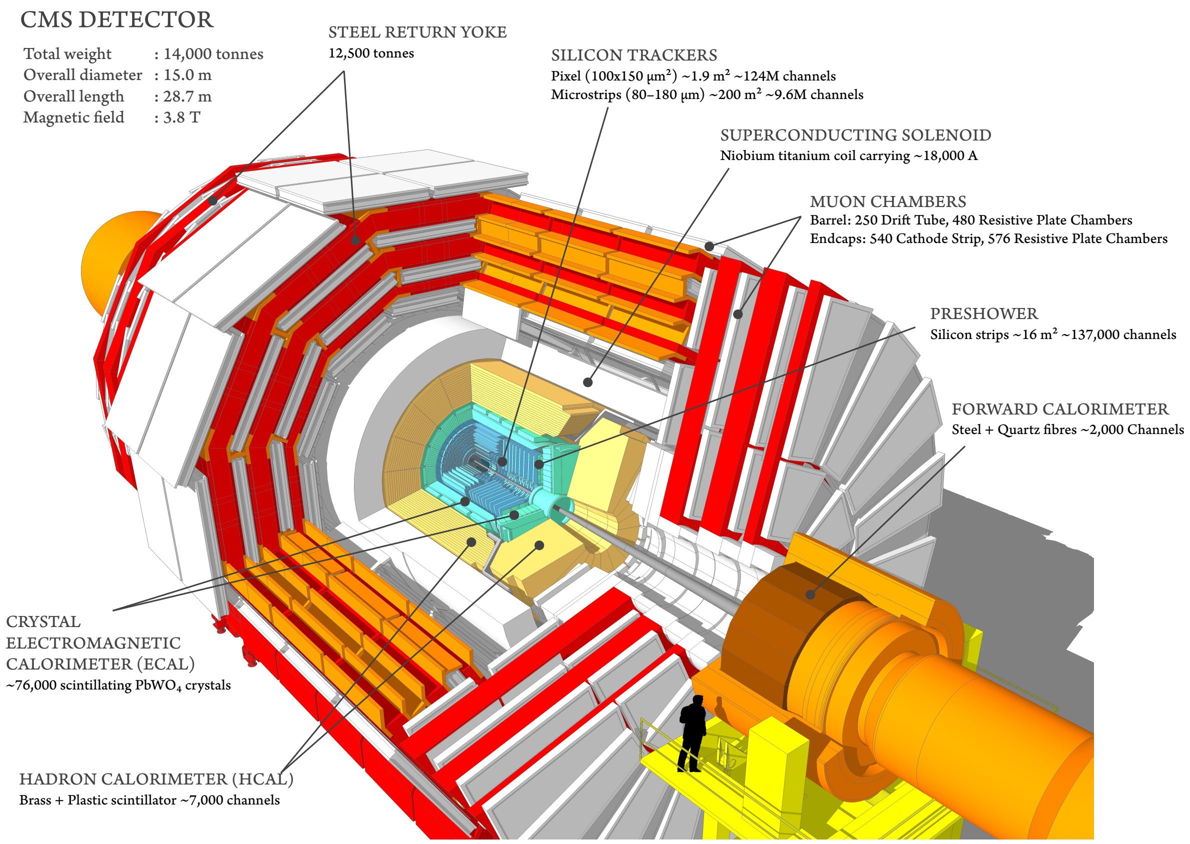

Figure 3.3: Illustration of the CMS detector and its sub-detectors. Image courtesy of the CMS collaboration.

The rapidity y is given by

$$y = \frac{1}{2}\ln\left(\frac{E+p_z}{E-p_z}\right) \tag{3.6}$$

In the limit that the particles are moving close to the speed of light, or if they are massless, the rapidity can be approximated by the **pseudorapidity** η, given by

$$\eta = -\ln\left[\tan\left(\frac{\theta}{2}\right)\right] = \frac{1}{2}\ln\left(\frac{p+p_z}{p-p_z}\right) = \tanh^{-1}\left(\frac{p_z}{p}\right) \tag{3.7}$$

Determining η depends only upon the trajectory of the particle, namely the polar angle θ, which is a much simpler observable than E and p_z. It is therefore used instead of the rapidity in high energy collision experiments, at high momenta and when the mass is unknown. In B-physics, typically, the rapidity y is used. Figure 3.4 shows the CMS coordinate system, as well as the polar angle (and corresponding pseudorapidity) angles. The spatial separation between two particles is denoted by ΔR, defined to be

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} \tag{3.8}$$

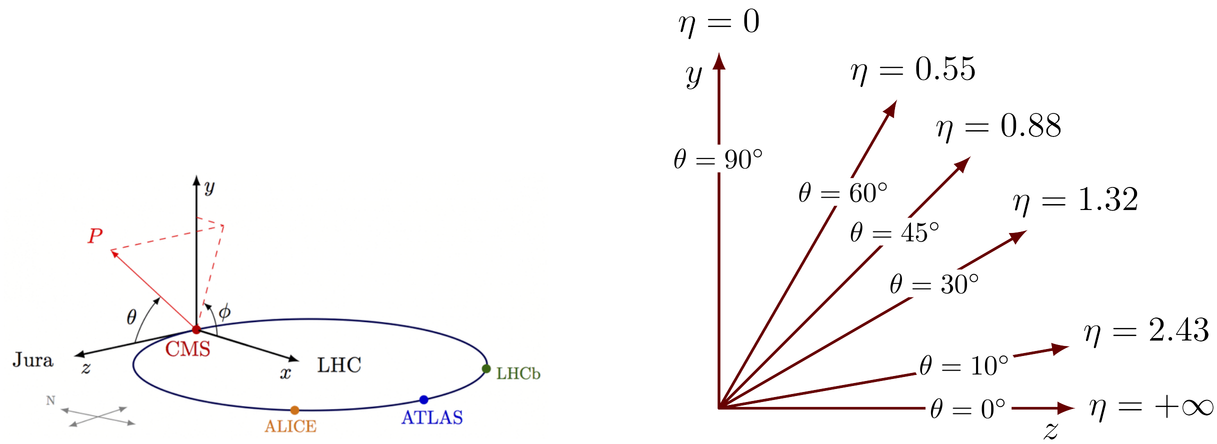

Figure 3.4: The CMS coordinate system (left), and the polar angles corresponding to different pseudorapidity η values (right).

3.3.2 The Solenoid Magnet

One of the central components of the CMS detector is the solenoid magnet [80, 81]. It is the largest superconducting solenoid magnet ever built, measuring 6.3 m in diameter and 12.5 m in length. It generates a powerful 3.8 T magnetic field that is essential to bend the trajectories of charged particles produced in proton-proton collisions, thereby measuring their momenta. To generate the field, a current of 18,500 A passes through the cylindrical coil of 4-layered windings of stabilized reinforced Niobium-Titanium (Nb-Ti) conducting wires embedded in an aluminium structure, that composes the solenoid, which operates at around 4.65 K (-268.5°C) to minimize resistance. The total energy stored in this solenoid during operation is roughly 2.6 GJ.

The magnetic flux is guided and returned through a steel return yoke surrounding the magnet that features five wheels, and two endcaps composed of three disks each. This return yoke has an outer diameter of 14 m and a barrel length of 13 m and weighs 12,500 tons, providing the structural support that is required for the CMS detector to withstand the massive forces generated by the magnetic field. The tracker and calorimeter (ECAL and HCAL) detectors fit inside the magnet coil, while the muon detectors are interleaved with the return yoke. Thus, the steel acts as a hadron absorber, filtering out the particles that traverse the inner tracking system and calorimeters and allowing only muons (and the invisible neutrinos) to pass through to the muon system.

3.3.3 The Tracker

The CMS Tracker [82, 83] is a sub-detector system dedicated to precisely measuring the trajectories of charged particles traversing the detector origination from the collisions at or near the LHC **beam spot**, which is defined as the pp interaction region. It is composed of fine granular silicon pixel detectors near the collision point, and silicon microstrip detectors afterwards. The tracker barrel is 5.8 m in length, with a radius of 1.2 m. The tracker has very little passive material to minimize particle interactions. Figure 3.5 shows the geometry of the tracker system and the various tracker regions, which are briefly described below.

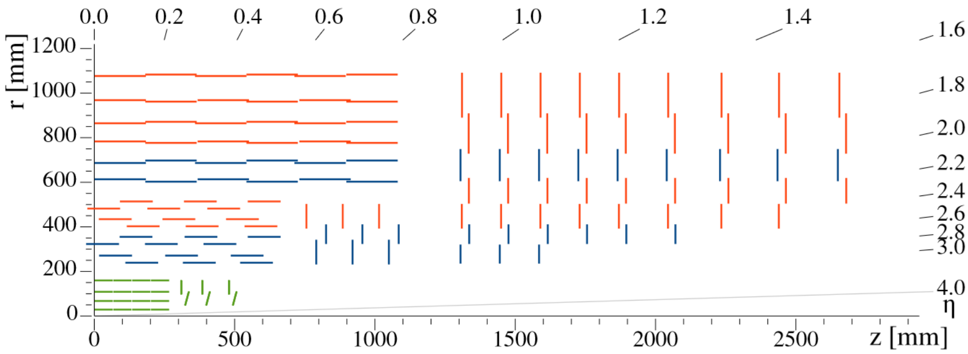

Figure 3.5: Schematic cross-section of one quarter of the CMS tracker [82]. The lines indicate detector modules.

The innermost part of the tracking system consists of silicon **pixel detectors**. It was upgraded [84] after the data-taking period in 2016, and consists of an arrangement of four concentric barrel layers (L1-L4) at radii r of 29, 68, 109, and 160 mm, and three endcap disks (D1-D3) on each end at distances of 291, 396 and 516 mm from the center of the detector. The innermost layer is located as close to the LHC beam pipe as possible to maximize the accuracy of measurement and extrapolation to the beam spot. The individual pixel sensors have a size of $100 \times 150\ \mu\text{m}^2$, corresponding to a total of 128 million readout channels, and provide a granularity fine enough to cope with the high particle flux densities expected during typical LHC running conditions.

In the outer parts of the tracking system, the particle flux densities are smaller, and silicon microstrip sensors provide adequate hit resolution in these regions. The microstrip sensors have a total of 9.3 million readout channels and 198 m^2 of active silicon area distributed over 15,148 modules [85]. These modules are arranged in four distinct zones.

- The four concentric layers in the **Tracker Inner Barrel (TIB)** are located at $r =$ 255, 339, 418.5, and 498 mm from the beam axis and span from $z = -700$ mm to +700 mm. The two inner TIB layers host back-to-back detector modules with an 80 μm separation between the modules (strip pitch), and the outer two layers feature single-sided with a strip pitch of 120 μm.

- The **Tracker Inner Disk (TID)** consists of three disks of detector modules placed between $z = \pm$ 800 mm and $z = \pm$ 900 mm on either side of the barrel. Each disk consists of three rings spanning radii of 200 mm $\leq r \leq$ 500 mm. The inner two rings hold back-to-back modules, while the other ring holds single-sided modules. Their strip pitches vary between 80 μm and 140 μm.

- The **Tracker Outer Barrel (TOB)** consists of six concentric cylindrical layers located between 555 mm $< r <$ 1160 mm and $|z| <$ 1180 mm.

- The **Tracker Endcaps (TEC)** extend between $\pm 1240 < z < \pm 2800$ mm on either side of the barrel, with a radial extent from 220 mm to 1135 mm. The two endcaps are called TEC+ and TEC-, depending on their z position in the CMS coordinate system. Each endcap consists of nine disk layers carrying up to 7 rings of silicon micro-strip detectors (320 μm thick on the inner four rings, and 500 μm thick on the outer three rings), with radial strips of 97 μm to 184 μm average pitch.

The total material budget of the tracker must be kept at a minimum to prevent particle energy losses as they traverse the tracker volume. Ultimately, this is determined by the silicon detectors themselves, as well as the readout electronics and the support and cooling structures. The tracker material budget is not isotropically distributed around the interaction point but varies as a function of the pseudorapidity, from 0.4 X_0 in the barrel region to a maximum of 1.8 X_0 in the overlap region, to 1.0 X_0 in the endcaps (where X_0 is a radiation length).

Charged particles traversing the tracking system sensors induce electron-hole pairs in the semiconductors, which are recorded as particle **hits** in the sensor. Multiple hits across the tracker layers serve as input for the reconstruction of the full particle trajectory - called a particle **track**. This reconstruction is done by means of dedicated pattern recognition and fitting algorithms. The trajectories of charged particles are bent by the magnetic field, allowing momenta, as well as their electric charges, to be reconstructed with high accuracy.

- The transverse momentum resolution, $\sigma(p_T)/p_T$, provided by the inner tracker is $< 1\%$ at 10 GeV, and about 2% at 100 GeV in the barrel region ($|\eta| < 1.2$). In the endcap region ($|\eta| \approx 2.5$), it is around 2% at 10 GeV and 10% at 100 GeV.

- The resolution of the transverse impact parameter d_{xy}[3], defined with respect to the beam line, ($\sigma(d_{xy})$), is between 20-200 μm at 10 GeV, and around 10 μm at 100 GeV. In the longitudinal direction, the position resolution ($\sigma(d_z)$) is 40-100 μm at 10 GeV and 100-1000 μm at 100 GeV.

- The tracking efficiency is around 99% for muons, and around 75-90% for charged hadrons ($\pi^{\pm}$).

These measurements were extensively studied using muons from Z boson decays and charged pions from $D^{\pm}$ meson decays in both data and simulation.

Most of the technical challenges for the design of the CMS tracker arise from the harsh collision environment due to the LHC run conditions. The detector and readout electronics have to operate reliably at frequencies of 40 MHz (corresponding to a 25-ns bunch-crossing interval). All tracker components have to endure a radiation-heavy environment over the full operational period.

[3] Also denoted by d_0

3.3.4 The Calorimeters

The CMS calorimeter systems are housed outside the tracking system, but still mostly within the confines of the solenoid. These sub-detector systems are designed to stop all particles except muons and neutrinos to measure their total energy. The accurate measurement of the particle energies requires that the components be highly granular. The calorimeter systems are divided into two sub-detectors, each designed to measure specific types of particles.

The Electromagnetic Calorimeter

The CMS Electromagnetic Calorimeter (ECAL) [86] is designed to measure the energies of lighter charged particles (mainly electrons and positrons), as well as photons. It is a hermetic, homogeneous, and transparent calorimeter made up of fine-grained lead tungstate ($PbWO_4$) crystals.

The ECAL is subdivided into a cylindrical **ECAL Barrel (EB)** and two disk-shaped **ECAL Endcap (EE)** modules. The EB has an inner radius $r_{in} = 124$ cm and an outer radius $r_{out} = 175$ cm, extending up to $|\eta| < 1.479$. The EE is located at $z = \pm 314$ cm and covers a pseudorapidity range of $1.479 < |\eta| < 3.0$. Precision measurements are possible up to $|\eta| < 2.6$, while the remaining coverage augments the measurement of energy flow in the forward direction. The structure of the ECAL is illustrated in Figure 3.6.

There are 75,848 crystals in total, with a crystal length of 230 mm in the barrel and 220 mm in the endcap. The total EB crystal volume is 8.14 m^3 with a weight of 67.4 tons, while the EE crystals are 2.90 m^3 with a weight of 24.0 tons. To compensate for the coarser granularity in the EE and improve the spatial resolution of the energy measurements, an additional ECAL Preshower Detector is installed in front of the endcap calorimeters at $1.65 < |\eta| < 2.6$, consisting of two silicon microstrip modules (4288 sensors with 137216 strips) and two lead absorber plates, with a thickness of 20 cm ($\approx 3\ X_0$). The lead radiator initiates the EM shower, and the silicon sensor chips, placed behind the lead block, are used for the readout. The PS detector improves the separation between photons and neutral pions (π^0), and can distinguish $e^\pm$ and photons against minimum ionizing particles (MIPs).

The optical signal produced by particles interacting with the crystal scintillators is converted into an electrical signal by photodetectors, which are attached to each crystal, and further digitized for signal processing purposes. Avalanche photodiodes (APDs) are used in the EB, while vacuum phototriodes (VPTs) are used in the EE.

The Hadron Calorimeter

The CMS Hadron Calorimeter (HCAL) [87] measures strongly interacting particles (such as protons, neutrons, pions and kaons) that create showers extending beyond the ECAL,

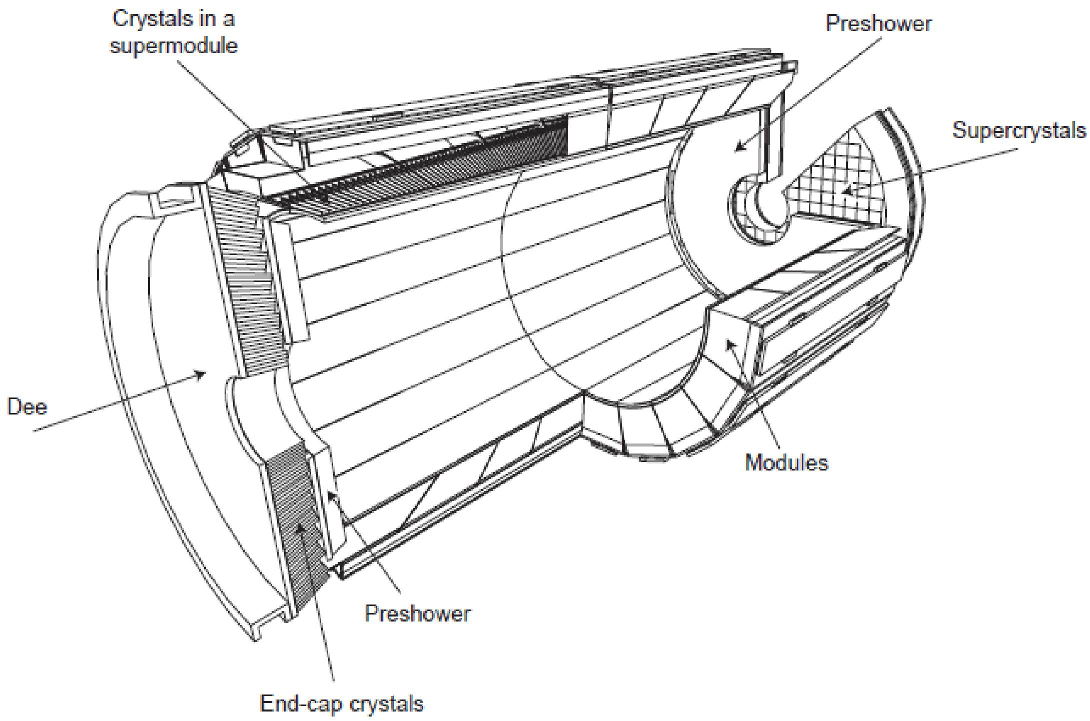

Figure 3.6: Schematic drawing of the CMS Electromagnetic Calorimeter (ECAL) [62].

by measuring the energy they deposit as they pass through. The HCAL practically covers the full solid angle and is hermetic..

The HCAL is a sampling detector composed of alternating layers of plastic scintillators and non-magnetic brass absorbers. Particles traversing the HCAL can produce hadronic showers in the brass layers, leading to measurable scintillation light in the interleaved scintillator layers. The HCAL is largely located between the ECAL (r_{in} = 1.77 m) and the solenoid magnet (r_{out} = 2.95 m). Figure 3.7 shows a schematic cross-section of the sub-detector and its various geometric regions.

- The **HCAL Barrel (HB)** constitutes the innermost part of the HCAL system, extending radially from $r \approx 1.8$ m to 3 m, with a length of about 9 m, extending in pseudorapidity to $|\eta| < 1.4$, with a granularity in $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$.
- The **HCAL Endcap (HE)** covers $1.3 < \eta < 3.0$, with a granularity of 0.087 in $\Delta\eta$, and about 5° to 10° in ϕ.
- The **HCAL Outer (HO)** is installed outside the solenoid and wraps around the HB. It provides coverage up to $|\eta| < 1.26$, complementing the HB and catching the 'tail' end of highly energetic showers that pass through both the HB and the magnet system.
- The **HCAL Forward (HF)** is placed outside the muon system endcaps on either side of the detector, at $z = \pm 11.2$ m. With a length of 165 cm, it provides coverage over $3.0 < |\eta| <5.0$. The HF needs to be radiation hard - each pp collision deposits

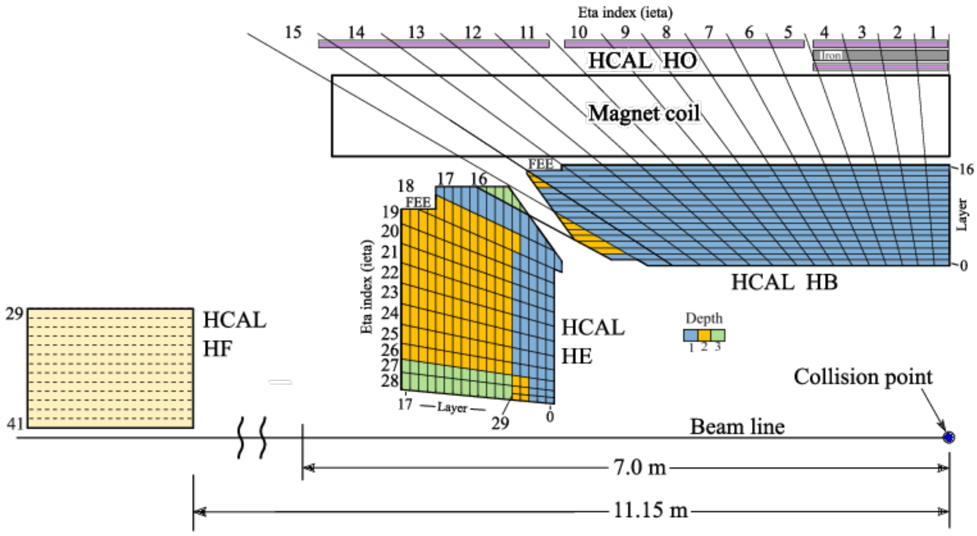

Figure 3.7: Schematic of one quarter of the CMS Hadron Calorimeter (HCAL), shown in the $r - z$ plane [62]. The overlaid acronyms indicate the HCAL sub-modules.

an energy of 760 GeV on average in the HF, compared to 100 GeV in the rest of the calorimeter system. Quartz fibers (polymer cladding with a fused silica core) are chosen as the active medium, and are expected to endure the radiation-heavy conditions for at least a decade.

The output is read out by wavelength-shifting fibres embedded in the scintillator tiles, and is sent to photodetectors via clear fibres. The thickness of the HCAL is represented in terms of the interaction length, λ_I, which is the mean distance travelled by hadronic particles before undergoing inelastic nuclear interactions. The total thickness of HCAL layers is around 7-11 λ_I up until the HF, and around 10-15 λ_I including the HO.

3.3.5 Muon system

Muons are 200 times heavier than electrons, minimizing the bremsstrahlung energy loss from interactions with the detector material. At the energies available at the LHC, they are MIPs and retain most of their energy when traversing the detector, making them favorable final-state particles for precision measurements due to the general discovery potential of decays featuring muons. The CMS detector was designed to identify and measure muons efficiently.

Figure 3.8 shows a schematic of the CMS muon system and its subdetectors, which make up the outermost layer of the detector.

The CMS muon system utilizes three main detector technologies – **Drift Tube (DT) chambers, Cathode Strip Chambers (CSCs), Resistive Plate Chambers (RPCs)**. Additionally, the CMS detector now also uses **Gas Electron Multipliers (GEMs)**, which are needed for the HL-LHC operations. The muon system is structured as a central cylindrical barrel section and two planar endcap regions, covering up to $|\eta| < 2.4$ without

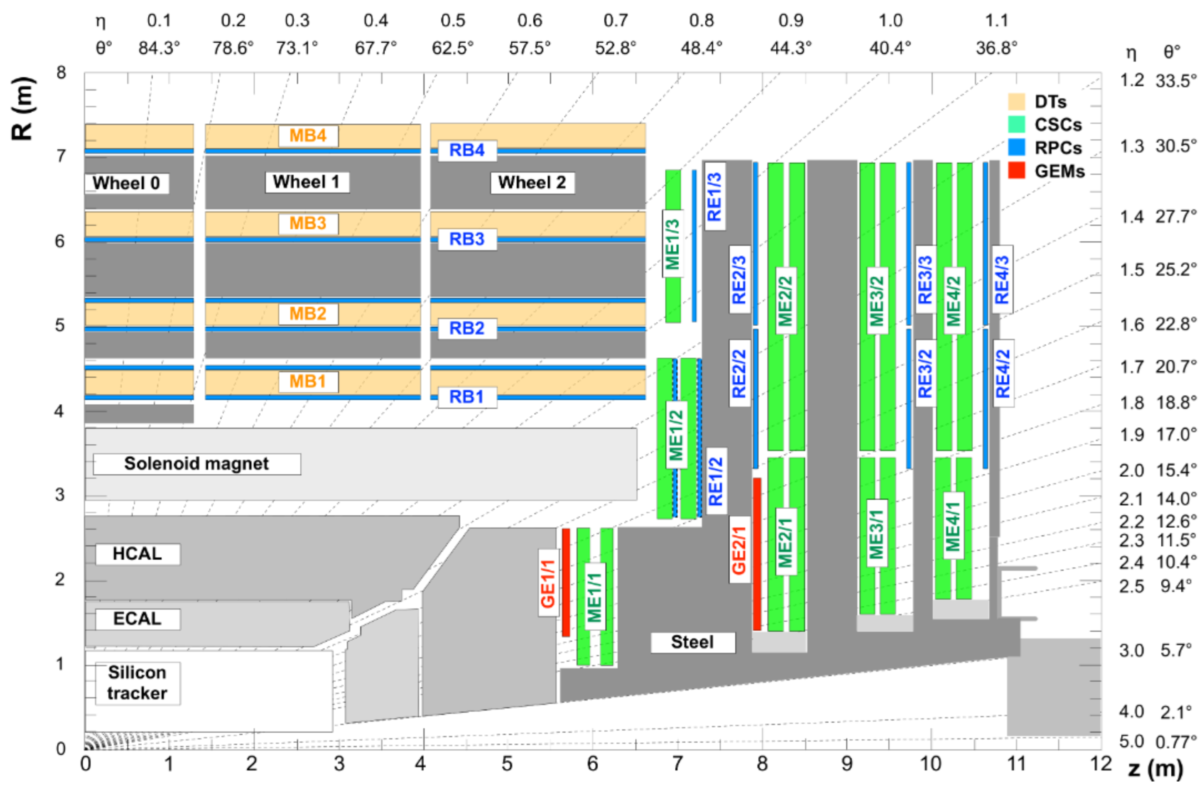

Figure 3.8: Schematic showing the CMS muon system [88, 89], and its sub-detectors.

any gaps, allowing for high reconstruction efficiency of prompt muons over the full η range. The muon system can also reconstruct muon tracks in real-time and trigger based on the measured p_T of the muons.

The different detector technologies used are briefly described below.

Drift Tubes

The **muon system barrel** covers $|\eta| < 1.2$ and is suffused with a uniform magnetic field of 3.8 T. The flux in this region is relatively low due to the small neutron-induced background. Thus, it is feasible to use drift chambers with rectangular drift cells to record muon hits and reconstruct tracks, as no other charged particles, other than the few muons, reach these layers.

Four muon barrel layers, called **stations**, are interspersed among the flux return yoke layers, and house 60 **Drift Tube (DT)** chambers in each of the three inner stations, while 70 DT chambers are housed in the outermost station.

A drift chamber consists of three detector units called **super-layers (SLs)**. Each SL consists of four layers of rectangular drift cells. In the outer two SLs, the charge-collecting wires are arranged parallel to the beamline and record hits in the magnetic bending plane (r, ϕ). The inner two SL wires are perpendicular to the beamline and provide measurements of the muon track z position along the beamline. This z measuring layer is missing in the fourth drift chamber, and thus, it only has two SLs. Each chamber has its drift cells offset by a half-cell width relative to neighboring cells to remove any dead spots. Hit information from multiple stations is combined (when possible) to reconstruct muon tracks with improved precision and minimize backgrounds. The cross-section of a drift tube in the DT chambers is 13 mm $\times$ 42 mm, optimized for mechanical stability, space constraints, and redundancy requirements.

Muons are often accompanied by electromagnetic cascades. To cope with such cascades, as well as uncorrelated background hits from neutrons or photons, the DT chambers were designed with several layers of separated drift cells per station, allowing for good tracking efficiency and measurement redundancies. The DT chambers are also capable of providing a muon time measurement with a resolution of a few nanoseconds, which is used for triggering and beam crossing time identification.

Cathode Strip Chambers

The **muon system endcaps** are situated on either side of the muon barrel, providing coverage for $0.9 < |\eta| < 2.4$. Each endcap contains four stations, positioned perpendicular to the beamline and embedded between the flux return plates. These layers experience higher rates and backgrounds than the barrel and are suffused by a large and non-uniform magnetic field. Thus, the detector technology is required to have a fast response time, fine segmentation, radiation resistance, and the ability to operate in large inhomogeneous magnetic fields. This is achieved by **Cathode Strip Chambers (CSCs)**.

CSCs are multiwire proportional chambers featuring six anode wire planes interleaved among seven cathode panels. An azimuthal alignment of wires allows for accurate measurement of the r coordinates of the hits. The hit ϕ coordinates are obtained by interpolating the charges induced in the strips, which run lengthwise at constant $\Delta\phi$ widths.

Resistive Plate Chambers

An additional complementary trigger system is installed in the muon barrel and endcap regions, consisting of **Resistive Plate Chambers (RPCs)**. These RPCs have fast response times, providing excellent bunch-crossing association and timing resolution, even in the high collision and background rate conditions expected during high luminosity LHC runs. Additionally, the RPCs also assist in resolving track reconstruction ambiguities when there are multiple hits in the DTs or CSCs. There are six RPC layers in the barrel, and four in the endcaps.

In total, there are 1400 muon chambers (250 DTs, 540 CSCs, and 610 RPCs). Altogether, the muon system contains nearly 1 million readout channels.

Between $0.9 < |\eta| < 1.2$, the muon system layers are aligned such that a muon from the interaction point traverses both the barrel and the endcap detectors. To optimize track reconstruction algorithms, the full η coverage is divided into three regions:

- The **barrel region**, with $|\eta| < 0.9$.

- The **overlap region**, with $0.9 < |\eta| < 1.2$.
- The **endcap region**, with $|\eta| > 1.2$.

3.3.6 Trigger and Data Aquisition

During peak luminosity runs, the LHC can provide over a billion pp collisions per second in the CMS detector. Each event corresponds to a bunch crossing, with multiple simultaneous collisions, called pileup (PU). For an event size of $O(1$ MB), this would result in a total output rate of $O(15$ PB/s). This much data is simply beyond the ability of the CMS **data acquisition (DAQ)** system to process and store, which has a maximum rate of about 10 GB/s. Most of the beam particle collisions are either not inelastic pp scattering collisions or do not have a high momentum transfer (called 'hard' collisions). Therefore, the number of events expected to contain 'interesting' physics (at least, in the context of the CMS physics programme) is much smaller.

The CMS trigger system [75] is an event filter that selects only relevant events at a rate that is manageable for the DAQ pipeline. It is structured in two sequential tiers – a hardware and a software trigger, respectively – that reduce the rate sequentially to manageable levels by rejecting events that fail the trigger requirements. At each level, a **trigger menu**, consisting of a set of individual trigger selections, is implemented in a logical 'OR'. These tiers are described below :

- The **Level-1 (L1) Trigger** is a hardware level trigger, an array of field-programmable gate arrays (FPGAs) that can implement up to 512 different trigger algorithms, called **L1 seeds**. Data recorded by the detector is buffered for about 3.2 μs at each sensitive element before it gets overwritten by fresh incoming data. Most L1 seeds are designed to process the data and reach a true/false decision in about 1 μs. However, the strict time constraints do not allow for arbitrarily complex trigger algorithms. Instead, the L1 trigger uses information from the muon system and the calorimeters only (with coarse granularity) to construct **trigger primitives (TPs)** for common particles, such as muons, electrons or photons, as well as other detector objects such as jets or local transverse energy sums. The L1 trigger reduces the event rate from 40 MHz (LHC rate) to about 100 kHz. The L1 seeds serve as pre-filters for algorithms in the next tier of the trigger. Figure 3.9 shows the architecture of the CMS Level 1 trigger system.
- The **High Level Trigger (HLT)** is a software-level trigger that implements trigger algorithms (called **paths**) executed on a computing cluster at the experiment site. Paths are sequences of algorithmic steps designed to reconstruct objects and make selections based on specific physics requirements. The HLT receives only those events that pass the L1 trigger, and due to the greatly reduced data input, it is possible to implement sophisticated track reconstruction algorithms, and the use of full detector

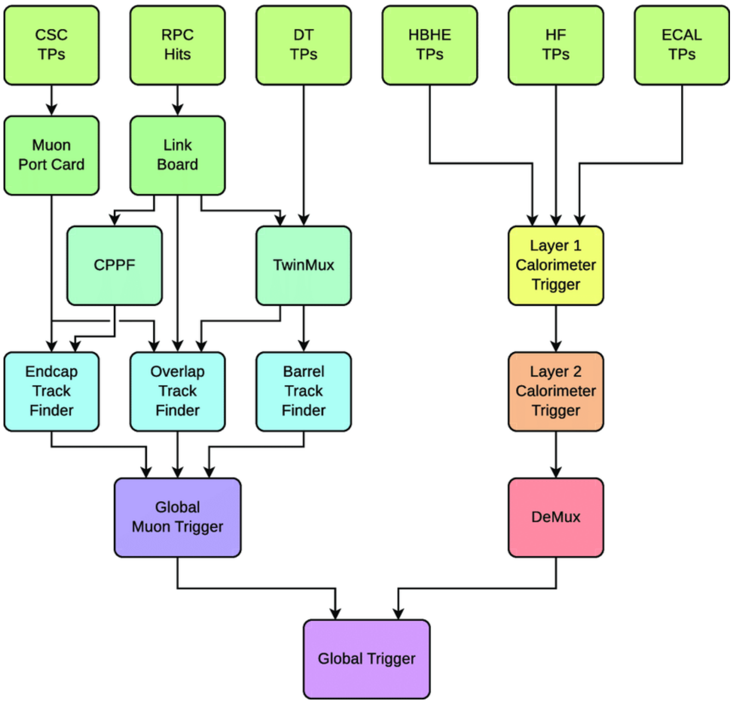

Figure 3.9: A diagram of the CMS Level-1 trigger system after its upgrade for Run 2.

readout data at the HLT, making for a better and more accurate performance that is comparable to the full reconstruction (also called **offline** reconstruction) performed at a later stage. The HLT further reduces the data-taking rate to a few kHz and about 10 GB/s – a manageable level for transfer to permanent storage and offline reconstruction.

At either stage, an event is accepted if it passes at least one of the seeds/paths, otherwise it is discarded.

Several factors determine the total rate of a given L1/HLT menu, such as the instantaneous luminosity and pileup. These conditions vary within a given data-taking run or LHC fill as collisions (along with beam loss effects) deplete the particles in a bunch, and as LHC parameters are adjusted. Thus, to keep the trigger rate roughly constant over time, the L1 physics triggers are usually available in several versions ('main' and 'backup' seeds), each tailored to specific intervals of instantaneous luminosity. The seeds and paths can also be 'prescaled' with certain numerical factors depending on the instantaneous luminosity – a prescale of 2 means that only one out of every two triggered events is stored. Standard multi-purpose CMS physics triggers are usually unprescaled and accept every triggered event.

The triggers constitute an important part of this thesis. New triggers were developed at both L1 and HLT, to benefit from new and improved reconstruction algorithms developed during the LS2 as well as from the increased rate budget due to the upgraded processing and storage capabilities of CMS. These triggers (described in Section 5.2) were designed keeping in mind lessons learned from the previous analysis [22], and resulted in a significant gain in sensitivity.

3.3.7 Overview of the Reconstruction of Physics Objects in CMS

All CMS subdetectors collect hit information and energy deposits. Different particles show different signatures in the detectors depending on their charges and the nature of their interactions, which can be used to identify them. Figure 3.10 shows these interactions and signatures in a schematic diagram.

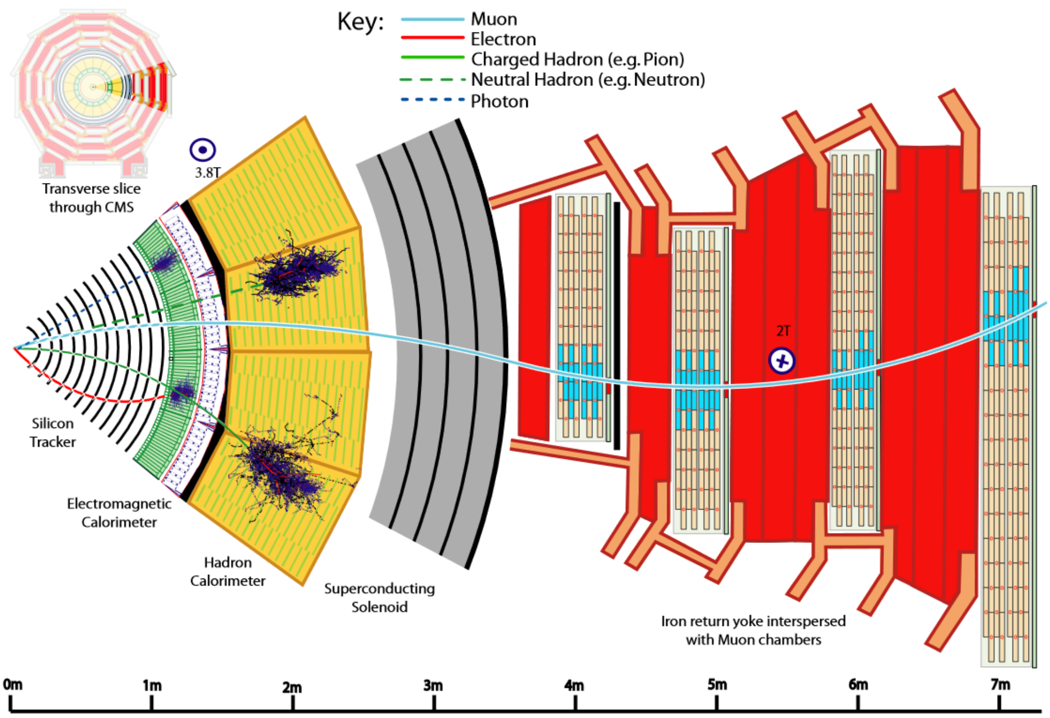

Figure 3.10: Interactions of different particle types with CMS sub-detectors. The hit/shower patterns across the sub-detectors are combined in the Particle Flow (PF) algorithm [90] to form hypotheses of final state particles.

Particles produced in the beam collisions initially pass through the tracker. Charged particles produce hits in the tracker layers, which are used to reconstruct their trajectories ('tracks') and their points of origin ('vertices'). Neutral particles, however, traverse the tracker without producing any hits.

The Electromagnetic Calorimeter (ECAL) is the next sub-detector to be encountered by the particles. Electrons and photons produce electromagnetic showers in the ECAL, which are reconstructed electronically as clusters of energy deposits. These particles are fully absorbed in the ECAL volume, allowing for a measurement of their energy, including an estimate of the direction of the particles' propagation.

Hadrons, both charged and neutral, traverse the tracker as well as the ECAL. While they might deposit some of their energy in these layers, they are generally too heavy for significant bremsstrahlung and pair-production-induced energy losses in the material, and most of their energy is deposited in the Hadronic Calorimeter (HCAL). The HCAL stops all particles (except muons and neutrinos) and measures their energies and directions through the energy clusters of the induced hadronic showers. Depending on the size of the showers and whether there are corresponding tracker hits and ECAL depositions aligned with the HCAL showers, it is possible to identify the particles.

Muons interact minimally with the calorimeters and only leave hits in the tracker system due to their small interaction cross-section. As a result, they are among the only particles to pass through to the muon system, and being charged, they leave hits therein. This makes the identification and measurement of muons unambiguous and highly accurate, making muons an excellent choice of object to trigger upon, and a useful handle in analyses.

Neutrinos have an extremely small interaction cross-section with matter. Therefore, they traverse the entire detector without interacting with any of the layers, escaping detection entirely, resulting in an imbalance in the total measured energy in the transverse plane. This imbalance is called missing transverse energy, E_T^{miss}, and is a useful handle in analyses.

These particle hits and energy deposits are used to produce 'physics objects' – hypotheses of stable particles along with their kinematic properties, such as estimated trajectories, energies, momenta, and types, using reconstruction algorithms. This reconstruction of physics objects and the global event is known as "particle flow' (PF) reconstruction [91]. While PF is not discussed in this thesis, muons are of particular importance to this work. The reconstruction of muon tracks, both at the trigger and in the offline analysis, is done using the Kalman filter technique [92] – a recursive algorithm that iteratively builds tracks, performing pattern recognition layer by layer, while simultaneously updating the trajectory parameters.

The following sections briefly describe the reconstruction of muons at various stages of data taking.

3.3.8 Online Muon Reconstruction

The identification and measurement of muon candidate objects at the CMS trigger stages is called the **'online'** muon reconstruction. This section briefly describes these objects and reconstruction algorithms. Further details may be found in Refs. [74, 75, 88, 93, 94].

L1 muon objects

At the Level 1 Trigger, muon candidates are reconstructed using trigger primitives – short track segments – from the four overlapping subdetectors in the muon system only: DT, CSC, RPC, and GEM. The GEM detector was added as part of the Phase 1 upgrade and was used for the first time in Run 3. The TPs carry rudimentary position coordinates (θ and ϕ), as well as information about the muon direction and timing (with respect to a collision bunch crossing). The position coordinates – mainly ϕ – differ between consecutive stations due to the deflection caused by the magnetic field. From this, it is possible to measure the angular deflection $\Delta\phi$, and thus obtain a measurement of the track p_T.

Different track-finding algorithms are used in three distinct pseudorapidity regions of the detector :

- The **Barrel Muon Track Finder (BMTF)** uses hits or trigger primitives (TPs) (short track segments) in the barrel muon stations (DTs and RPCs) to provide muon candidates in $0 < |\eta| < 0.83$ region. An approximate Kalman filter [95] based algorithm (called **'KBMTF'**) reconstructs muon tracks from detector hits beginning from the outermost muon stations and propagating inwards while updating the track parameters. It was updated during LS2 leading up to Run 3, and retains high efficiency for prompt muon tracks up to $d_{xy} \approx 50$ cm. A version of KBMTF developed for displaced muons called displaced KBMTF was introduced in Run 3, increasing the efficiency for muon tracks with displacements above 50 cm.

- The **Overlap Muon Track Finder (OMTF)** produces muon candidates in the $0.83 < |\eta| < 1.24$ region, using TPs from the DTs and RPCs in the barrel, and CSCs and RPCs in the endcap. The algorithm uses a Bayes classifier based on precomputed patterns generated from simulated events to associate hits in each station with the reference hit in the pattern. In Run 3, the OMTF finds the best matching prompt and displaced patterns in parallel for each muon track, with efficiency of around 80% up to displacements of 200 cm.

- The **Endcap Muon Track Finder (EMTF)** produces muon candidates in the $1.24 < |\eta| < 2.4$ region, using TPs from CSCs, RPCs and GEMs in both endcaps. GEM inputs are used in the EMTF algorithm for the first time in Run 3. The algorithm uses one TP per endcap muon station, with CSCs having priority over RPCs, to build muon tracks based on a set of predefined patterns. The differences in ϕ and θ angles between hits in different stations are used in a boosted decision tree (BDT) to estimate the p_T of the muon track.

All three muon track finders provide up to 36 muons each per bunch crossing to the **Global Muon Trigger**, which resolves duplicates and provides a maximum of 8 muon tracks per bunch crossing to the **Global Trigger**.

Due to the strict time constraints of the L1 trigger, the computational complexity of the L1 reconstruction algorithms is limited, and is thus generally less accurate than the offline algorithms, which are used at a later stage and are able to reproduce track parameters as closely as possible to their true values. Still, the performance of L1 muon reconstruction for prompt muons is $> 90\%$ up to a transverse distance d_{xy} of 40 cm (after which it falls rapidly), as shown in Figure 3.12.

The track finding algorithms up to Run 2 used the beam spot as an extra point in the track fitting (called the **'beam spot constraint'**), giving a longer lever-arm and much better precision in p_T measurements for prompt muons. However, as Figure 3.11 illustrates, for displaced muons, the beam spot constraint would cause the fitted track to have an artificially larger curvature, and thus assign a smaller p_T to the track, causing the muons to fail the trigger p_T thresholds and lead to an inefficiency in the trigger. In 2022, a new Kalman filter based track finding algorithm [95] was introduced in the barrel region of the muon system, specifically to improve efficiency for displaced muons. The displaced KBMTF can recover this efficiency by omitting the beam spot constraint, resulting in a better track fit, and therefore, a better p_T measurement for displaced muons. However, the trade-off is a slightly worse p_T resolution for prompt muons. The L1 muon p_T reconstruction efficiency for displaced muons is above 90% up to a d_{xy} of 60 cm, and falls slowly beyond that to about 80% at d_{xy} = 80 cm, as shown in Figure 3.12. The performance in Run 3 of the prompt and displaced L1 muon triggers was measured in a cosmic ray muon dataset recorded in 2018 [85].

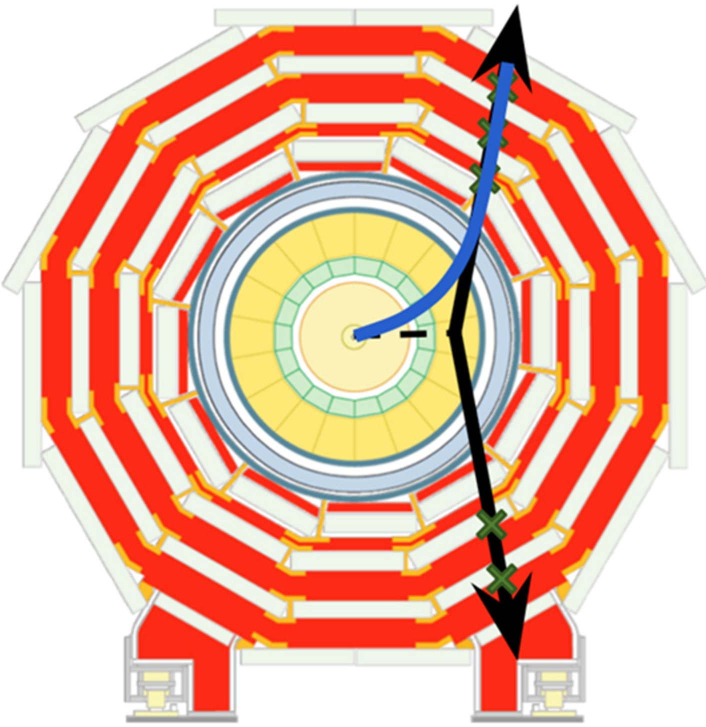

Figure 3.11: Illustration of the beam spot constraint in muon track-finding, on a transverse cross-section of the detector. The black arrows represent the true path of the muon leaving hits in the muon system. The blue curve is the fitted track assuming that the muon originates from the beam spot, resulting in a trajectory that has a stronger curvature and, consequently, a lower p_T measurement.

L2 muon objects

At the HLT, more sophisticated reconstruction algorithms are employed, due to the more relaxed time constraint, to more accurately identify and measure muons. Consequently, this further reduces the data-taking rate. At this stage, muons are reconstructed in two

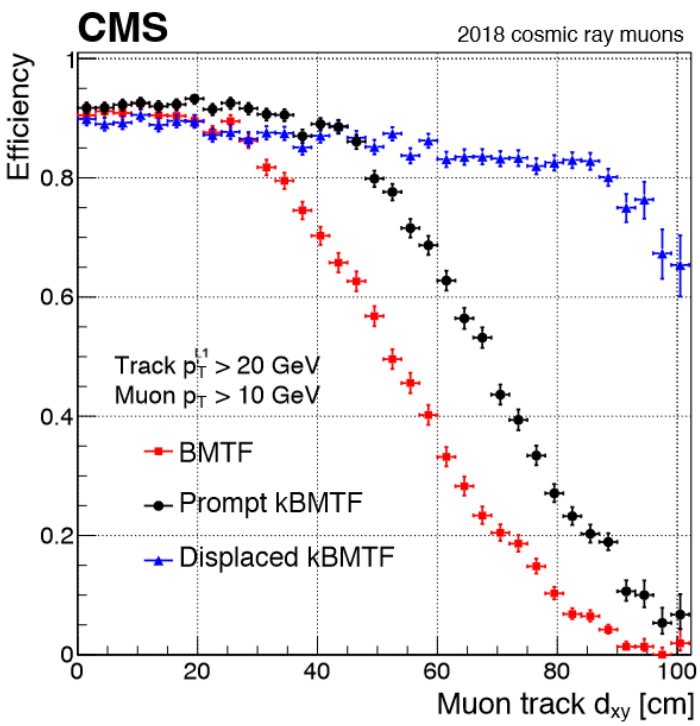

Figure 3.12: Displaced (blue) and prompt (black) KBMTF trigger efficiencies compared to the prompt BMTF (red) trigger efficiency with respect to the muon track d_{xy}, obtained using a sample of cosmic ray muons from 2018 data. The efficiencies are measured using muon candidates with $p_T > 10$ GeV. Figure sourced from Ref [85].

steps: first using only hits in the muon system, called **L2 reconstruction**, followed by a combination with hits in the inner tracking system, called **L3 reconstruction**.

L2 reconstruction is equivalent to the offline **'standalone'** muon reconstruction. At the beginning of the L2 reconstruction, initial track states called the **'L2 seeds'**, are generated from segment patterns in and across the DT and CSC stations. These 'states' refer to a 5-dimensional parametrization of a trajectory on a specific muon station surface, comprising of: the ratio of muon charge to its p_T, the transverse and longitudinal angles with respect to the beam direction, and the transverse and longitudinal impact parameters with respect to the beam spot. While this initial seeding step assumes that the muon originates in the vicinity of the beam spot, an alternate version of the L2 muon reconstruction designed for cosmic ray muons replaces the initial seeding step with one free of this assumption – this is called the **'cosmic-seeded'** L2 muon reconstruction, and has a higher efficiency for displaced muons not pointing to the beam spot, but is slightly less efficiency for prompt muons (especially in events with two close-by muons).

For these L2 seeds, geometric matches ($\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} < 0.3$) to a L1 muon are required. They are then assigned a direction and a p_T^{L2} value obtained from a $\Delta\phi$ parametrization based on the bending angle of the segment with respect to a neighboring segment. Standalone muon tracks are then built, starting from the L2 seed state, using the Kalman filter method [92] of iteratively constructing tracks.

In one iteration, track reconstruction is performed ‘inside-out’, where the tracks are extrapolated outwards to the muon system and correlated with measurements in the muon chambers. In another iteration, the track is built in the opposite direction - ‘outside-in’. A combination of these two iterations allows the reconstruction to remove potential biases originating from the initial seed.

In the final step, a beam spot constraint may be applied to the muon tracks (similar to the L1 track finders), improving track parameter resolution from prompt muons and yielding an overall L2 reconstruction efficiency of $> 99.5\%$ with respect to the L1 muons. However, this introduces trigger inefficiencies for displaced muons when $p_T^{(\mathrm{L2})}$ thresholds are imposed. The L2 muon reconstruction is thus available in two configurations, with and without the beam spot constraint. These correspond to the offline **‘standalone muon reconstruction’** and the **‘displaced standalone muon reconstruction’** respectively, and are discussed in more detail in Section 3.3.9. The latter version is suitable for the reconstruction of muons that are produced away from the pp collision point, such as cosmic-ray muons (in dedicated data-taking runs), or displaced muons resulting from LLP decays in hypothesized BSM processes.

L3 muon objects

Lastly, the L3 muon reconstruction algorithm at the HLT is not restricted to measurements from the muon system alone, but utilizes all available information, including from the inner tracking detector.

In Run 3, the HLT paths used in our search build L3 muons using the following two distinct ways.

The ‘cascade algorithm’ : L3 muon seeded by an L2 muon

The ‘cascade’ algorithm reconstructs muon trajectories in the inner tracker based on L2 muons, using three sequential algorithm steps, such that the later (more precise and computationally demanding) steps are only performed if the previous steps were not successful in reconstructing a track for a given L2 muon – hence, the name ‘cascade’. Muon tracks reconstructed in the earlier steps with $d_{xy} > 0.2$ cm are rejected to allow the later steps to reconstruct them again properly and ensure that poorly reconstructed tracks are not used in the final fit. These are :

- The ‘outside-in state’ algorithm propagates the track state of the L2 muon to the outermost tracker layer. Using a "stepping helix" method, which accounts for the magnetic field, energy losses, and scattering, the algorithm identifies detector modules compatible with this track state. If no compatible module is found, the algorithm continues inward layer by layer until one is located or all layers are checked. The pattern recognition uses a Kalman filter, propagating layer-by-layer towards the interaction point and adding compatible hits to the trajectory, followed by the final track fit.

- The ‘outside-in hit’ algorithm tracks the L2 muon to the outermost layer of the tracker. Instead of using the propagated state as a track seed, it first looks for compatible hits in that layer. Up to five of these hits are used to update the track state and create one seed per hit. The algorithm then moves inward layer by layer, looking for compatible hits until it finds one or exhausts all layers. If the seeds are within 25 cm of the detector center along the beam axis, they are used to start the same pattern recognition and track fitting steps as the outside-in state algorithm.
- The ‘inside-out hit’ algorithm uses the L2 muon information to define a region of interest (ROI) on the innermost pixel layer. The size of this region is determined by either a parametrization of the η and ϕ of the L2 muon, or the uncertainty in the L2 parameters at the pixel detector, with a minimum size of 0.1×0.1 in η and ϕ. Within these regions, compatible triplets or doublets of pixel hits are formed, depending on whether three or two pixel layers have compatible hits respectively. These hit multiplets serve as seeds to start the pattern recognition from the pixel detector outward.

The tracker tracks created from these three cascading algorithms are then matched to the L2 muons by propagating both the tracker tracks and L2 muons onto a common surface and comparing using η, ϕ, and χ^2 compatibility. A combined fit is then performed with the Kalman filter using information from both the selected tracker track and the matched L2 muon, and the fit with the best χ^2 per degree of freedom is used as the final L3 muon.

The ‘iterative’ algorithm : L3 muon seeded by both an L2 and L1 muon

The installation of the new pixel detector in 2017 required updated reconstruction techniques to take advantage of the new pixel layer. The iterative algorithm consists of three steps: an ‘outside-in’ step seeded by an L2 muon, an ‘inside-out’ step seeded by an L2 muon, and an ‘inside-out’ step seeded by an L1 muon.

- The ‘outside-in’ reconstruction step uses an L2 muon to search for track reconstruction seeds in the inner tracker. Using parameters from the L2 muon, the trajectory is propagated to the outermost tracker layers, similar to what is done in the outside-in steps of the cascade algorithm. The iterative algorithm creates two trajectory states on the outer tracker surface for seeding, one with the beam spot constraint and one without. Additionally, up to five seeds are generated by finding a silicon strip detector hit that is compatible with the extrapolated state.

- The ‘inside-out’ steps make use of iterative tracking techniques, generating ROIs using either L2 or L1 muons, depending on the step. The first iteration is seeded using pixel quadruplets, the second iteration is seeded by pixel triplets, and one final iteration is seeded by pixel doublets to recover efficiency in situations where fewer than three hits might be available for seed building.

Each reconstruction step produces a collection of tracker tracks and combines them, removing duplicates. In the L2 muon-seeded steps, the tracker tracks are matched to the L2 muon, and the L3 muon is built from a combined fit of the tracker and the muon track. In the L1 muon seeded step, the tracker tracks are matched to hits in the muon system to create tracker muons. After removing duplicate tracks that are already contained in the L3 muons collection after the previous steps, all muons are merged into a single collection. This matching and refitting procedure is identical to that used in the offline reconstruction.

3.3.9 Offline muon reconstruction

The offline muon reconstruction is conceptually the same as that used in the HLT. In the standard CMS muon reconstruction for pp collisions [77], muon tracks are first built independently in the inner tracker and the muon system. The final reconstructed muon objects are of three major types, depending on what information they use :

- **Standalone (SA)** muons are reconstructed using information from the muon system only. This is done by gathering all available information from the DT, RPC, and CSC detectors along the muon trajectory using a Kalman filter technique. Groups of DT or CSC segments are the seeds for the SA muon reconstruction.

  Several variants of standalone muons exist in the CMS reconstruction software. The original SA muon reconstruction, developed prior to the start of the LHC operations, is available with and without the beam spot constraint in the track fit. The latter version is more suitable for the reconstruction of displaced muons and was used as a starting point for two subsequent modifications.

  The **‘refitted standalone’ (RSA)** muon reconstruction was used in the Run 1 search for displaced dimuons [58]. It performs an additional track refit in an attempt to mitigate biases caused by the beam spot requirements in the standalone muon seed generator. The RSA muons provided better spatial and momentum for displaced muons than SA muons [96].

  The **‘displaced standalone’ (DSA)** muon reconstruction replaces the seed generator of the SA reconstruction with the one used in the dedicated reconstruction of cosmic-ray muons and improves spatial and momentum resolution for displaced muons. This is a more recent attempt to get rid of the beam spot-related biases in the SA muon reconstruction and is used in both the Run 2 search [22] as well as this search.

- **Tracker muons** are reconstructed starting from tracks in the tracker, which are propagated outwards to the muon system if they have $p_T > 0.5$ GeV and $p > 2.5$ GeV. Loose matching to DT and CSC segments is then attempted for these extrapolated tracks. If a match is found, the tracker track is called a tracker muon.

  The tracker muon reconstruction thus relies on the tracker measurements, and uses the muon system only for tagging the particle type.

- **Global muons** are constructed by propagating SA muon tracks from the muon system inwards to search for matching tracker tracks. A combined fit is implemented using the Kalman filter technique, using information from both the tracker and the SA muon tracks to build a final global muon track.

The full analysis uses a holistic approach, using an optimal mix of the different types of reconstruction, depending on the targeted detector region, to benefit from the improved resolution of the tracker and global muons close to the beam spot but also from the extensive coverage of the standalone muons. For promptly produced muons within the geometrical acceptance of the muon system, about 99% of such muons are reconstructed as either tracker or global muons, and very often as both [94]. The efficiencies for displaced muons have been studied and documented in [29, 77, 88, 94].

Due to the fact that only the displaced version of the standalone reconstruction is used, some clarification is needed regarding the terminology used in subsequent sections: Displaced standalone muons (DSA) muons are referred to as **STA muons**. Muons using the tracker information (i.e., tracker and global muons) will be referred to as **TMS muons**.

Chapter 4

Data and Simulation

4.1 First Year of Run 3 Data (2022)

The search is based on the analysis of pp collision data collected at a center-of-mass energy $\sqrt{s} = 13.6$ TeV during the 2022 data-taking year in Run 3 of the LHC. The events of interest are recorded by a set of dedicated double muon triggers, described in Section 5.2. The data are collected into the `DoubleMuon` and `Muon` datasets, listed in Table 4.1. The data are divided into 'Eras', each of which corresponds to a specific set of conditions, such as luminosity profiles, trigger configurations, detector conditions, and data quality, and are optimized for specific physics analyses. The most recent version of the data reprocessing, global tags, and JSON files are chosen according to the recommendations of the Physics Data and Monte Carlo Validation (PdmV) group at CMS. As the analysis does not rely on calorimeter information, we can maximize the number of valid data-taking runs by including even runs or lumi-sections[1] in which the calorimeter information is sub-optimal. Thus, we obtain a total integrated luminosity of 36.6 ± 0.8 fb^{-1} [97].

Dataset	$\mathcal{L}_{\text{int}}[\text{fb}^{-1}]$
/DoubleMuon/Run2022C-10Dec2022-v1/AOD	0.61
/Muon/Run2022C-10Dec2022-v1/AOD	5.56
/Muon/Run2022D-10Dec2022-v1/AOD	3.28
/Muon/Run2022E-10Dec2022-v2/AOD	5.93
/Muon/Run2022F-PromptReco-v1/AOD	18.12
/Muon/Run2022G-PromptReco-v1/AOD	3.08
Total	36.60

Table 4.1: A list of the 2022 data sets used in the analysis as certified by CMS, with their corresponding integrated luminosities.

The mean pileup in 2022 was 44 interactions per bunch crossing, with a peak of up to 60 interactions per bunch crossing.

[1] A lumi-section corresponds to 23 s of data taking

4.2 Signal Simulation

To motivate the characteristics and properties of displaced dimuons from BSM processes, we use two signal models in this search, introduced in Section 2.3. These are the **Hidden Abelian Higgs model (HAHM)**, the **BSM heavy scalar** model, and an **R-parity violating Supersymmetric (RPV SUSY)** model. Monte Carlo (MC) samples were simulated for these models, the details of which are provided in the next few sections.

4.2.1 HAHM Model

The HAHM signal samples are produced on a discrete grid of dark photon mass $m(Z_D)$ and kinetic mixing parameter ϵ, the values of which are listed in Table 4.2 and illustrated in Figure 4.1. The values of the corresponding mean proper lifetime and the mean decay vertex position are shown in the last two columns.

$m(Z_D)$ [GeV]	$\mathcal{B}(Z_D \to \mu\mu)$	ϵ	$c\tau(Z_D)$ [mm]	$\langle L_{xy} \rangle$ [cm]
10	0.15	$1 \cdot 10^{-6}$	1.14	0.77
		$5 \cdot 10^{-7}$	4.55	2.89
		$1 \cdot 10^{-7}$	114	72.0
		$3 \cdot 10^{-8}$	1264	800
20	0.15	$5 \cdot 10^{-7}$	2.17	0.72
		$2 \cdot 10^{-7}$	13.6	4.14
		$5 \cdot 10^{-8}$	217	66.4
		$1 \cdot 10^{-8}$	5425	1665
30	0.14	$3 \cdot 10^{-7}$	3.90	0.80
		$1 \cdot 10^{-7}$	35.1	6.73
		$3 \cdot 10^{-8}$	390	74.7
		$7 \cdot 10^{-9}$	7165	1374
40	0.15	$2 \cdot 10^{-7}$	6.21	0.86
		$8 \cdot 10^{-8}$	38.8	5.10
		$2 \cdot 10^{-8}$	621	81.4
		$5 \cdot 10^{-9}$	9937	1301
50	0.13	$2 \cdot 10^{-7}$	4.42	0.47
		$6 \cdot 10^{-8}$	49.1	4.37
		$1 \cdot 10^{-8}$	1768	158
		$4 \cdot 10^{-9}$	11049	982
60	0.11	$1 \cdot 10^{-7}$	11.8	0.70
		$4 \cdot 10^{-8}$	73.7	3.93
		$7 \cdot 10^{-9}$	2405	129
		$2 \cdot 10^{-9}$	29464	1591

Table 4.2: HAHM signal sample parameters used in the search. For each $m(Z_D)$, the assumed branching ratio $\mathcal{B}(Z_D \to \mu\mu)$ of the dark photon decay into muons is listed, along with values of the kinetic mixing parameter ϵ, and the corresponding proper decay length (in mm) and mean transverse decay position (in cm).

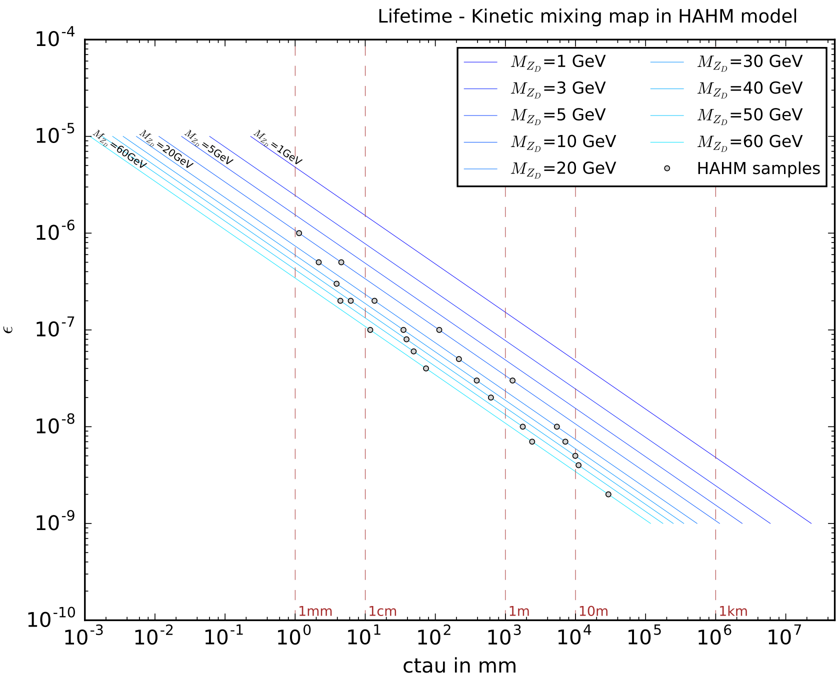

Figure 4.1: Kinetic mixing of a dark photon as a function of its lifetime for different values of the dark photon mass. The open circles correspond to the signal samples generated for the analysis and are listed in Table 4.2.

The production of dark photons is modeled at leading order by `MadGraph5_aMC@NLO` [98] version 2.9.9. The samples use the NNPDF3.1 NNLO PDF set [99]. About 500,000 events are generated at each signal $c\tau$ point, using the CP5 tunes [100] to model the underlying event. The dominant gluon-fusion production mechanism is used for the samples. The Higgs boson production cross-section is then normalized to the most recent theoretical prediction of the sum of all production modes. For $m(h) = 125$ GeV at $\sqrt{s} = 13.6$ TeV, this value is 59.8 pb [101]. The decays of the dark photon are simulated by PYTHIA 8.306 [102]. To reduce processing time and storage, only the events with at least one muon in the final state are considered. The filter efficiency depends on the $\mathcal{B}(Z_D \to \mu\mu)$ and ranges between 20% to 28%. The large number of generated events together with a grid in kinetic mixing allows us to interpolate between samples with different kinetic mixing using a lifetime reweighting method, described in Appendix C.

4.2.2 Heavy Scalar model

The SM-like Higgs is not necessarily the only particle that can initiate an LLP decay chain featuring displaced muons in the final state. In principle, the LHC energies allow for the production of even heavier mediators. To explore a larger phase-space of signal kinematic properties and topologies than the HAHM, the BSM ‘Heavy Scalar’ benchmark model features generic heavy scalar particles H that decay to spin-0 mediators X, with a broad spectrum of masses $m(H)$ and $m(X)$. Two sets of benchmark signals are produced, depending on whether one or both LLPs were forced to decay into muons.

- For $H \to XX \to 4\mu$, both scalars X decay to final state muon pairs.
- For $H \to XX \to 2\mu 2j$, only one X is required to decay to a muon pair, while the other decays into a pair of light quarks ($u\bar{u}, d\bar{d}, s\bar{s}$) with equal probabilities, resulting in a final state with only one muon pair. As the quarks hadronize quickly, the resulting final states are collectively referred to as 'jets', and not further considered in the analysis.

$m(H)$ [GeV]	$m(X)$ [GeV]	$c\tau(X)$ [mm]	$N_{events}(2\mu 2j)$	$N_{events}(4\mu)$
125	20	13, 130, 1300	150,000	75,000
	50	50, 500, 5000		
200	20	7, 70, 700	75,000	35,000
	50	20, 200, 2000		
400	20	4, 40, 400	40,000	20,000
	50	8, 80, 800		
	150	40, 400, 4000		
1000	20	2, 20, 200	20,000	12,000
	50	4, 40, 400		
	150	10, 100, 1000		
	350	35, 350, 3500		

Table 4.3: Model parameters for the $H \to XX$ benchmark signal samples used in the search. Both $H \to 2X \to 2\mu 2j$ and $H \to 2X \to 4\mu$ samples were generated at each combination of $m(H)$, $m(X)$ and $c\tau(X)$, with the last two columns showing the number of events generated per lifetime point for each final state.

The 2022 $H \to XX$ signal samples are generated with Pythia 8.306 [102] using the CP5 tune [100] and the NNPDF3.1 NNLO PDF set [99]. Samples in each set are generated with different combinations of $m(H)$ and $m(X)$. The simulated width of the H boson has a negligible effect on the analysis. Each combination of masses, $(m(H), m(X))$, was produced with three lifetime ($c\tau(X)$) points, corresponding to mean transverse decay lengths spanning three orders of magnitude (e.g., 3 cm, 30 cm, and 300 cm). These chosen values, listed in Table 4.3 allow us to study a wide range of signal displacements, event kinematics and event topologies. The dimuon acceptance and selection efficiencies strongly depend on the mass of the signal, therefore the number of generated events varies with the H boson mass.

4.2.3 RPV SUSY Model

The search uses benchmark signal samples generated assuming $\mathcal{B}(\tilde{\chi}_1^0 \to \mu\mu\nu) = \mathcal{B}(\tilde{\chi}_1^0 \to ee\nu) = 0.5$, giving rise to events with up to two displaced dimuons. The chosen $m(\tilde{q})$, $m(\tilde{\chi}_1^0)$, and $c\tau(\tilde{\chi}_1^0)$ grid is listed in Table 4.4. To explore a wide range of kinematic variables and event topologies, for each chosen $m(\tilde{q})$, we generate sets of samples with $\Delta m = m(\tilde{q}) - m(\tilde{\chi}_1^0)$ of 25, 200 and 650 GeV, and sets with constant $m(\tilde{\chi}_1^0)$ values of 50 and 500 GeV such that $m(\tilde{\chi}_1^0) < m(\tilde{q})$. Some of the $m(\tilde{\chi}_1^0), m(\tilde{q})$ combinations have already been used in previous CMS [57, 58] and ATLAS [59] searches performed at lower $\sqrt{s}$. All other SUSY particles (e.g., gluinos and sleptons) are assumed to be too heavy to be produced, and their masses are set to 10 TeV. As with the Heavy Scalar model, the generated $c\tau$ values span three orders of magnitude.

$m(\tilde{q})$ [GeV]	$m(\tilde{\chi}_1^0)$ [GeV]	$c\tau(\tilde{\chi}_1^0)$ [mm]	$\sigma^{13.6\text{ TeV}}(\tilde{q}\bar{\tilde{q}})$ [pb]	N_{events}
125	50	20, 200, 2000	7323 ± 458	150,000
	100	35, 350, 3500		
200	50	15, 150, 1500	835 ± 51.9	75,000
	175	40, 400, 4000		
350	50	8, 80, 800	50.5 ± 3.21	50,000
	150	25, 250, 2500		
	325	45, 450, 4500		
700	50	5, 50, 500	0.941 ± 0.0719	35,000
	500	40, 400, 4000		
	675	50, 500, 5000		
1150	50	3, 30, 300	0.0313 ± 0.00343	35,000
	500	30, 300, 3000		
	950	50, 500, 5000		
	1125	60, 600, 6000		
1600	50	2, 20, 200	0.00215 ± 0.000393	35,000
	500	20, 200, 2000		
	950	40, 400, 4000		
	1400	60, 600, 6000		
	1575	70, 700, 7000		

Table 4.4: Model parameters for the simulated RPV SUSY samples ($\tilde{q} \to q\tilde{\chi}_1^0$) used in the analysis, and cross sections for squark-antisquark pair production at $\sqrt{s} = 13.6$ TeV. The last column shows the number of events generated per lifetime point for a given $m(\tilde{q})$.

The samples are generated with PYTHIA 8.306. The squark-antisquark production cross sections [60] are calculated with `NNLL-fast` version 2.0 to approximately next-to-next-to-leading order (NNLO) in the strong coupling constant and include the resummation of soft gluon emission at next-to-next-to-leading logarithmic accuracy [103]. The computation uses the NNLO PDF4LHC21 parton distributions functions (PDFs) [104]. Since the dimuon acceptance and selection efficiency strongly depend on the mass of the signal, the number of generated events varies with the $m(\tilde{q})$.

4.3 Background Simulation

The search presented in this thesis targets LLPs with a mass above 10 GeV. There are no genuine LLPs in this mass range in the SM, and most pp collision backgrounds are instrumental due to reconstruction mistakes, either from fake particles (which are not matched to the true generated particles) or from misidentification of properties due to mismeasured parameters. As simulations do not accurately reproduce the rate of these effects, the optimization of the analysis selection and the estimation of backgrounds is performed using data. The official Run 2 MC backgrounds samples were used to help guide our understanding of the behavior and properties of the background, however, they were not actually used in the background estimation.

To complement performance measurements in data, studies were also conducted using simulated samples of B-hadrons decaying into $J/\psi \to \mu\mu$ in a cascade, listed below.

- $B_s \to J/\psi\phi; \phi \to KK; J/\psi \to \mu\mu$
- $B_u \to \psi(2S)K_s; \psi(2S) \to J/\psi\pi\pi; J/\psi \to \mu\mu$
- $B_u \to X(3872)K_s; X(3872) \to J/\psi\rho; \rho \to \pi\pi; J/\psi \to \mu\mu$

Chapter 5

The Search for Displaced Dimuons

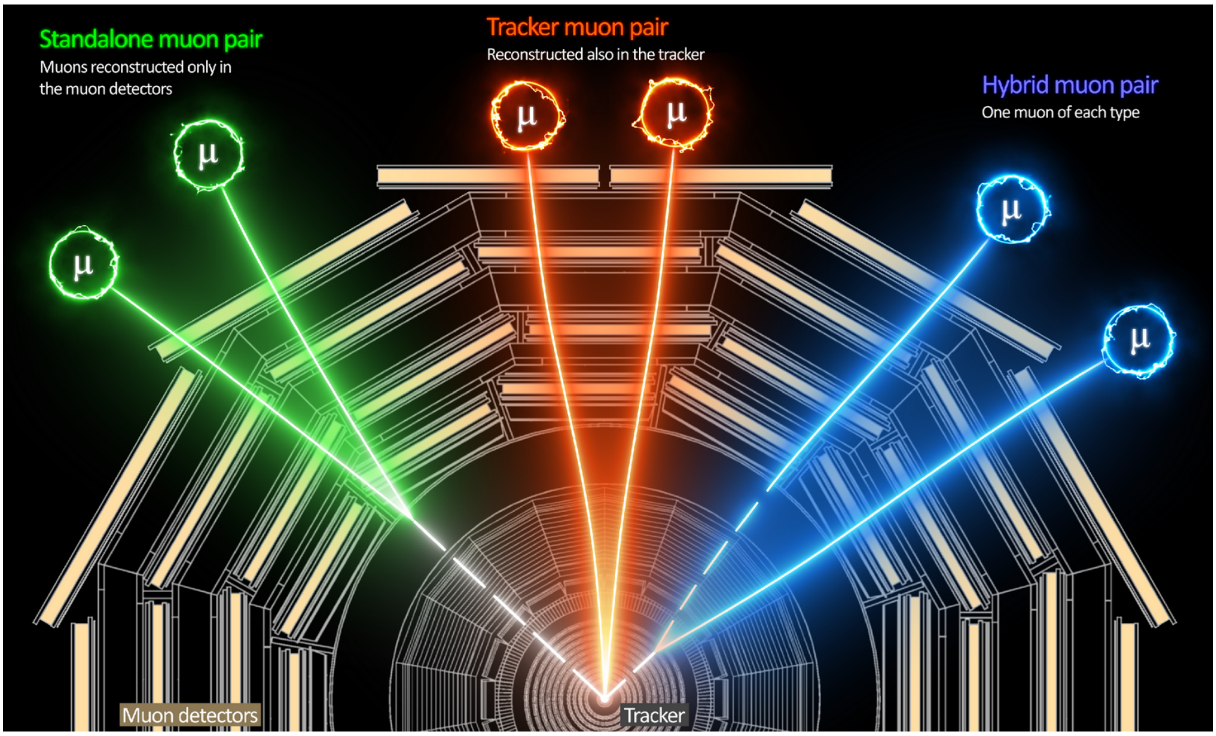

Figure 5.1: Illustration showing two long-lived particles with different lifetimes decaying into a pair of muons, depicting how the signals of the muons can be traced back to the long-lived particle decay point using data from the tracker and muon detectors.

The work presented in this thesis was carried out at the Compact Muon Solenoid (CMS) detector at the LHC, in proton-proton collisions with a center-of-mass energy of $\sqrt{s} = 13.6$ TeV, the highest achieved till this point, making it the first search for BSM physics at this energy [23]. A brief outline of the search is described below.

Particles produced in the hard interaction of colliding protons are usually associated with a vertex at the beam spot position, known as the **primary vertex (PV)**. The PV is chosen to be the vertex with the largest momentum transfer – i.e., the largest scalar sum of particle tracks associated with it. If the particles are produced immediately, forming tracks with PVs at the beam spot, they are **'prompt'** particles. A long-lived particle produced at the PV will travel some distance in the detector before decaying; in the context of our analysis, the decay products are muons. Consequently, a pair of reconstructed muon tracks are fit to a common **secondary vertex (SV)**, which, depending on the average decay length, may be displaced with respect to the PV. Muon pairs originating from an SV are **'dimuons'**, and if their common origin is separated from the beam spot by more than a few hundred micrometers, they are 'displaced dimuons'. The individual muons of a given dimuon are referred to as the muon 'legs' of the dimuon. One of the core capabilities of the CMS detector is its ability to accurately detect muons and measure their properties. As described in Section 3.3.9, the CMS detector is capable of reconstructing muons in two ways:

- In the first method, we can use the hits in the muon system alone, giving us what we call 'standalone' **(STA)** muons; all muons are first reconstructed as STA muons, and therefore muons that are produced even beyond the inner detector layers (i.e., displaced muons) are reconstructed by this method.
- In the second method, we are able to use the information from the tracker in addition to the muon system **(TMS)** to reconstruct muons to obtain much more precise measurements of the muon trajectory and momentum. Muons that originate in the tracker, close to the pp collision point, are reconstructed in this way.

Depending on whether the muons are produced within the tracker volume, they are reconstructed as either STA or TMS muons. From the extrapolated tracks of these muons, we can then reconstruct dimuon SVs of three kinds: **STA-STA** where both muons are reconstructed in the muon system, **TMS-TMS** where both muons are reconstructed in the CMS tracker, and **STA-TMS** where we have one STA and one TMS muon. An illustration of these three categories is shown in Figure 5.1.

In the Run 2 search [22], these muons were extensively studied and robust identification and selection criteria were developed to distinguish between signal and background events with high sensitivity. This Run 3 search [23] builds on those well-proven ideas, using lessons learned from the former to develop new triggers and refine the selection criteria, greatly improving the sensitivity to signal events. The focus of this thesis is on the development of the new triggers and on the analysis of the STA-STA category of dimuons. A detailed description of the TMS-TMS category is not provided in this thesis; however, measurements performed using TMS-TMS dimuons that are relevant to the trigger performance, or the final results, are described in this chapter, and a brief overview of the TMS muon identification and the selection criteria used to define the TMS-TMS signal region is provided in Appendix B. Lastly, while the STA-STA category provides the most sensitivity to LLPs decaying beyond 50 cm in the transverse place, and the TMS-TMS category provides the best sensitivity for LLPs decaying within the tracker volume, up to 20 cm in the transverse plane, the STA-TMS category provides only modest sensitivity gains over the combination of these two categories in the intermediate range between 20-50 cm.

In this search, we focus on the effectiveness of the new triggers to improve sensitivity to the signal using only the STA-STA and TMS-TMS dimuons, even with a limited dataset, while the STA-TMS category will be included in a future publication with the full Run 3 data.

To design the search, we use three common BSM benchmark signal models, introduced in Section 2.3: the Hidden Abelian Higgs Model (HAHM) [42, 43], featuring an added dark sector Higgs boson that mixes with the SM Higgs boson and decays into two long-lived dark photons, which further decay into a pair of muons; a generic model featuring a heavy scalar boson that decays into two spin-0 long-lived mediator particles [61], which can then decay into muons; and an R-parity violating (RPV) SUSY model [55], featuring a long-lived neutralino decaying into a pair of muons and a neutrino. Finally, we interpret the results of this search in terms of the above two models.

5.1 Main discriminating variables

Figure 5.2 illustrates a typical LLP decay to muons and the key variables involved. The key variables used in the analysis of STA-STA dimuons are described in the next section. Figure 5.2 illustrates these key variables and their relationship to each other.

Transverse impact parameter d_0

The impact parameter, d_0, is a key discriminant in the analysis and plays an important role in both the trigger design and the analysis. It is defined as the shortest distance in the transverse plane between the PV and the track of the muon extrapolated towards the PV. Promptly produced muons have small values of d_0, and their tracks are referred to as 'pointing' tracks, while muons produced with a displaced vertex tend to have large d_0, and have 'non-pointing' tracks. Figure 5.3 (top left) shows the distributions of the minimum d_0 of the two muons, normalized to unity to facilitate the comparison of the shape, for some selected simulated $H \to XX \to 2\mu 2j$ signal points.

The transverse decay length, L_{xy}

The distance between the PV and the SV in the plane transverse to the beam line is the transverse decay length, denoted by L_{xy}. We prefer to use the 2D measurement rather than the 3D L_{xyz} measurement due to the fact that the precision of the measurement is much better in the transverse direction than along the longitudinal direction. LLPs decaying at different L_{xy} can have significantly different properties. For instance, decays with $L_{xy} \leq 10$ cm take place in the innermost part of the tracker volume, leaving hits in most tracker layers, and therefore, the muons are likely to be reconstructed in the high-precision tracker. This class of events is subject to high rates of backgrounds from processes such as dimuons from Drell-Yan decays, or muons emerging from the decays

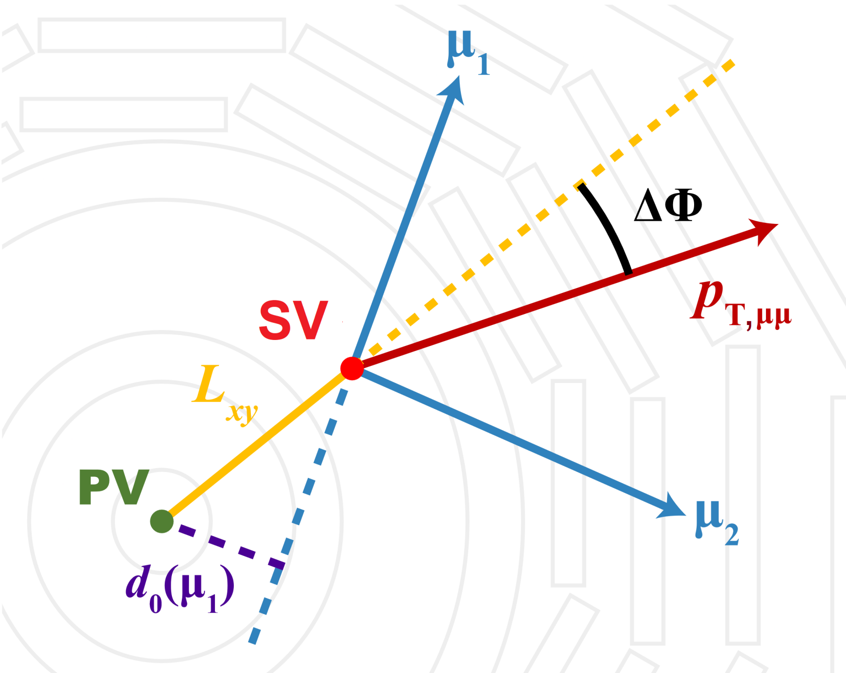

Figure 5.2: Diagram of the key variables used in the analysis. A long-lived particle is produced at the primary vertex (PV) and travels some distance in the detector before decaying into two muons (μ_1 and μ_2), forming a common secondary vertex (SV). The distance between these two vertices in the plane transverse to the beam line is referred to as the transverse decay length, L_{xy}. The angle between the $\vec{L}_{xy}$ vector and the dimuon transverse momentum vector, $\vec{p}_{T,\mu\mu}$, is the transverse collinearity angle $\Delta\Phi$. The transverse distance between the PV position and the point of closest approach of a muon to the primary vertex is the transverse impact parameter d_0. The diagram is reproduced from Ref. [96].

of quarks produced in strong interactions. Decays with $L_{xy} > 100$ cm, however, occur outside the tracker acceptance. Muons from such decays can only be reconstructed using the muon system. They have coarser resolution than tracks reconstructed with the tracker but are also subject to a smaller background from particles produced at or very close to the pp collision point. The normalized L_{xy} distributions for some selected simulated $H \to XX \to 2\mu 2j$ signal points are shown in Figure 5.3 (top right).

L_{xy} significance, $L_{xy}/\sigma(L_{xy})$

The uncertainty in L_{xy}, denoted by $\sigma(L_{xy})$, is calculated by combining the uncertainties in the transverse positions of both the SV and the PV. The quantity $L_{xy}/\sigma(L_{xy})$ is then referred to as the 'L_{xy} significance', as it is a measure of the significance of the observed displacement compared to zero in a Gaussian approximation. This quantity is large for a well-reconstructed dimuon that corresponds to a long-lived particle decaying at some distance away from the PV, as in signal events. In contrast, promptly produced background events, such as those from Drell-Yan processes, are expected to result in dimuons with small L_{xy} significance. It is, therefore, one of the main signal-background discriminating variables used in our analysis. Figure 5.3 (bottom left) shows the normalized L_{xy} distributions for some selected simulated $H \to XX \to 2\mu 2j$ signal points.

Transverse collinearity angle $\Delta\Phi$

In general, the momentum vectors of the muons point in different directions than that of the LLP, whose momentum is along the $\vec{L}_{xy}$ vector. For a typical two-body LLP decay as a consequence of momentum conservation, as expected in the signal, the *dimuon* transverse momentum vector, $\vec{p}_{T,\mu\mu}$ is expected to be aligned with $\vec{L}_{xy}$. The collinearity angle in the transverse plane, $\Delta\Phi$, is thus defined to be the angle between the dimuon transverse momentum vector ($\vec{p}_{T,\mu\mu}$) and the transverse decay length $\vec{L}_{xy}$. For a pair of muons produced in the decay of a long-lived particle originating at the PV, the $\vec{p}_{T,\mu\mu}$ and $\vec{L}_{xy}$ vectors are expected to point away from the PV, and in the same direction, resulting in a small $|\Delta\Phi|$. In contrast, in most types of prompt background events, the $\vec{L}_{xy}$ measurement is expected to be uncorrelated with $\vec{p}_{T,\mu\mu}$, resulting in a $|\Delta\Phi|$ distribution that is symmetric around $\pi/2$. A $|\Delta\Phi|$ distribution that peaks at zero is therefore indicative of a signal and, therefore, serves as a key discriminating variable between signal and background. This variable is further explored in detail in Sections 5.5.8 and 5.6.1, showcasing its power in discriminating between signal and background, as well as providing a method for estimating prompt backgrounds in the signal region.

Other important discriminating variables

We also classify selected dimuons as opposite-sign (OS) or same-sign (SS) dimuons depending on the observed muon charges. While well-reconstructed signal dimuons are OS, many of the background processes featuring small-$|\Delta\Phi|$ OS dimuons also yield small-$|\Delta\Phi|$ SS dimuons, either because the processes are charge symmetric or due to incorrectly assigning of the muon charge. Most of these muons are embedded in jets from QCD processes, and B-meson cascade decays, and thus, the charge correlation between the muons is a useful discriminating variable.

Muons emerging from cosmic ray showers in the atmosphere are effectively suppressed by requirements on the 3D opening angle between the two muons, α, as well as requirements on the timing of the muon with respect to the bunch crossing, $t_{\text{in-out}}$, and the direction of flight of the muon. These selections are described in more detail in Section 5.5.7.

For signal events, the spectrum of observables depends on the model parameters. For example, the masses of the particles, m_H and m_X, and the LLP lifetimes, $c\tau_X$, are free parameters in the models we consider, and different values will produce vastly different signatures in terms of the observables. This point is illustrated in Figure 5.3 for the following scenarios:

- A heavy long-lived mediator, $m_X \approx m_H/2$ with a large $c\tau_X$, for example with $(m(H), m(X), c\tau(X)) = (125, 50, 5000)$, results in a large L_{xy}, a large d_0 (non-pointing muons), and a large 3D opening angle α between the two muons.

- A heavy but short-lived mediator with $m_X \approx m_H/2$ and small $c\tau_X$, for example with $(m(H), m(X), c\tau(X)) = (125, 50, 50)$, would result in small L_{xy}, large α, and generally smaller d_0 than the previous scenario.

- A light long-lived mediator (producing boosted particles) with $m_X << m_H/2$ and large $c\tau_X$, for example with $(m(H), m(X), c\tau(X)) = (125, 20, 1300)$, results in large L_{xy}, small α and small d_0 (pointing muons).

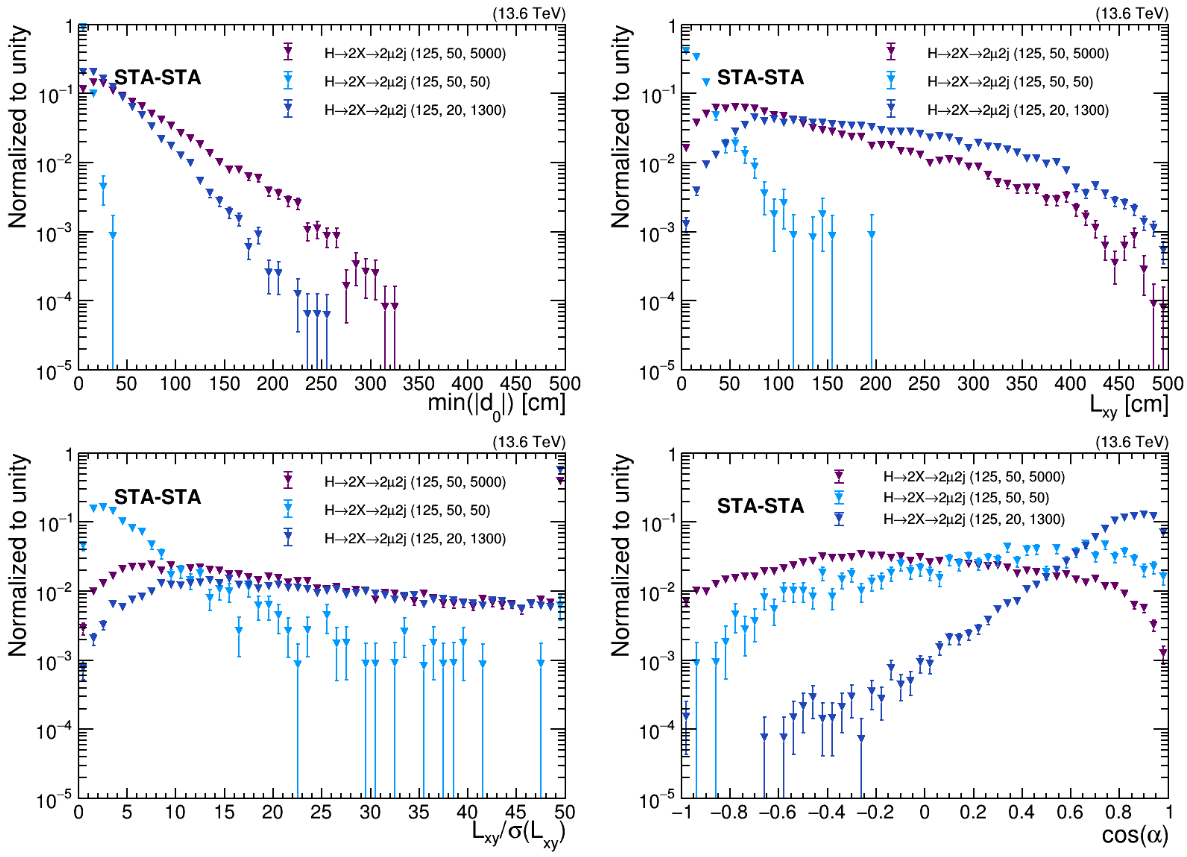

Figure 5.3: Normalized distributions of min(d_0) (top left), L_{xy} (top right), $L_{xy}/\sigma(L_{xy})$ (bottom left) and $\cos\alpha$ (bottom right) in STA-STA dimuons, shown for three illustrative scenarios in the $H \to XX \to 2\mu 2j$ signal model.

Figure 5.4 shows an event display of a typical simulated $H \to XX \to 2\mu$ signal event with an LLP decaying outside the tracker.

These variables do not by themselves characterize the signal processes completely while rejecting every kind of background. However, a carefully optimized combination of selection thresholds on these variables allows us to define signal regions (SRs) that are highly sensitive to signal events while at the same time being nearly background-free. We are also able to, by inverting some of the selections, obtain control regions (CRs) that

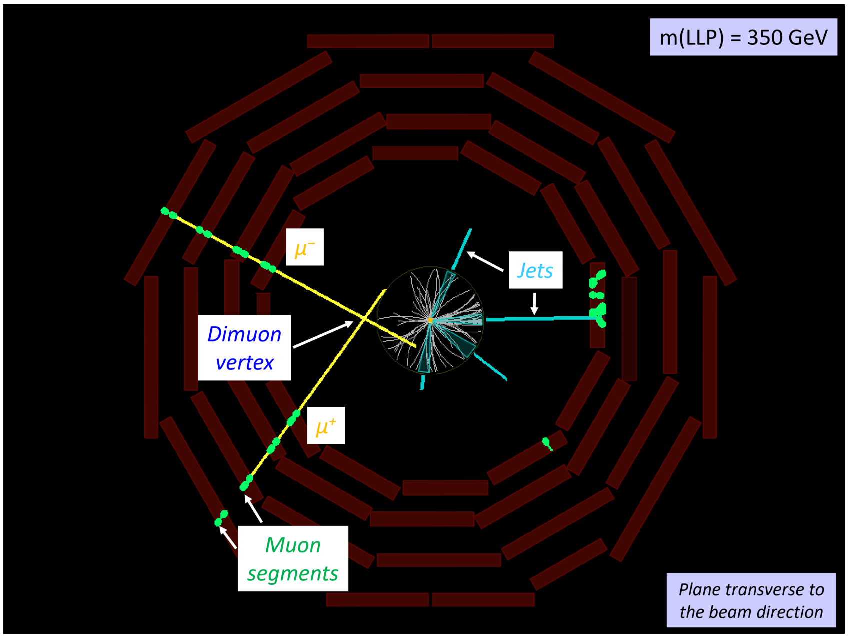

Figure 5.4: Event display of a typical simulated $H \to XX \to 2\mu$ signal event (x-y view). A long-lived particle with $m(X) = 350$ GeV decays into two muons with generated transverse momenta $p_T^{gen}(\mu^-) = 336$ GeV and $p_T^{gen}(\mu^+) = 184$ GeV outside of the tracker, at $(x, y) =$ (-169.7 cm, 4.6 cm), giving $L_{xy} = 169.9$ cm. Both muons are reconstructed as STA muons (shown in yellow) with reconstructed transverse momenta $p_T^{reco}(\mu^-) = (448 \pm 180)$ GeV and $p_T^{reco}(\mu^+) = (149 \pm 29)$ GeV. These muons form a common vertex with $L_{xy}^{rec} = 170.1 \pm 0.3$ cm, and $L_{xy}/\sigma(L_{xy})) = 597$. The reconstructed dimuon invariant mass $m_{\mu\mu}$ is 364 ± 80 GeV. The reconstructed value of the collinearity angle in the transverse plane is $|\Delta\Phi| = 0.16$.

are enriched in specific classes of backgrounds while being almost free of signal, allowing us to estimate the background contribution that can leak into the signal regions due to mismeasurement or misreconstruction effects.

It is possible to utilize these variables already at the trigger level to improve sensitivity to the signal while rejecting prompt backgrounds. The following section describes the design strategy, implementation, and performance of the triggers used in this search.

5.2 New and Improved Displaced Dimuon Triggers

Triggers are the starting point for data collection in any analysis at the CMS.

The LHC can provide approximately a billion proton-proton collision events per second to the CMS detector. It is impossible with the current technology to record and store all the data that is produced. To facilitate high-quality data acquisition, the CMS experiment has a two-tiered trigger system that utilizes both hardware (L1 Trigger, or L1T) and software (High Level Trigger, or HLT) that captures the interesting events while filtering out the vast majority of events that mostly consist of well-known physics, thus reducing the data taking rate to manageable levels. The CMS trigger system is briefly described in Section 3.3.6.

Well-performing triggers that are sensitive to the displaced dimuons signature are particularly important to this search. Muon triggers in CMS have traditionally been developed for muons that originate at, or very close to, the pp collision point. Their performance would rapidly deteriorate for muons produced a significant distance away from the beam spot, and are not efficient for L3 muons produced in the outer half of the tracker. While events recorded by the triggers can be later studied and rejected in the offline analysis selection, a signal event that is not recorded by the trigger is lost forever.

For the Run 2 search for displaced dimuons [22], we used dedicated HLT paths that trigger on L2 muons, reconstructed using information only from the muon systems as described in Section 3.3.8. In 2016, the HLT path (`HLT_L2DoubleMu28_NoVertex_2Cha_Angle2p5_Mass10`) used requirements largely inherited from the triggers used in the Run-1 displaced dimuon searches [57, 58]. The trigger was then reintroduced in 2018 in a revised form (`HLT_DoubleL2Mu23NoVtx_2Cha`), using lowered p_T thresholds (from 28 GeV to 23 GeV) and removing requirements (that the 3D angle between the muons is less than 2.5 rad, and that the invariant mass of the two L2 muons is $>$ 10 GeV) to make the triggers less model dependent, which significantly increased the trigger efficiency in the Heavy Scalar model. This HLT path uses the standard reconstruction of L2 muons. A complementary HLT path (`HLT_DoubleL2Mu23NoVtx_2Cha_CosmicSeed`) was also introduced that used an alternate L2 muon reconstruction with a cosmic-seeded initial state. The 2018 HLT paths took as input the following L1 Trigger double muon and triple muon seeds, `L1_DoubleMu_15_7`, `L1_DoubleMu_15_5_SQ` (which imposes a higher quality requirement on the muon track fit), and `L1_TripleMu_5_3_3`.

For Run 3 data taking, the primary goal was to further improve the signal acceptance by lowering the trigger p_T thresholds as much as possible, allowing us to be model-independent without increasing the trigger rate significantly. We designed and optimized our trigger strategy using lessons learned from the Run 2 search.

5.2.1 Trigger Strategy

The main limitation of the Run 2 triggers is that the beam spot constraint in the L1 triggers caused a rapid degradation in the L1 trigger efficiency as the muon displacement increased, as described in Section 3.3.8. Figure 5.5 shows this degradation, as a function of d_0, in cosmic muon data and signal simulation for each of the L1 trigger p_T thresholds used in 2016 and 2018.

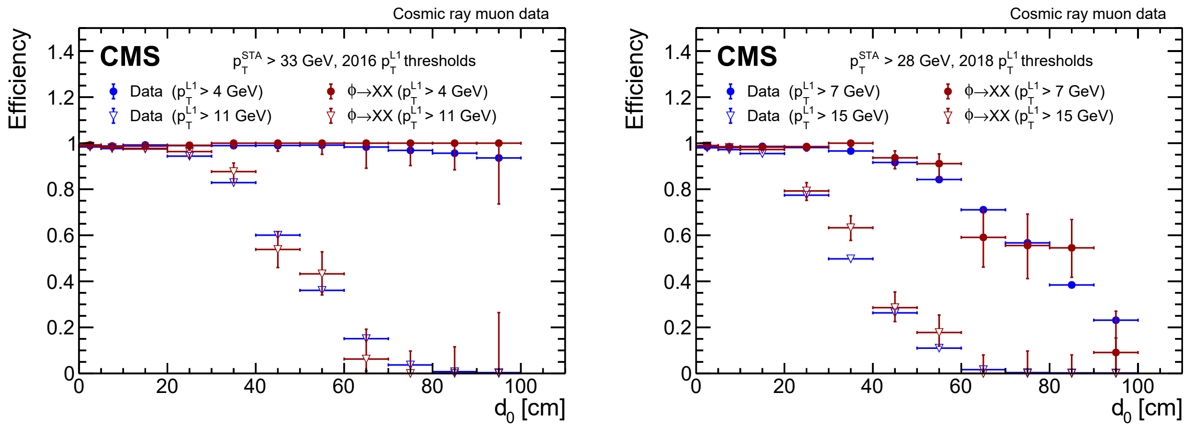

Figure 5.5: L1 muon trigger efficiencies in cosmic muon data (blue) and signal simulation (red) as a function of d_0 for the L1 trigger p_T thresholds used in (left) 2016 and (right) 2018 data. A detailed study of Run 2 efficiencies can be found in Ref. [29].

The efficiency fall-off is steeper for higher p_T thresholds, mainly because of the beam spot constraint in the L1 trigger, as it reconstructs displaced muons with a lower p_T than they actually have, thus falling below the threshold and failing the event. This inefficiency can be addressed and improved in two ways:

- By removing the beam spot constraint in the track fit to give a more accurate p_T measurement for displaced muons.
- By lowering the L1 p_T thresholds in prompt-muon seeds so that even muons with underestimated p_T can pass the trigger. However, this would give substantially higher rates and introduce more background events.

We employ both these strategies in the design of the triggers for Run 3. To determine the signal efficiency and compare the gain with respect to the 2018 suite of triggers, we use signal samples generated for the Hidden Abelian Higgs Model and the benchmark Heavy Scalar model.

5.2.2 L1 Triggers

As a first measure, we employ two existing L1 double muon seeds, namely

- `L1_DoubleMu0er1p5_SQ_OS_dR_Max1p4`

- L1_DoubleMu4p5_SQ_OS_dR_Max1p2

These are designed for low-mass dimuons that originate in heavy flavor decays. These seeds have much lower p_T thresholds, and make use of the L1 muon candidate quality (based on the number of hits in different muon stations) and kinematic constraints on $|\eta|$, ΔR between the two muons, and their charge to keep the rate manageable. These seeds remain efficient for muons in our p_T range even at high d_0, as their p_T thresholds are tolerant to the p_T underestimation.

The second method was to include new L1 double muon seeds,

- L1_DoubleMu0_Upt15_Upt7
- L1_DoubleMu0_Upt6_IP_Min1_Upt4

These L1 seeds are developed specifically for Run 3. These triggers use, for the first time, L1 muons reconstructed without applying the beam spot constraint in the L1 muon track finders. During 2022 data taking, both seeds use a Kalman filter [92] based track finding in the barrel only (KBMTF, with $|\eta| < 0.83$). The KBMTF provides two new quantities – the unconstrained p_T (Upt), and the muon impact parameter d_0 (IP) – obtained from a Kalman fit [105]. Due to the bit-size constraints in the hardware, the IP measurement is available in four d_0 bins of 25 cm – [0, 25), [25,50), [50, 75) and [75, ∞). The L1_DoubleMu0_Upt15_Upt7 seed is a replica of L1_DoubleMu_15_7 using unconstrained p_T instead of the regular (beam spot constrained) p_T, while the latter L1 seed utilizes the IP information by restricting $d_0 > 25$ cm in order to lower the Upt to 6 GeV for the leading muon, and 4 GeV for the sub-leading muon.

These two sets of L1T seeds were added in a logical OR combination to the seeds already used in 2018. The Run 3 L1T configuration used in 2022 is listed in Table 5.1 and illustrated in Figure 5.6.

The gain at L1 relative to the 2018 L1 triggers from the two sets of triggers is shown in Figure 5.7, with gain from the displaced KBMTF-based L1T seeds added in OR to the 2018 L1T seeds shown in blue, and the gain from adding also the low p_T L1T seeds in OR (resulting in the full 2022 L1T configuration) shown in purple. Here, we apply the 2018 HLT paths and the Run 2 analysis selection (detailed in Table 1 of Ref. [22]) to illustrate the impact of the new L1 trigger strategy. The triggers recover significant efficiency at larger lifetimes, with the UPT seeds improving the trigger efficiency by over 40%, while the low p_T dimuon seeds further improve the trigger efficiency by over 80%.

5.2.3 High Level Trigger

These L1 Trigger seeds serve as the input for the software-level trigger of CMS, the HLT. To benefit from the lowered L1 p_T thresholds and the higher efficiency for displaced muons, we use complementary trigger strategies at the HLT as well, together with the 2018 HLT paths. This strategy is illustrated in Figure 5.8.

L1 Trigger Seeds	Requirements
L1_DoubleMu_15_7	(Constrained) muon p_T: 15, 7 GeV
L1_DoubleMu_15_5_SQ	(Constrained) muon p_T: 15, 5 GeV
	High Quality Muons
L1_TripleMu_5_3_3	(Constrained) muon p_T: 5, 3, 3 GeV
L1_DoubleMu0er1p5_SQ_OS_dR_Max1p4	No muon p_T thresholds
	Muon $\|\eta\| < 1.5$
	Opposite sign muons
	$\Delta R_{\mu\mu} < 1.4$
	High Quality Muons
L1_DoubleMu4p5_SQ_OS_dR_Max1p2	(Constrained) muon p_T: 4.5, 4.5 GeV
	Opposite sign muons
	$\Delta R_{\mu\mu} < 1.2$
	High Quality Muons
L1_DoubleMu0_Upt15_Upt7	(Unconstrained) muon p_T: 15, 7 GeV
L1_DoubleMu0_Upt6_IP_Min1_Upt7	(Unconstrained) muon p_T: 6, 4 GeV
	Highest p_T (leading) muon IP > 25 cm

Table 5.1: Description of the L1 seeds used in the Run 3 triggers. The first set of triggers (top) were inherited from the Run 2 displaced dimuon triggers. The second set of triggers (middle) was used in a related search in Run 2 [54] and has now been introduced in Run 3. The last set of triggers (bottom) was designed for Run 3, using new displaced muon track finders in the barrel (KBMTF) to measure muon p_T and impact parameter without the beamspot constraint.

5.2.4 Run 2 (2018)

We retain the HLT paths used in the Run 2 analysis with updated L1T seeds, listed here for posterity:

- `HLT_DoubleL2Mu23NoVtx_2Cha`
- `HLT_DoubleL2Mu23NoVtx_2Cha_CosmicSeed`

These HLT paths require L2 muons that leave hits in at least two chambers of the muon system (indicated by `2Cha`) and have been reconstructed with a p_T threshold of 23 GeV (without the beam spot constraint, as indicated by `NoVtx`). They do not have any requirements on displacement and are thus sensitive to muons that emerge close to the beam spot, as well as for muons produced beyond the tracker. The latter path uses the cosmic-seeded variant of the L2 muon reconstruction (indicated by `CosmicSeed`). They are labelled ‘Run 2 (2018)’ in Figure 5.8.

5.2.5 Run 3 (2022, L2)

The first strategy, labelled ‘Run 3 (2022, L2)’ in Figure 5.8, implements some features of the Run 2 offline analysis into the online trigger selection. Similar to the ‘Run 2 (2018)’

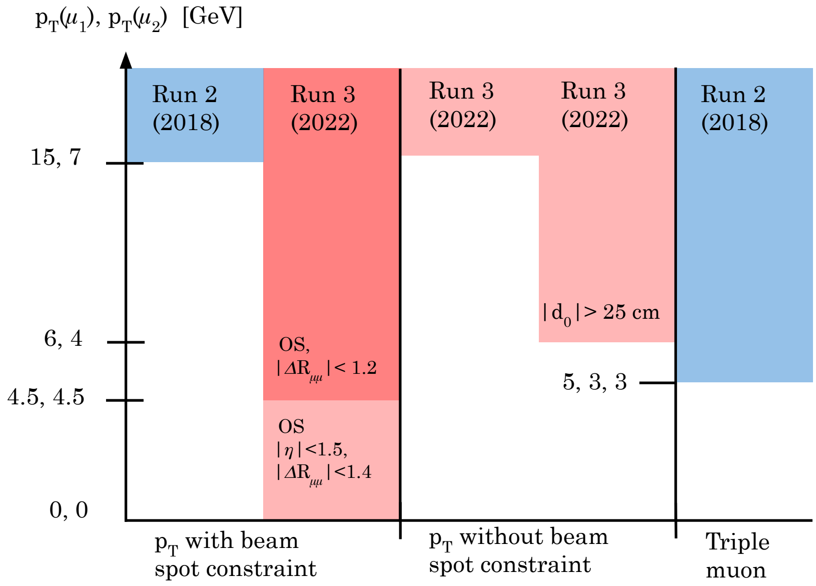

Figure 5.6: The p_T and phase space coverage of the 2018 Run 2 (blue) and 2022 Run 3 (red) L1 trigger seeds used in the search.

HLTs, these require two L2 muons in the event. However, we do not stop the online reconstruction at the L2 stage but further attempt to reconstruct L3 muon candidates as well (analogous to offline global muons), allowing us to take advantage of the superior resolution of L3 muons to determine the transverse impact parameter d_0 of the muon using tracker information. If any of the two L2 muon candidates can be reconstructed as an L3 muon with $d_0 < 1$ cm with respect to the beam spot, the event is discarded (as indicated by `VetoL3Mu0DxyMax1cm` in the path name). Such an L2 muon candidate is very likely to originate from prompt background processes.

The resulting trigger paths,

- `HLT_DoubleL2Mu10NoVtx_2Cha_VetoL3Mu0DxyMax1cm` and
- `HLT_DoubleL2Mu10NoVtx_2Cha_CosmicSeed_VetoL3Mu0DxyMax1cm`

have much lower p_T thresholds of 10 GeV for both muons and significantly improve the signal efficiency in the STA-STA category while only adding about 10 Hz of trigger rate. Here, too, the latter path uses the cosmic-seeded variant of the L2 muon reconstruction, indicated by `CosmicSeed`.

5.2.6 Run 3 (2022, L3)

The second strategy, labelled 'Run 3 (2022, L3)' in Figure 5.8, introduces new paths relying on the online L3 reconstruction, described in Section 3.3.8. It was designed to complement the 'Run 3 (2022, L2)' triggers, bringing elements of the offline TMS-TMS category selection to the trigger and allowing us to use the higher precision measurements

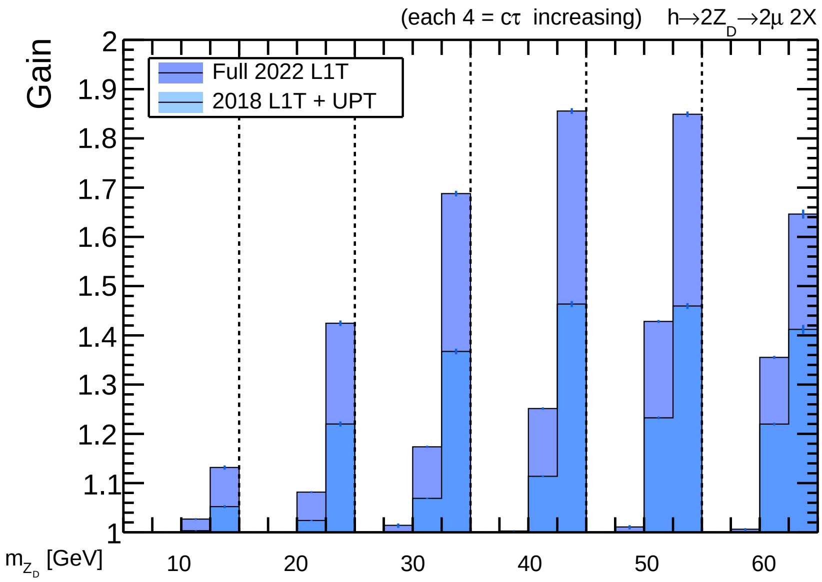

Figure 5.7: Gain in efficiency relative to 2018 from the sequential addition of the new displaced KBMTF-based L1 triggers (blue) and the low p_T L1 triggers (purple) in the HAHM signal samples. The latter is the full Run 3 L1 trigger configuration. The gain is defined as the ratio of simulated events passing the OR of the L1 triggers added in 2022 with the 2018 L1 triggers to the events passing only the 2018 L1 triggers, after applying the full Run 2 event selection [22]. The four divisions in each mass bin show $c\tau(Z_D)$ in increasing order of magnitude, as listed in Table 4.2.

of the tracker to improve sensitivity to LLPs decaying within the tracker while at the same time suppressing prompt backgrounds. We begin by requiring that the event have two L2 muons with p_T thresholds of 15 GeV and 7 GeV for the leading and subleading muons, respectively. We then require that these muons are also reconstructed as L3 muons, with a moderate threshold on the impact parameter of each muon of $d_0 > 0.01$ cm with respect to the beam spot. The L3 muon p_T thresholds are kept relatively low, at 16 GeV for the leading and 10 GeV for the subleading muon. The resulting trigger path,

- `HLT_DoubleL3Mu16_10NoVtx_DxyMin0p01cm`

significantly improves the signal efficiency in the TMS-TMS category while adding only 10 Hz of rate to the CMS trigger. Figure 5.9 illustrates the gain in HAHM signal at low p_T, achieved due to the new triggers.

The trigger efficiencies from the three sets of triggers, ‘Run 2 (2018)’, ‘Run 3 (2022, L2)’ and ‘Run 3 (2022, L3)’, in the HAHM signal samples with $m(Z_D) = 20$ GeV are shown in Figure 5.10. The ‘Run 2 (2018)’ set of triggers are most efficient at the shortest lifetimes due to the fact that they do not cut on the displacement at all. The ‘Run 3 (2022, L3)’ triggers provide a large gain in trigger efficiency, of about a factor of 2.5, between 0.01 cm and 20 cm, compared to the ‘Run 2 (2018)’ configuration. At lifetimes longer than 20 cm, the improvement in trigger efficiency is driven by the ‘Run 3 (2022, L2)’ set

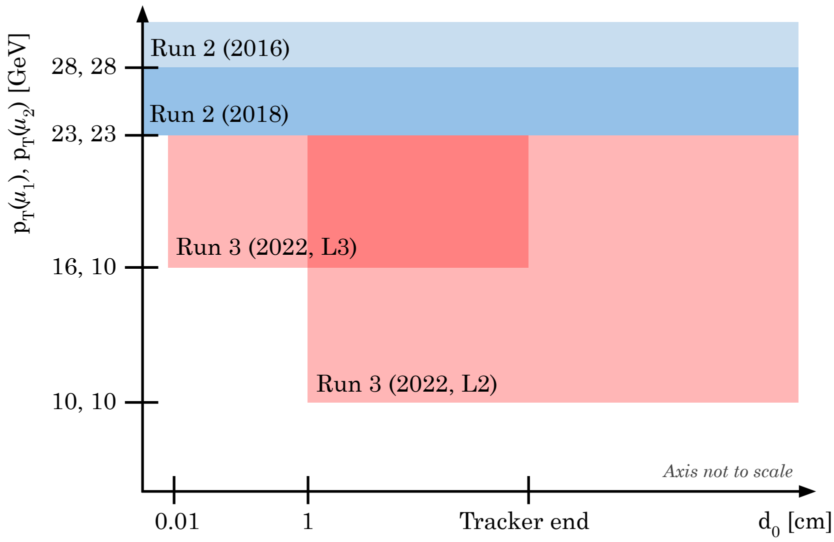

Figure 5.8: The p_T and d_0 coverage of the 2016 Run 2 triggers (light blue), 2018 Run 2 triggers (blue), and the new 2022 Run 3 triggers (red). The two values of the p_T refer to the trigger thresholds for the leading and subleading muon, respectively. The illustration is reproduced from Ref. [23].

of triggers, providing a gain of up to a factor of 4, depending on $m(Z_D)$, relative to the 'Run 2 (2018)' configuration.

A similar gain is observed for the other HAHM mass points. Figure 5.11 shows the trigger efficiencies for the 'Run 2 (2018)' and 'Run 3 (2022)' set of triggers for each of the mass points, for two lifetime points at $c\tau(Z_D) = 1$ cm and 10 m.

5.3 Trigger Performance

The 2018 triggers were extensively studied during the Run 2 analysis [22, 29]. For the new triggers, we studied the performance with prompt and displaced dimuons in 2022 data. To verify that the triggers were effective at rejecting prompt dimuons while remaining highly efficient in selecting displaced dimuons, we used a sample of TMS-TMS dimuons recorded by the 'Run 2 (2018)' set of HLT paths that rely exclusively on information from the muon detectors. As these paths do not impose a requirement on displacement or use information measured in the tracker, they are well suited to measure the impact of the online d_0 requirements applied to the new paths, as well as the correspondence of online and offline d_0 measurement.

This can be illustrated in a known background sample: Drell-Yan events, with prompt dimuons that have a reconstructed invariant mass distribution peaking at the Z boson mass ($M_Z = 91.1880 \pm 0.0020$ [106]). Thus, we look at the invariant mass distributions for TMS-TMS dimuons with $L_{xy}/\sigma(L_{xy}) < 1$ in events recorded with the 'Run 2 (2018)'

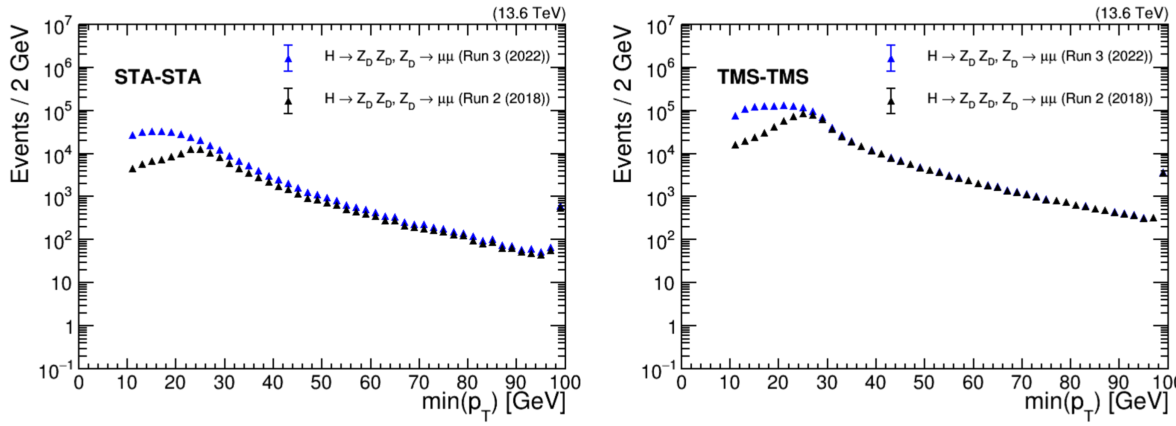

Figure 5.9: The blue curve shows the distributions of the minimum p_T of the two muons for STA-STA dimuons (left) and TMS-TMS dimuons (right) in HAHM signal events that pass the 'Run 3 (2022)' trigger configurations as well as the Run 2 analysis selection. The black curve shows the subset of these events that pass the Run 2 trigger configuration, illustrating the gain in signal sensitivity at low p_T achieved due to the new triggers.

trigger, and in the subset of events that also fired 'Run 3 (2022, L3)' trigger, shown in Figure 5.12 (left). For this study, we use the HLT paths that use the standard L2 muon reconstruction. Figure 5.12 (right) shows the efficiency of the latter trigger relative to the former, which is defined as the ratio of these two distributions and is measured in prompt events. As events in this selection are considered to be prompt background, the background rejection factor is defined as the reciprocal of this efficiency. Due to the veto on dimuons with online d_0 smaller than 0.01 cm, it is possible to reduce the prompt dimuon background by a factor larger than 20, consequently keeping the total trigger rate below 16 Hz during 2022 data taking.

We measure the efficiency of the 'Run 3 (2022, L2)' trigger in prompt and non-prompt dimuons. To factor out possible inefficiencies arising from the L2 muon reconstruction in the study, both HLT paths use the L2 muon reconstruction designed for cosmic-ray muons (`CosmicSeed`). The veto – on L3 muons with $d_0 < 1$ cm matched to one of the L2 muons – allows us to suppress events with prompt dimuons by a rejection factor better than 100 at low mass, to 2000 at high mass, thus keeping the total trigger rate to below 10 Hz during the 2022 data taking period, as shown in Figure 5.13.

To evaluate the efficiency of the L3 muon reconstruction, and the effect of the online $d_0 > 0.01$ cm requirement for displaced muons, we use a similar strategy of studying the $\min(d_0)$ distribution in events recorded by the standard version of 'Run 2 (2018)' trigger, and the subset of these events recorded by the 'Run 3 (2022, L3)' trigger, as shown in Figure 5.14 (left). For this study, we use TMS-TMS dimuons enriched in $J/\psi \to \mu\mu$ candidates with an invariant mass of $3.0 < m_{\mu\mu} < 3.2$ GeV. These events contain prompt as well as displaced dimuons, and illustrate the efficiency of the trigger as a function of the displacement as well as the effect of the d_0 turn-on at 0.01 cm, shown in Figure 5.14 (right). The efficiency is higher than 90% for displaced dimuons with $d_0 > 0.012$ cm in all

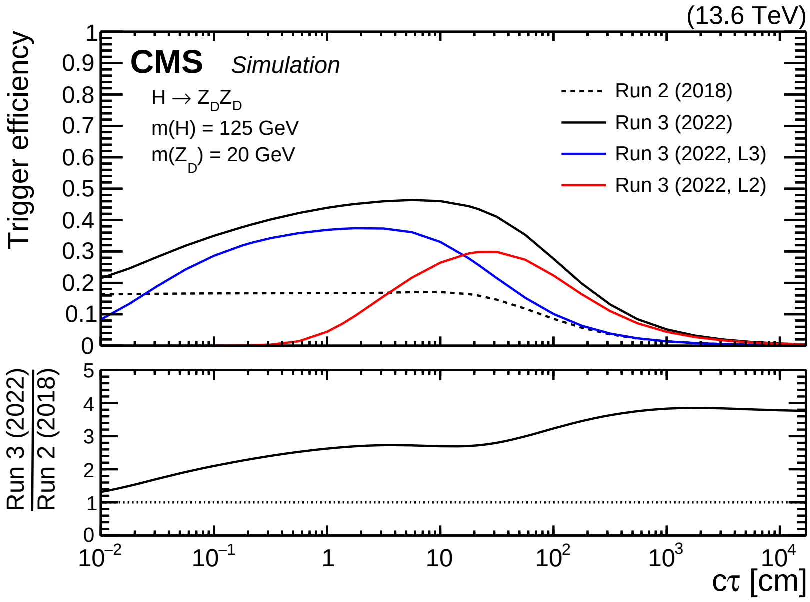

Figure 5.10: Efficiencies of the various displaced dimuon trigger paths and their combination as a function of $c\tau$ for the HAHM signal events with $m(Z_D) = 20$ GeV. The efficiency is defined as the fraction of simulated events that satisfy the detector acceptance and the requirements of the following sets of trigger paths: the Run 2 (2018) triggers (dashed black); the Run 3 (2022, L3) triggers (blue); the Run 3 (2022, L2) triggers (red); and the OR of all these triggers - Run 3 (2022) (black). The lower panel shows the ratio of the overall Run 3 (2022) efficiency to the Run 2 (2018) efficiency. Reproduced from Ref. [23].

data-taking eras, demonstrating its stability over time, with a background rejection factor of up to 20 in prompt dimuons having $\min(d_0)$ close to zero.

We also performed an efficiency measurement independent of the Run 2 dimuon triggers to eliminate any biases from that preselection. We repeated the measurement using $J/\psi \to \mu\mu$ events in a dataset comprised of events recorded by triggers based on jets and E_T^{miss} ('JetMET' dataset). The efficiencies are defined as the subset of $J/\psi \to \mu\mu$ events in the full 2022 JetMET dataset that fired the 'Run 3 (2022, L3)' trigger. The measurement is compared with the efficiency in simulation, defined as the fraction of events in this dataset that fired the same trigger path, obtained by combining all simulated samples of B hadrons that decay into $J/\psi \to \mu\mu$, listed in Section 4.3. About 80% of the events in the combined sample are from the $B_s \to J/\psi\phi$ sample. Figure 5.15 shows trigger efficiencies as a function of $\min(p_T)$ (top left), $\max(p_T)$ (top right), and $\min(d_0)$ (bottom). In each measurement, the other two variables are required to be at the trigger efficiency plateau. The efficiencies reach a plateau close to 65% in data and 70% in simulation, for $\min(p_T) > 12$ GeV, $\max(p_T) > 18$ GeV, and $\min(d_0) > 0.015$ cm, which corresponds to what is expected given the p_T and d_0 thresholds in the trigger. This inefficiency is enhanced in the case of $J/\psi \to \mu\mu$ events, as they typically have $0.05 < \Delta R(\mu\mu) < 0.15$. The data-to-simulation efficiency scale factor and the corresponding systematic uncertainty derived from this measurement are discussed in Section 5.7.

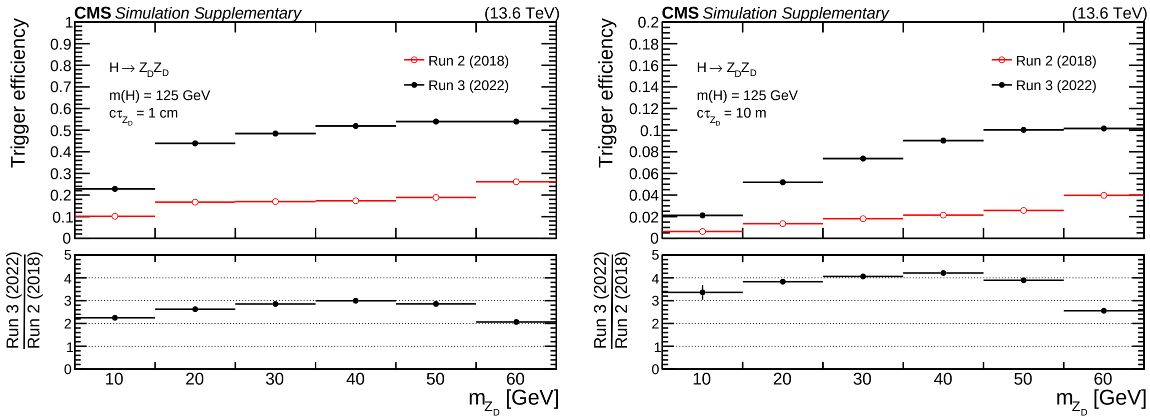

Figure 5.11: Efficiencies showing the Run 2 (2018) (red) and Run 3 (2022) (black) sets of displaced dimuon triggers as a function of $m(Z_D)$ for the HAHM signal events with $c\tau(Z_D) = 1$ cm (left) and 10 m (right). The efficiency is defined as the fraction of simulated events that satisfy the detector acceptance and pass the indicated set of triggers. The lower panel shows the ratio of the Run 3 (2022) efficiency to the Run 2 (2018) efficiency. Reproduced from the supplementary material provided for Ref. [23].

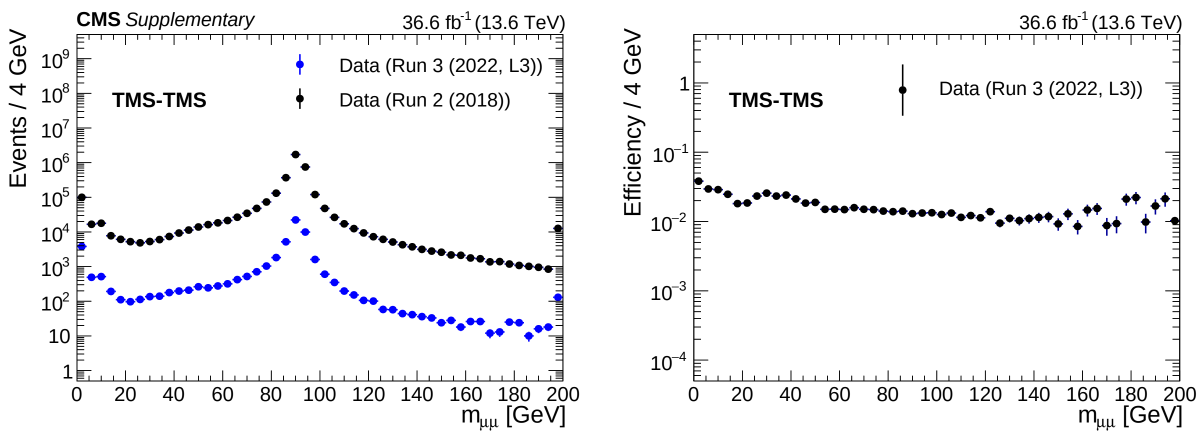

Figure 5.12: The invariant mass distributions (left) of TMS-TMS dimuons with $L_{xy}/\sigma(L_{xy}) < 1$ are shown in events recorded by 'Run 2 (2018)' trigger (black), and the subset of these events that also fired the 'Run 3 (2022, L3)' trigger (blue) – both of which employ the standard version of the L2 muon reconstruction at the HLT. The ratio of these two distributions (right) shows the rejection of prompt background events by the 'Run 3 (2022, L3)' trigger. Reproduced from the supplementary material provided for Ref. [23].

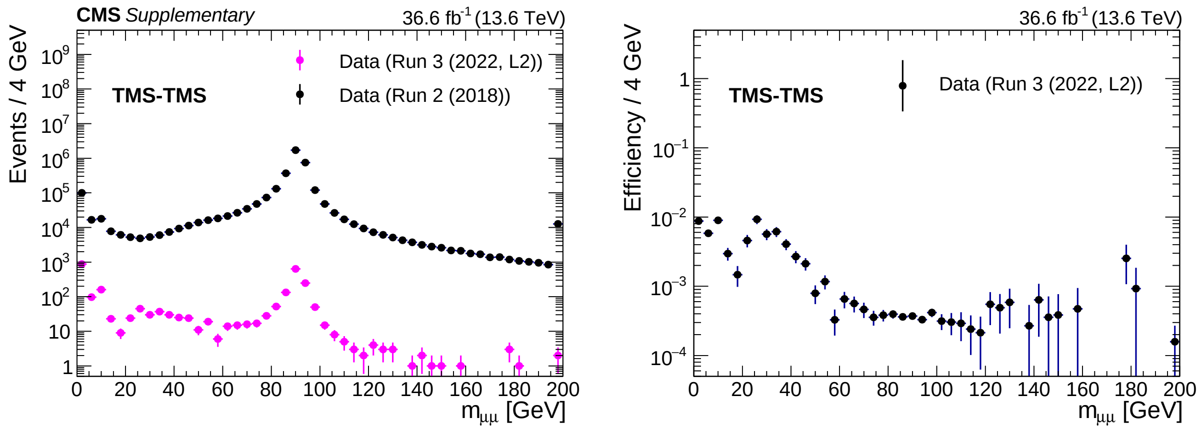

Figure 5.13: The invariant mass distributions (left) of TMS-TMS dimuons with $L_{xy}/\sigma(L_{xy}) < 1$ are shown in events recorded by the 'Run 2 (2018)' triggers (black), and the subset of these events that also fired the 'Run 3 (2022, L2)' triggers (pink), both of which employ the cosmic-ray muon seeded versions of the L2 muon reconstruction. The ratio of these two distributions (left) shows the rejection of prompt background events by the 'Run 3 (2022, L3)' trigger. Partially reproduced from the supplementary material provided for Ref. [23].

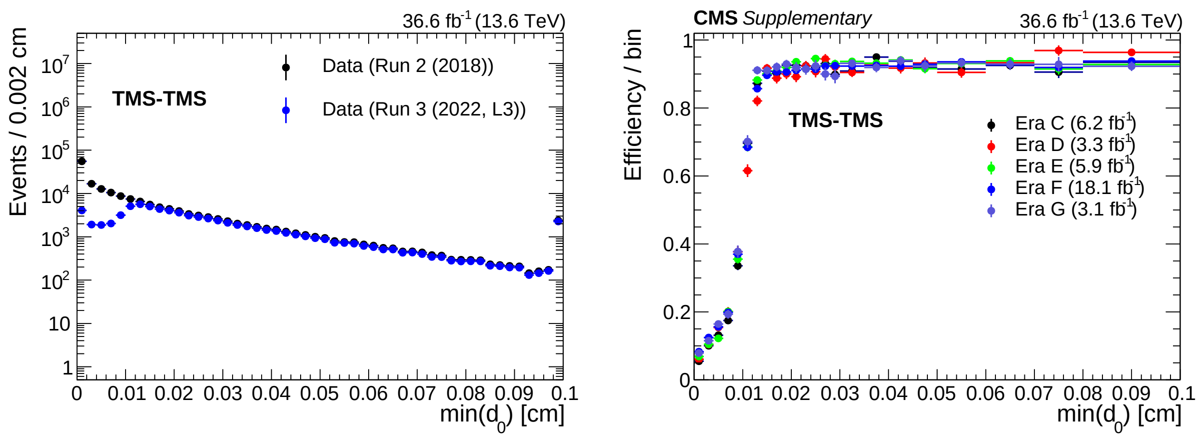

Figure 5.14: The min(d_0) distributions (left) in TMS-TMS dimuons enriched in $J/\psi \to \mu\mu$ events with $3.0 < m_{\mu\mu} < 3.2$ GeV are shown, recorded with the standard version of the 'Run 2 (2018)' trigger (black), and the subset of events that also fired 'Run 3 (2022, L3)' trigger (blue). The efficiency of the 'Run 3 (2022, L3)' trigger path (right), defined as the ratio of these two distributions, is shown in the various data-taking eras represented by different colors. Partially reproduced from the supplementary material provided for Ref. [23].

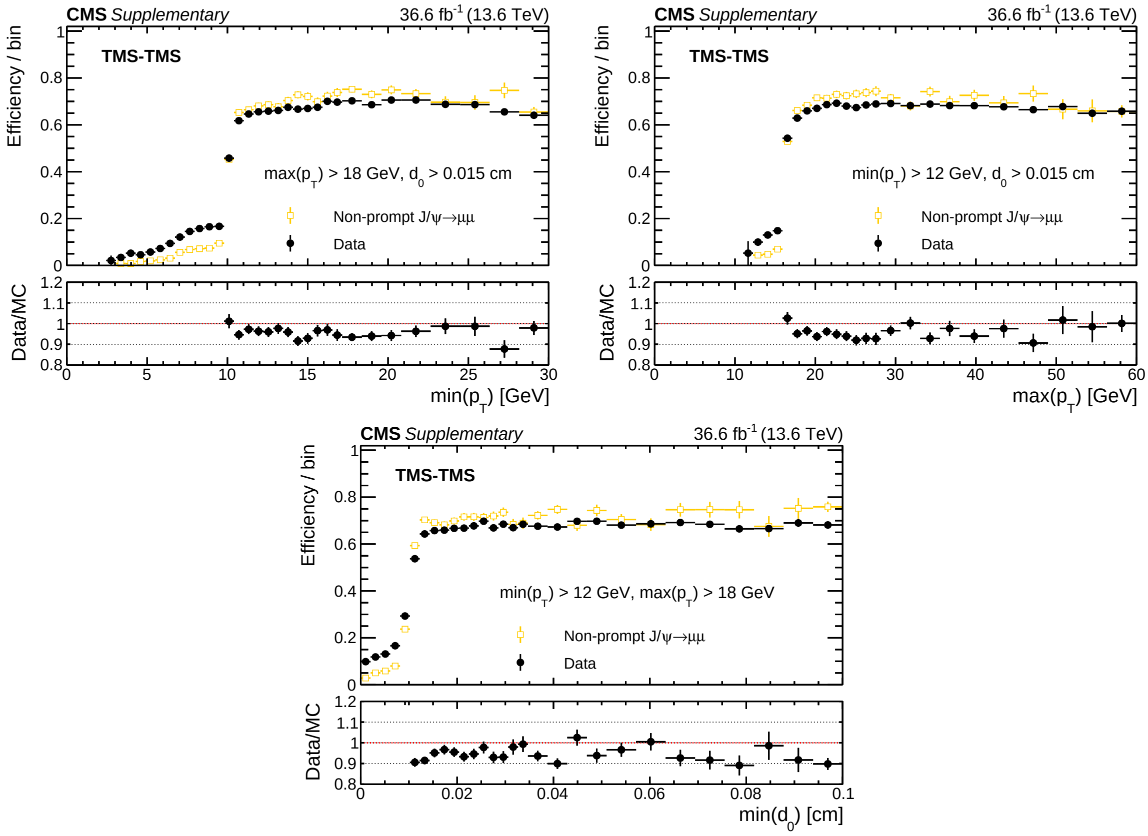

Figure 5.15: The efficiency of the 'Run 3 (2022, L3)' trigger as a function of the min(p_T) (top left), max(p_T) (top right), and min(d_0) (bottom) for TMS-TMS dimuons with $3.0 < m_{\mu\mu} < 3.2$ GeV, recorded in the 2022 JetMET dataset (black), and in a combination of simulated samples of non-prompt $J/\psi \to \mu\mu$ events produced in B hadron decays (yellow). Reproduced from the supplementary material provided for Ref. [23].

To evaluate the performance of the 'Run 3 (2022, L2)' set of HLT paths, namely the efficiency of selecting L2 muons that are not associated to L3 muons with $d_0 < 1$ cm, we use STA-STA dimuons that have a 3D opening angle compatible with a cosmic-ray muon traversing the CMS detector. The veto on the matched L3 muons is sensitive to activity in the tracker. Thus, we only use cosmic-ray muons taken in collision runs recorded by the 'Run 2 (2018)' triggers. The STA-STA dimuon min(d_0) distributions for this reference trigger, and the subset of events recorded by the 'Run 3 (2022, L2)' trigger, are shown in Figure 5.16 (left). Here, too, we use the cosmic-ray muon seeded versions of the two triggers. The efficiency of the latter in displaced dimuons, defined as the ratio of these two distributions, is shown in Figure 5.16 (right). It is higher than 95% for dimuons up to min(d_0) $\approx$ 350 cm in all data-taking eras. The sharp drop in efficiency in the reference trigger is a feature of the cosmic-ray trajectories reaching CMS. For larger d_0, the cosmic-ray muons usually enter parallel to the CSC, or they do not reach the DTs in the lower hemisphere. The residual inefficiency is dominated by the accidental matching between an L2 muon produced by a cosmic ray and a collision L3 muon and is essentially independent of min(d_0).

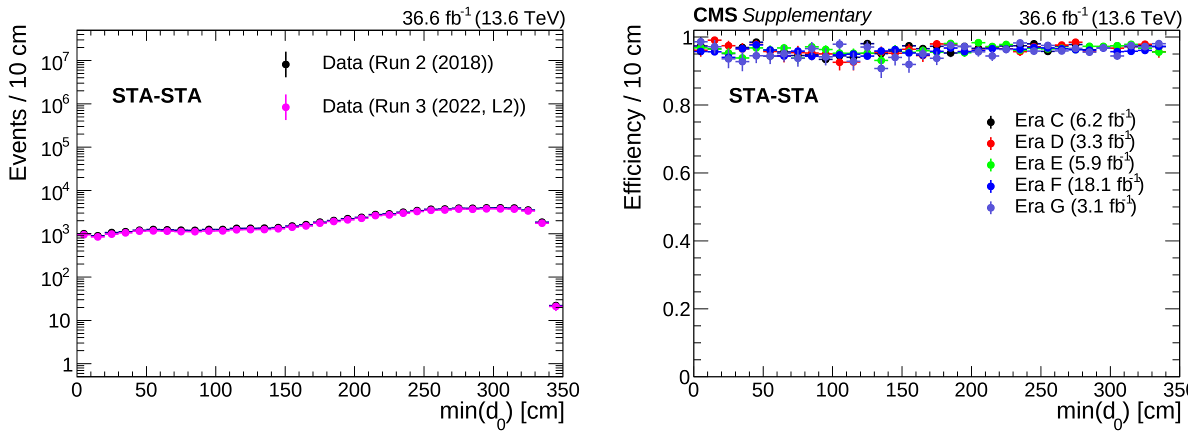

Figure 5.16: The min(d_0) distributions (left) in STA-STA dimuons enriched in cosmic-ray muon events with back-to-back reconstructed muons are shown, recorded with the 'Run 2 (2018)' trigger (black), and the subset of events that also fired 'Run 3 (2022, L2)' trigger (pink) - both of which employ the cosmic-ray muon seeded version of L2 muon reconstruction. The efficiency of the 'Run 3 (2022, L3)' trigger path (right), defined as the ratio of these two distributions, is shown by the various data-taking eras represented by different colors. Partially reproduced from the supplementary material provided for Ref. [23].

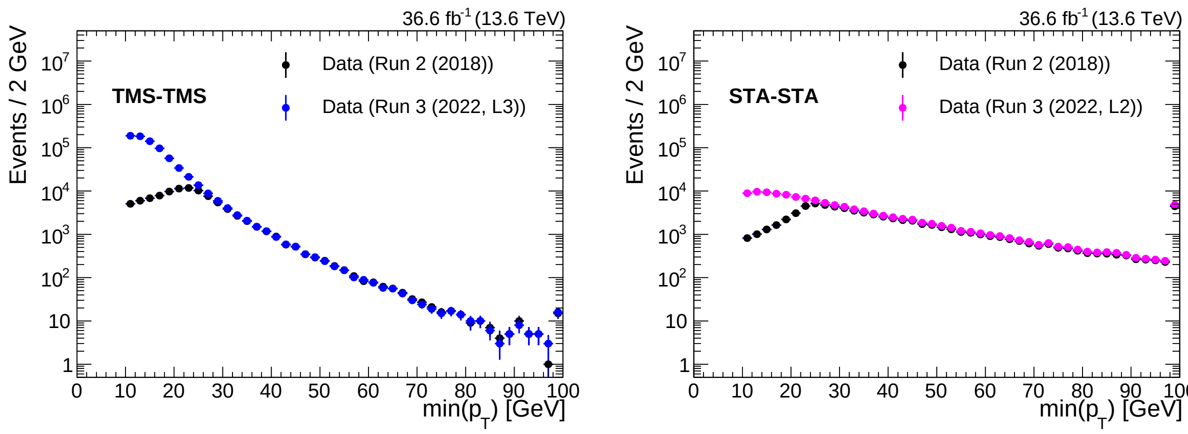

Figure 5.17: The min(p_T) distributions (left) in TMS-TMS dimuons enriched in $J/\psi \to \mu\mu$ events, with $3.0 < m_{\mu\mu} < 3.2$ GeV are shown, recorded with the standard version of the 'Run 2 (2018)' trigger (black) and the 'Run 3 (2022, L3)' (blue) trigger. The min(p_T) distributions (right) in STA-STA dimuons compatible with a cosmic-ray muon track are shown, recorded with the cosmic-ray muon seeded version of the 'Run 2 (2018)' (black) and the 'Run 3 (2022, L2)' (pink) triggers.

Finally, we verify that the new triggers extend the muon p_T spectrum from 23 GeV down to 10 GeV for displaced muons. Figure 5.17 (left) shows the min(p_T) distribution for TMS-TMS dimuons with min(d_0) > 0.02 cm and $3.0 < m_{\mu\mu} < 3.2$ GeV recorded by the standard version of the 'Run 2 (2018)' and the 'Run 3 (2022, L3)' triggers. Similarly, Figure 5.17 (right) shows the min(p_T) distribution for STA-STA dimuons compatible with a cosmic-ray muon track, as recorded by the cosmic-ray muon seeded versions of the 'Run 2 (2018)' and the 'Run 3 (2022, L2)' triggers. For $p_T > 23$ GeV, by design, the data recorded by the new Run 3 triggers overlap with the Run 2 triggers, while for $10 < p_T < 23$ GeV, the new triggers recover the natural shape of the min(p_T) spectrum, confirming the trigger complementarity illustrated in Figure 5.8.

5.4 Muon reconstruction and selection

We search for vertices formed by an intersection of the trajectories of two opposite-sign muons, separated from the PV by displacements between several hundred micrometers to several meters. Different muon reconstruction algorithms give optimal performance at different ranges of displacements, as described in Section 3.3.9. Muons produced near the PV can be accurately reconstructed by using commonly used algorithms developed for prompt muons that combine measurements in the tracker and the muon system, such as the global-muon and tracker-muon reconstruction algorithms. Such muons are referred to as TMS muons in this analysis. However, the efficiency of these algorithms decreases rapidly as the distance between the PV and the muon origin increases, falling to zero when L_{xy} exceeds 60-70 cm. In contrast, muons produced in the outer tracker layers and beyond can still be efficiently reconstructed by standalone (SA) muon reconstruction algorithms that use only hits in the muon system. SA algorithms can reconstruct muons with L_{xy} of up to a few meters, but they have poorer spatial and momentum resolution (with a p_T resolution between 20% and 30%) than muons reconstructed using more precise information from the silicon tracker (with a p_T resolution in the range of 1-3%).

To benefit from the advantages offered by both types of algorithms, and to follow the trigger strategy, we begin the muon selection with the algorithm that is most efficient for the whole range of displacements (i.e., STA muons) and substitute them with more accurately reconstructed global and tracker muons (i.e. TMS muons) whenever such muons are available. STA muons that satisfy a few basic quality requirements form the initial list of the muon candidates retained for the analysis. We then attempt to match each standalone muon in the list with a global or a tracker muon. If a matched global or tracker muon is found, it replaces the associated standalone muon in the list of muon candidates; otherwise, the original standalone muon is retained in the list. Therefore, the analysis uses the optimal combination of standalone and global/tracker muons.

5.5 Event Selection

The large backgrounds are reduced as much as possible by a careful optimization of the event, muon, and dimuon selection, which was performed for the Run 2 analysis as described in Section 4 of [22]. A summary of the selection criteria used in the Run 2 analysis is shown in Table 1 of the same reference. Most of the selection criteria are unchanged. The following sections provide insight into the selection criteria, notably the STA to TMS muon association (Section 5.5.2), the STA and TMS muon identification criteria (Section 5.5.3), the muon pairing algorithm (Section 5.5.5), and criteria used to reject cosmic ray muons (Section 5.5.7). The most important changes with respect to the Run 2 analysis are the addition of the STA muon isolation (Section 5.5.4), and the introduction of two different signal regions in terms of the $\Delta\Phi$ selection (Section 5.5.8). All changes made with respect to the Run 2 analysis are highlighted in the following sections and summarized in Section 5.5.11.

5.5.1 Event Preselection

The preselection begins by requiring that the event is selected by at least one of the trigger paths described in Section 5.2 and has at least two STA muons, each containing more than 12 valid hits in the DT or the CSC:

$$N(\text{hits}) = N(\text{DT hits}) + N(\text{CSC hits}) > 12 \tag{5.1}$$

The N(hits) requirement rejects muon tracks composed of one single DT segment in the barrel, or one or two CSC segments in the endcaps. Such short tracks often originate from hadron punch-through and other background sources. The fraction of signal muons failing the cut is in the range of 1-2%, depending on the signal sample.

Next, we loop over all pairs of HLT Level 2 muons that triggered the event and require that at least one pair of these match the STA muons reconstructed offline having $p_T > 10$ GeV and $|\eta| < 2.0$, with $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} < 0.4$. This requirement, identical to the one used in the Run 2 analysis [22], suppresses events that are triggered on due to muons not related to the signal and allows us to perform trigger efficiency measurements in the analysis. For HLT paths having both Level 2 and Level 3 muons, we only require matching of the STA muons with the Level 2 muons. The number of signal events that pass the trigger but are rejected by this cut varies between 1 and 5%; most of the efficiency loss occurs for the low-mass samples with one or two muons that have $p_T < 10$ GeV but are reconstructed online with a p_T exceeding the trigger thresholds.

The HLT matching requirement ensures that the trigger decision is confirmed using objects reconstructed offline for each event. However, we do not require the selected dimuons to be formed from pairs of trigger HLT muons, allowing the analysis to also be sensitive to 4μ final states in which all four are considered signal muons, but only two of the four muons pass the trigger.

5.5.2 STA to TMS Muon Association

We then try to establish whether a particle that gave rise to a preselected STA muon was also seen in the tracker and reconstructed as a TMS muon. The list of candidate TMS muons consists of all muons that are reconstructed as either global muons, arbitrated tracker muons, or both. Most muons are successfully reconstructed by both algorithms. Candidate TMS muons are associated with STA muons when they satisfy at least one of the two matching criteria:

- **Segment-based matching**: We consider a TMS muon matched to an STA muon if two-thirds or more of the muon segments used to build the STA muon are also used to reconstruct the TMS muon.
- **Proximity match**: We consider a TMS muon matched to an STA muon if $\Delta R = \sqrt{(\eta_{\text{hit}} - \eta_{\text{pca}})^2 + (\phi_{\text{hit}} - \phi_{\text{pca}})^2} < 0.1$, where
  - η_{hit} and ϕ_{hit} are the coordinates of the innermost hit of the STA muon
  - η_{pca} and ϕ_{pca} are the coordinates of the point of closest approach of the TMS muon to the innermost hit of the STA muon.

If an associated TMS muon is found, it replaces the corresponding STA muon in the list of muon candidates used for further analysis. The STA and TMS muons then have to satisfy additional quality selection criteria. Muons passing the quality cuts are used to build three categories of dimuons: STA-STA, STA-TMS, and TMS-TMS, which are analyzed separately. The STA to TMS muon association procedure eliminates most of the background to STA-STA dimuons from LLP decays outside of the tracker (in which case all STA to TMS muon associations can only be accidental) and significantly increases the sensitivity to LLP decays in the tracker, due to the far superior resolution of TMS muons compared to that of STA muons.

Ensuring that the matching procedure is both highly efficient (muons that should be associated actually are) and highly pure (muons that should not be associated are not) requires some care. The matching algorithm, developed for the Run 2 analysis, is the result of several iterations of tuning, making use of the simulated background and signal samples as well as Run 2 data in the signal-free regions. TMS muons possess a much larger number of handles for suppressing background than STA muons and, therefore, a higher purity. Thus, the matching criteria are kept loose to improve the efficiency of the matching. In particular, we do not apply any explicit p_T or η requirements to the candidate TMS muons used by the association algorithm.

Figure 5.18 illustrates the effect of the STA to TMS muon association procedure on the event selection, showing the fraction of $H \to 2Z_D; Z_D \to 2\mu$ signal events with zero (STA-STA), one (STA-TMS) and two STA muons matched to TMS muons (TMS-TMS) as a function of true L_{xy}. In the simulated signal samples, almost all dimuons produced close to the PV have both STA muons matched to TMS muons. The fraction of these events rapidly decreases with L_{xy}^{gen} following the L_{xy} dependence of the tracker reconstruction efficiency. Events with only one matched STA muon start to dominate at $L_{xy}^{\text{gen}} = 20$ cm and remain the dominant component up to approximately 50 cm, after which events with no STA to TMS muon matches dominate, constituting the vast majority of events outside the tracker acceptance. For LLP decays outside the tracker acceptance, $L_{xy}^{\text{gen}} > 65$ cm, the fraction of events removed from the STA-STA category is 10% or less, indicating that the rate of accidental STA to TMS muon matches does not exceed 5%. As a result, the association procedure gives rise to three categories of dimuons, each dominating in a certain L_{xy} range:

- **TMS-TMS**: two STA muons matched to TMS muons, dominating at small SV displacements, $L_{xy} \lesssim 20$ cm.
- **STA-TMS**: one STA muon matched to a TMS muon, sensitive at intermediate SV displacements $20 \lesssim L_{xy} \lesssim 50$ cm.
- **STA-STA**: no STA muon matches to TMS muons, dominating at large SV displacements, $L_{xy} \gtrsim 50$ cm.

The full details of the STA to TMS muon association procedure remain unchanged since the Run 2 analysis [22] and are described in Ref. [96]. The procedure is highly effective in suppressing pp collision backgrounds in the STA-STA category. Studies using Run 2 MC background events in the tails of the $L_{xy}/\sigma(L_{xy})$ and L_{xy} distributions showed that practically all of them are due to prompt muons poorly measured by the STA reconstruction. Since these prompt muons are usually also reconstructed as TMS muons, the STA to TMS association procedure results in a dramatic suppression of pp collision backgrounds in the STA-STA category.

In 2022 data, there are 11.7 million STA-STA dimuons with $|\Delta\Phi| > \pi/2$ that are not rejected by the cosmic-veto requirements (Section 5.5.7), but only 2581 of them remain in the STA-STA category after the STA to TMS muon association procedure, as shown in Figure 5.19 The corresponding background-rejection factor obtained – defined as the ratio of these two numbers – is about 4500 and can be compared to that of the trigger background-rejection factor for the 'Run 3 (2022, L2)' set of triggers tailored for the STA-STA category, which is between 100 and 2000 as shown in Figure 5.13. As the STA muon reconstruction is independent of the TMS muon reconstruction, these events allow us to form a powerful control region in the STA-STA category consisting of STA-STA dimuons with associated TMS-TMS dimuons (denoted by 'STA*-STA*'). We benefit from the

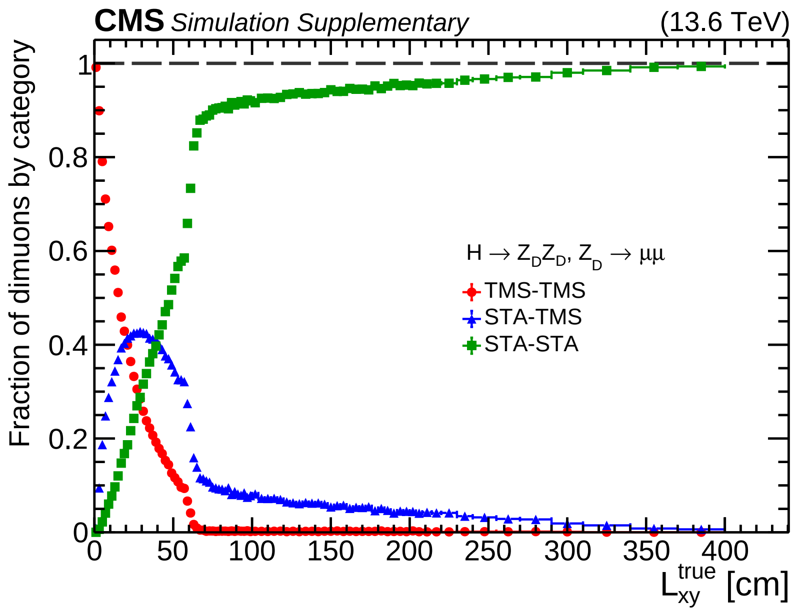

Figure 5.18: Fraction of signal events with zero (green), one (blue), and two (red) STA muons matched to TMS muons by the STA to TMS muon association procedure, as a function of the generated L_{xy}, in the combined $H \to 2Z_D \to 2\mu$ signal samples. The signal dimuons shown in the plot are required to be composed of STA muons with $N(\text{hits}) > 12$ and $p_T > 10$ GeV. In addition, both STA muons have to be matched to generated muons from LLP decays.

superior resolution of these associated TMS-TMS dimuons, and can thus obtain samples enriched in specific types of background and study the properties of the corresponding STA-STA dimuons.

Figure 5.20 compares the L_{xy} resolution of STA-STA dimuons before the STA to TMS muon association with that of TMS-TMS dimuons composed of TMS muons associated with STA muons by the matching procedure, in $H \to 2X \to 2\mu 2j$ events passing the preselection requirements. The quantity plotted is the standard deviation σ of a Gaussian fit to the $L_{xy}^{\text{rec}} - L_{xy}^{\text{gen}}$ distribution of dimuons with both reconstructed muons matched to generated signal muons. The L_{xy} resolution of TMS-TMS dimuons is better than that of STA-STA dimuons by about three orders of magnitude, thanks to the excellent resolution of the tracker. The L_{xy} resolution for STA-TMS dimuons (not shown) is very similar to that of STA-STA dimuons. In each dimuon category, the resolution depends mainly on the opening angle between the muons and on the dimuon transverse displacement.

The three dimuon categories have different spatial and momentum resolutions and are subject to different backgrounds. Therefore, they need to be analyzed separately. This thesis focuses on the STA-STA category only, which is sensitive to the longest lifetime LLP decays in a mostly background-free regime. Thus, the event selection, as well as the background estimation and prediction methods for the STA-STA category, are discussed in detail here. However, it is often instructive to discuss interesting observations in the

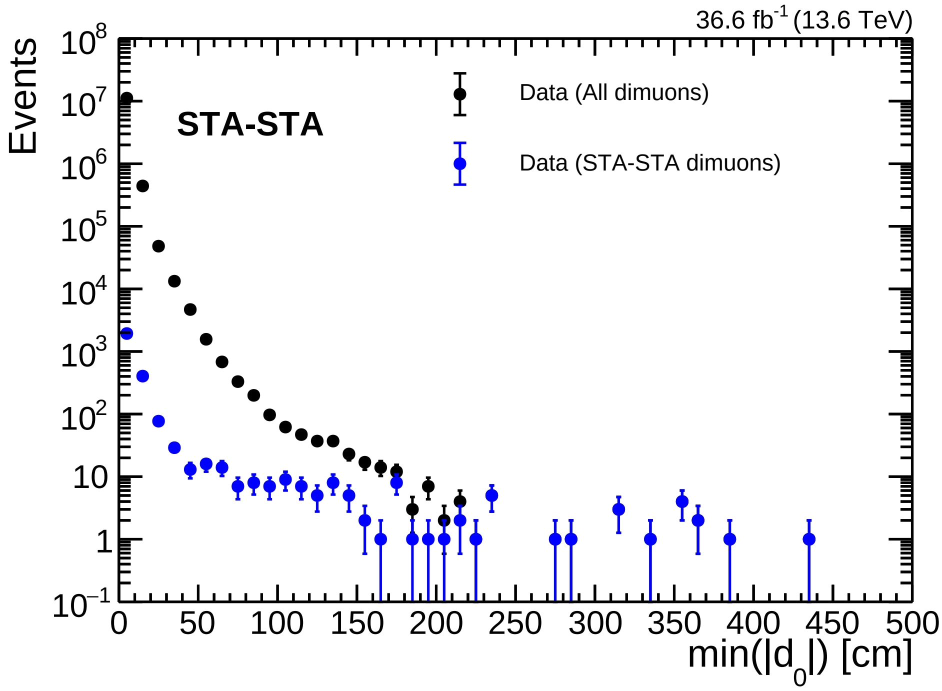

Figure 5.19: Distribution of the minimum impact parameter, $\min(d_0)$, of the two STA muons in all events that survive the selections that reject cosmic muons (black), as well as events that remain in the STA-STA category after the STA to TMS muon association and replacement procedure (blue). This demonstrates the effectiveness of moving the vast majority of prompt dimuon events from the STA-STA category to the TMS-TMS category, where we can benefit from the more precise measurements to discriminate between signal and background events.

TMS-TMS category as well. The results of the STA-STA and TMS-TMS categories are statistically combined for the final result.

The event selection starts by requiring that the event passes the preselection requirements described in Section 5.5.1. In the next step of the analysis, criteria on muon quality are applied to the STA and TMS muons for identification.

5.5.3 STA Muon Identification

The initial list of STA muon candidates entering the STA to TMS muon association procedure discussed in Section 5.5.2 consists of STA muons that have $p_T > 10$ GeV and $N(\text{hits}) > 12$, as described in Section 5.5.1.

To further suppress poorly reconstructed prompt STA muons, which represent the main source of background to STA muons from LLP decays, additional STA muon quality requirements are needed. A standard set of requirements to select good-quality STA muons did not exist prior to the Run 2 analysis, and the analysis team developed a custom set of identification criteria. These criteria were subsequently adopted by CMS as the recommended set of identification criteria for all analyses using STA muons and are described in this section.

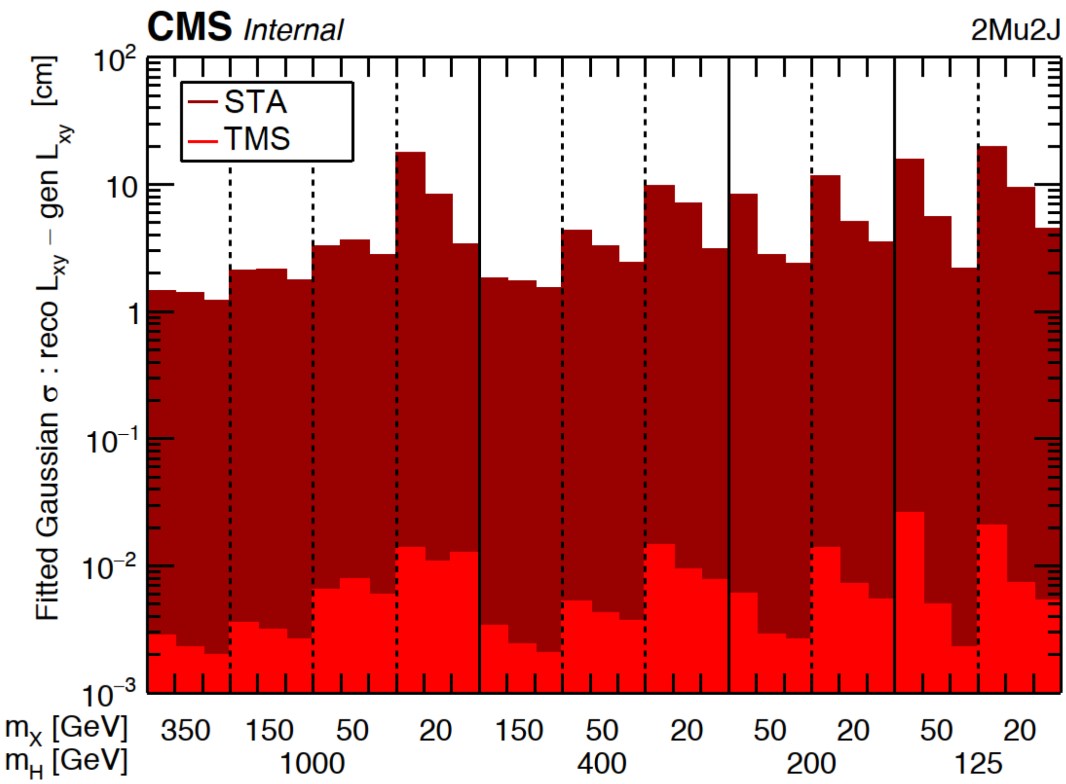

Figure 5.20: Sigma of a Gaussian fit to the $L_{xy}^{\text{rec}} - L_{xy}^{\text{gen}}$ distributions of STA-STA dimuons before the STA to TMS muon association (dark red) compared to that of TMS-TMS dimuons composed of TMS muons associated with STA muons by the matching procedure (bright red). Each bin includes dimuons passing the preselection requirements in one of the 33 $H \to 2X \to 2\mu 2j$ samples with a given combination of signal parameters $(m(H), m(X))$. The $c\tau$ values (not labeled) for each combination of $(m(H), m(X))$ are listed in Table 4.3 and are shown in decreasing order, i.e., longest to shortest.

The STA muon identification criteria were optimized using simulated signal and background samples, as well as samples of prompt STA-STA dimuons misreconstructed as displaced (with $L_{xy}/\sigma(L_{xy}) > 3$) in the signal-free regions in 2016 data, namely, dimuons with $|\Delta\Phi| > \pi/2$ and SS dimuons. Thus, the STA muons are required to meet the following:

- A relative p_T uncertainty less than 100%, where $\sigma(p_T)$ is the uncertainty of the muon track fit: $\sigma(p_T)/p_T < 1.0$
- χ^2/dof of the muon track fit less than 2.5: $\chi^2_{\text{track}}/\text{dof} < 2.5$.
- More than 18 DT hits for muons reconstructed only in the barrel: $N(\text{DT hits}) > 18$ for muons with $N(\text{CSC hits}) = 0$.
- Time with respect to bunch crossing less than 12 ns, obtained from the STA muon: $|t_{\text{in-out}}| < 12$ ns.
- The muon is identified as travelling outwards, setting its direction flag to `direction`=+1 ('inside-out').

Further details on the optimization of the above cut values in terms of signal efficiency and background are provided in [29, 107]. For the Run 3 analysis, the cut values remain the same as in Run 2.

5.5.4 STA Muon Isolation

In the Run 2 search, a custom isolation variable was introduced for the STA muons, I^{rel}_{STA}, defined as the scalar sum of the transverse momenta of tracker tracks in a cone of $\Delta R < 0.3$ around the STA muon (whose trajectory is extrapolated to the point of closest approach to the primary vertex), divided by the STA muon p_T.

Due to the low backgrounds in the Run 2 search after the STA to TMS muon association, explicit requirements on the STA muon isolation were not needed, and thus, we retained sensitivity to models predicting non-isolated displaced muons. The introduction of the new HLT paths with lower p_T thresholds in Run 3 greatly increases the efficiency for the signal, however it also enriches our sample with low p_T muons originating from hadronic activity, for example from semileptonic b and c quark decays in multijet QCD events. To effectively suppress this background, for the search with 2022 data, we apply an isolation criterion to the STA muons.

Dimuons from such backgrounds are often reconstructed as SS dimuons embedded in jets; thus, selecting SS dimuons provides a CR to optimize a threshold on the STA muon isolation variable. This threshold was optimized using SS dimuons in data and events passing the full selection except for the requirements on I^{rel}_{STA} in MC signal samples as shown in Figure 5.21, and found to be $I^{rel}_{\text{STA}} < 0.15$. The signal efficiency of this requirement is higher than 90% while reducing the QCD background by about 50% in SS events. A tighter threshold, such as 0.1, would introduce a significant efficiency loss of up to 20%.

The 'N-1' efficiency of a given selection cut is defined as the ratio of the number of events passing the event selection including the cut, to the number of events passing the event selection except the cut. Figure 5.22 shows the 'N-1' efficiency for different values of STA isolation cut tested during the optimization, in the HAHM (left) and RPV SUSY (right) signal models.

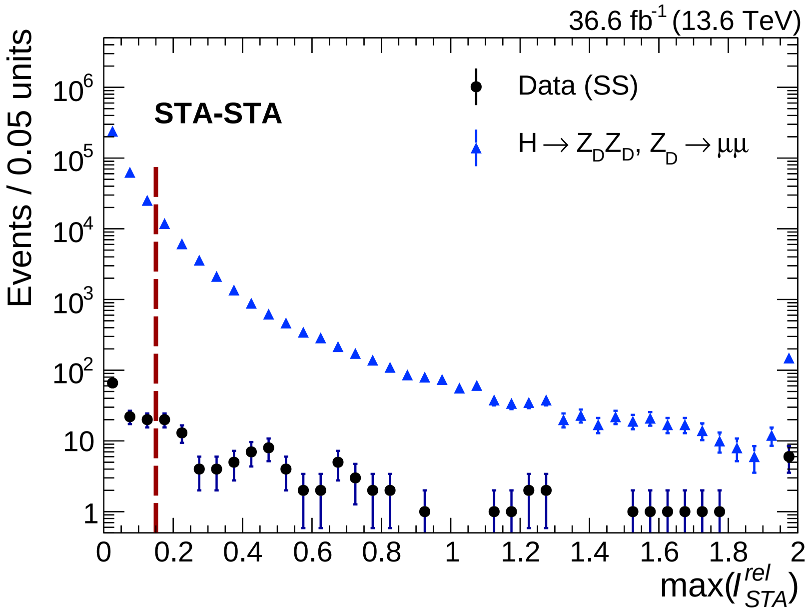

Figure 5.21: Distribution of the maximum of the two I^{rel}_{STA} values for STA-STA dimuons in the SS $|\Delta\Phi| < \pi/4$ control region in data and in all HAHM signal samples combined in the signal region.

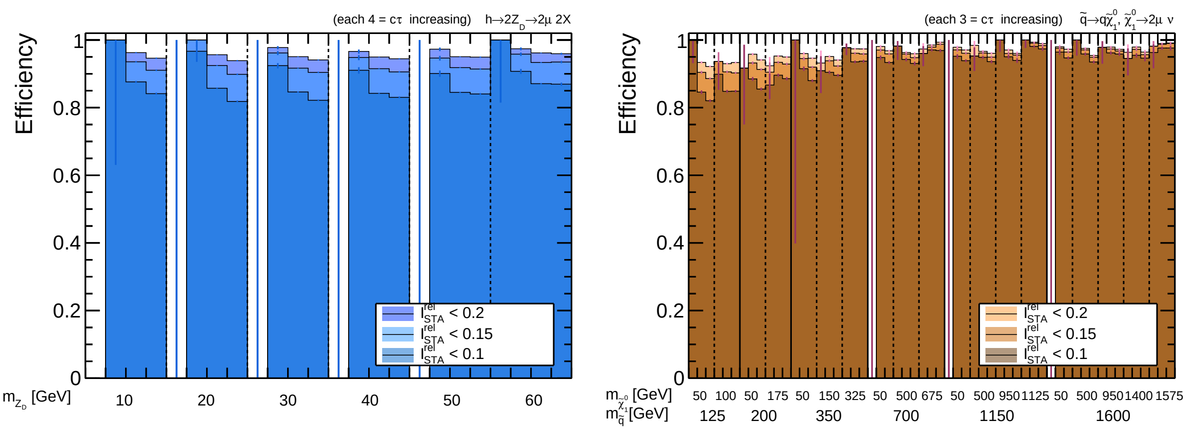

Figure 5.22: The 'N-1' cut efficiencies of different $\max(I^{rel}_{\text{STA}})$ requirements for STA-STA dimuons in each of the HAHM (left) and RPV SUSY (right) signal samples. The dimuons are required to pass the full set of the STA-STA selection criteria, except for the isolation.

The addition of the STA muon isolation criteria also allows us to construct a new QCD-enriched control region in our Run 3 analysis, described in detail in Section 5.6.3, consisting of opposite-sign (OS) STA-STA dimuons with $|\Delta\Phi| < \pi/4$ passing all analysis cuts and inverting the STA muon isolation requirement to evaluate the QCD background.

5.5.5 Common Vertex Fit and Muon Pairing Criteria

At the next step of the analysis, pairs of muons passing the muon identification criterie are used to form the dimuon candidates. The n selected STA muons combinatorially give rise to $n(n-1)/2$ possible pairs of distinct muons, all of which are initially considered. We then require that the distance of closest approach (DCA) of the two muon tracks helically extrapolated in the magnetic field is less than 15 cm. This is a loose requirement to ensure that the common vertex fit is not attempted on a pair of tracks that are far away from each other.

Each pair of selected muon tracks is fit to a common secondary vertex. This common SV fit is performed using the Kalman Vertex Fitter algorithm [92]. While the default implementation in the CMS analysis software restricts the location of the fitted vertex to be approximately within the boundary of the silicon tracker, we extended this range to the beginning of the muon system, allowing us to efficiently reconstruct a common SV for particles decaying outside the tracker, as was done in the Run 1 displaced dimuon analysis [57]. The common SV fit is required to converge, as it is used to define dimuon quantities such as transverse decay length.

The efficiency in the Run 2 search for the common SV fit to converge was close to 100% for STA-STA dimuons with L_{xy} up to ~ 330 cm. The highly displaced vertices beyond this point were affected by a bug in the vertex fitting that did not account for the fact that the magnetic field changes sign, thereby anomalously assigning the $\sigma(L_{xy})$ a value of 100 cm. In the Run 3 search, we updated and reoptimized the SV fit and achieved a near 100% efficiency for L_{xy} up to 500 cm. Figure 5.23 shows the anomalous $\sigma(L_{xy}) = 100$ cm assignment (left) for vertices beyond 330 cm before the update, and the improved SV fit (right) for vertices up to 500 cm after the update. Figure 5.24 shows the common SV fit efficiencies as a function of the generated (true) L_{xy} before (left) and after (right) the update.

The common SV fit refits the input tracks such that they are consistent with originating from a common vertex. These refitted tracks have better p_T resolution than the original tracks. Therefore, all kinematic quantities used in the analysis (such as muon p_T and dimuon invariant mass) are obtained from the refitted STA muon tracks unless specified otherwise. Since the parameters of these refitted tracks are defined at the common-vertex position, these kinematic quantities can be directly compared to the true signal quantities in the generated signal events.

If more than two muons pass the identification criteria, multiple dimuons can be formed, including those that share one of the two muons. Thus, it is important to efficiently choose the reconstructed dimuons that correspond to the signal dimuons. The best performance was achieved with the muon pairing algorithm developed for the Run 2 analysis, depicted

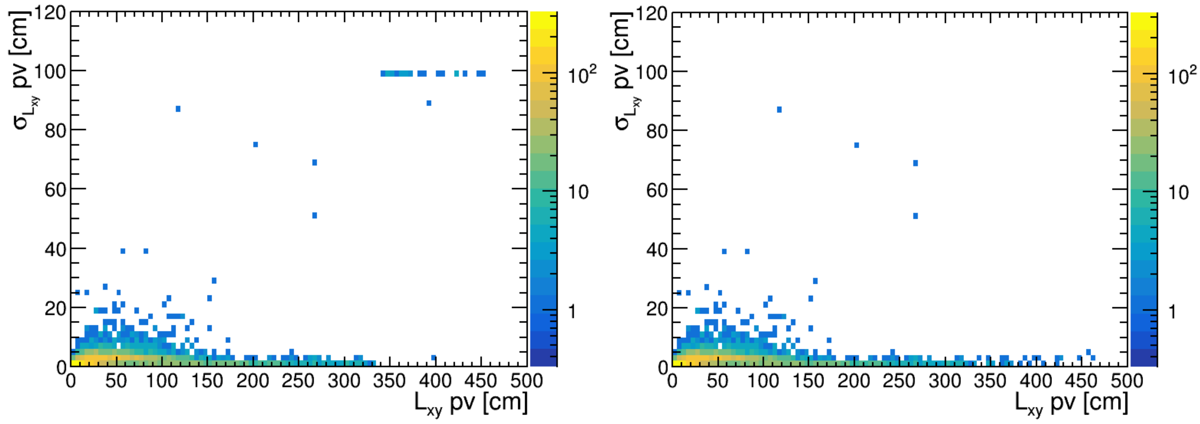

Figure 5.23: 2D distributions of $\sigma(L_{xy})$ vs. L_{xy} in $H \to XX \to 2\mu 2j$ signal events, showing fitted SVs beyond 330 cm reconstructed with an anomalous assignment of $\sigma(L_{xy}) = 100$ cm before the fix (left), in comparison to well-fitted SVs up to 500 cm after the fix (right).

in Figure 5.25. Up to four highest-p_T muons are selected among those passing all muon identification requirements (including $p_T > 10$ GeV). These muons can combinatorially yield up to 6 dimuons, after the DCA and vertex fit convergence requirements. The pairing criteria chooses up to 2 dimuons among the formed dimuons. In events with three muons, choosing a dimuon with the smallest vertex χ^2 yields the highest efficiency in selecting signal dimuons. In events with four or more muons, the highest efficiency is achieved by selecting the pair of dimuons with the smaller summed vertex χ^2 among all distinct pairs of dimuons that can be formed from the four highest p_T muons in the event.

The muon pairing algorithm is applied separately for each of the three dimuon categories, and only STA muons are considered for pairing in the STA-STA category. While up to 6 dimuons can be found in an event at this stage, in practice, the number of dimuons rarely exceeds 2 or 3. If more than two dimuons are found, only the two dimuons with the least vertex χ^2 are kept for further analysis, and the full selection criteria are applied to each of them independently. When choosing the best two dimuons in all categories combined, we ensure multiple dimuons do not include the same STA muons. To achieve this, we suppress dimuons that share a muon with another dimuon having a smaller vertex χ^2. Thus, the efficiency of picking up the right dimuon is close to 100% for the 2μ signal samples, and 85-95% for the 4μ signal samples.

5.5.6 Dimuon Vertex

Background from dimuons formed from unrelated muons are suppressed by requiring the common vertex fit χ^2 of the dimuons chosen by the pairing criteria to be below a certain threshold. Since χ^2 distributions in data are known to be difficult to reproduce accurately in simulation, the requirements on the vertex are deliberately chosen to be loose, leaving the signal efficiency almost intact. The optimal values of the thresholds

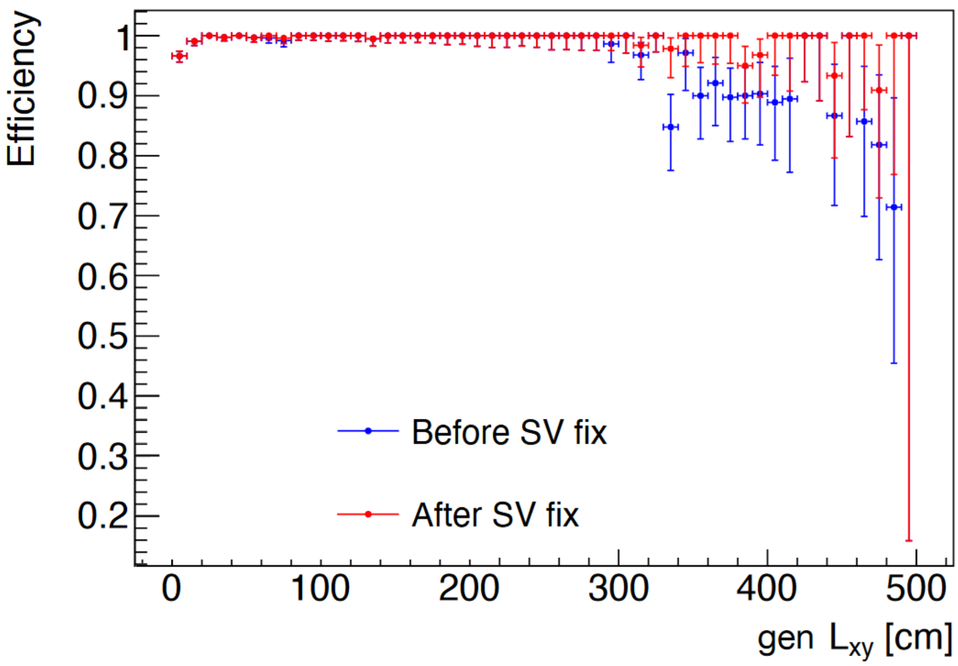

Figure 5.24: The common SV fit efficiencies as a function of gen L_{xy}, shown before (left) and after the fix (right), computed in $H \to XX \to 2\mu 2j$ signal events.

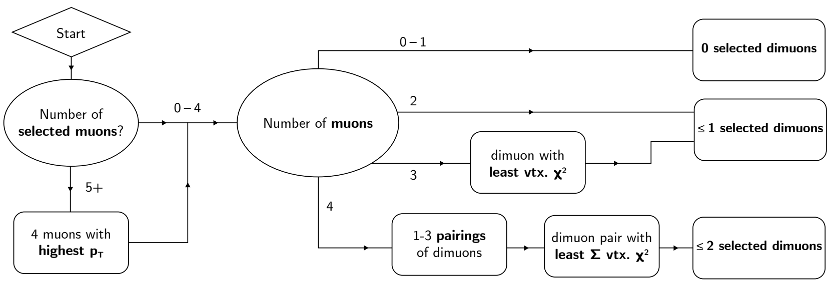

Figure 5.25: Flowchart depicted the muon pairing procedure. Reproduced from Ref. [96].

were chosen separately in each of the dimuon categories and were found to be $\chi^2_{\text{vertex}} < 10$ in the STA-STA category.

After extending the common vertex fit in the STA-STA category to $L_{xy} \sim 5$ m in Run 3, we observe STA-STA dimuons in the $|\Delta\Phi| > 3\pi/4$ control region with a large uncertainty in the transverse decay length $\sigma(L_{xy})$, shown in Figure 5.26. In order to suppress such dimuons, we add a new requirement, $\sigma(L_{xy}) < 20$ cm in this category. It is more than 99% efficient for signal events while removing most (10 out of 13) of the background events in this control region.

5.5.7 Rejection of cosmic-ray muons

Cosmic rays are high-energy particles originating from outer space, from highly energetic astrophysical events like supernovae. When these primary cosmic rays enter the Earth's atmosphere, they interact with the nuclei of atoms in the upper layers of the atmosphere,

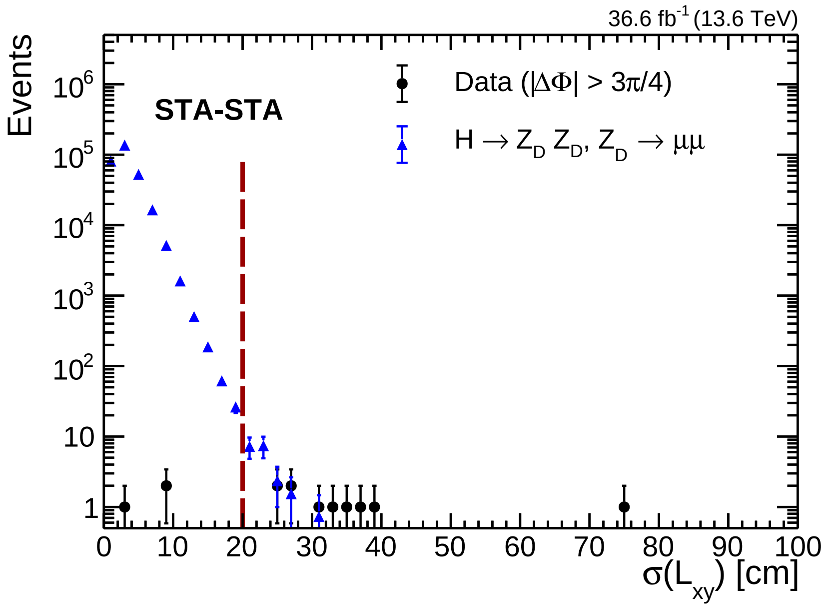

Figure 5.26: Distributions of $\sigma(L_{xy})$ for opposite-sign STA-STA dimuons with $|\Delta\Phi| > 3\pi/4$ in 2022 data, and in simulated HAHM signal events passing the full event selection, except for any requirement on $\sigma(L_{xy})$.

initiating a cascade of secondary particles, including pions, kaons, and photons. Charged pions and kaons further decay into muons and neutrinos. While other particles (except neutrinos) decay or are absorbed in the atmosphere, muons, due to their long lifetimes ($\tau(\mu) \approx 2.2\ \mu$s) and relatively weak interaction with matter, constitute a continuous flux of charged particles, towards the Earth's surface, and can even reach the CMS detector. A cosmic muon passing through the detector within the acceptance of the muon system is often reconstructed as two back-to-back muons, with one appearing in the upper half and the other in the lower half of the detector. Dimuons formed from such muons are usually displaced from the PV and may constitute an important source of background in this analysis, in particular in the STA-STA category.

Some of the cuts already described so far already help in reducing the cosmic-ray backgrounds. In particular, the timing and direction cuts described in Section 5.5.3 successfully reject most cosmic-ray muons due to two main factors: (a) at least one of the muon legs is typically early or late with respect to the nominal bunch-crossing time, and (b) one muon leg is typically facing in the 'outside-in' direction with respect to the center of the detector, in contrast to collision muons where both muons are reconstructed in the 'inside-out' direction. Further, the DCA requirement helps suppress dimuons that are formed from the two legs of different cosmic-ray muons in events with cosmic-ray-induced muon showers.

The remaining cosmic-ray background events are efficiently suppressed by requiring a threshold on the 3D opening angle between the two muons, $\alpha < 2.7$ rad (or equivalently, $\cos\alpha > -0.9$) in the STA-STA category. Figure 5.27 shows the distribution of $\cos\alpha$ for signal events passing the selection criteria in all HAHM samples combined and in data enriched in cosmic muons by requiring that the muons be out-of-time, in events that otherwise pass the selection criteria except for the threshold on $\cos\alpha$, demonstrating that the applied threshold effectively suppresses the vast majority of cosmic-ray muons.

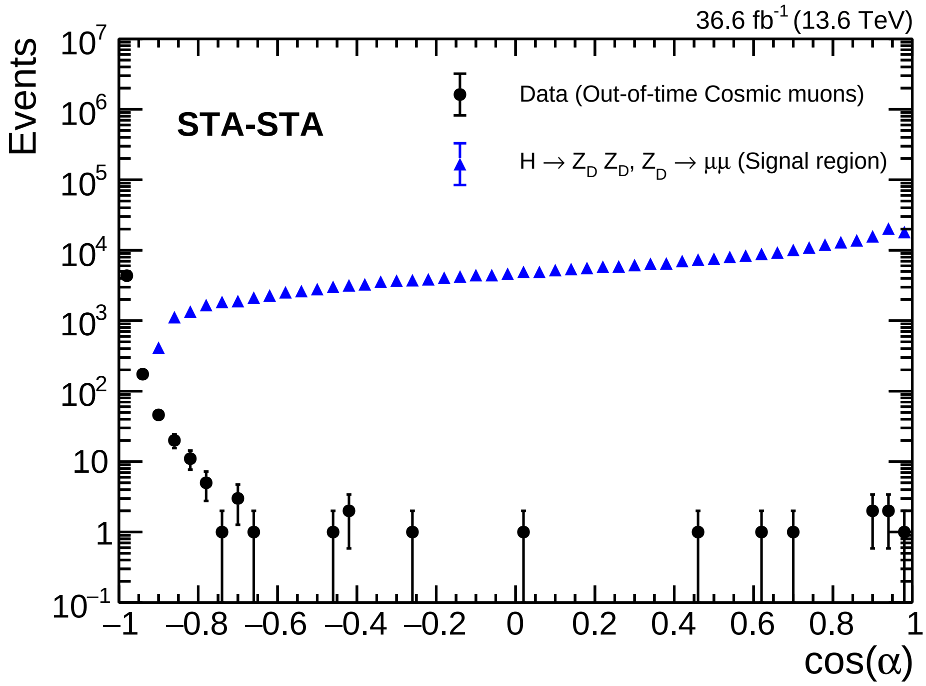

Figure 5.27: Distribution of $\cos\alpha$ in events passing the selection criteria in all HAHM signal events combined (blue) and in cosmic-enriched data consisting of out-of-time muons (black) in events that otherwise pass the selection criteria except for the $\cos\alpha > -0.9$ requirement.

This $\cos\alpha$ requirement does not, however, reject one particular class of cosmic muon background events – when a cosmic muon with a large incoming angle crosses the detector diagonally. A cosmic muon of this type can lead to multiple STA muons, each reconstructed from a small number (1 or 2) of segments. Such muons are poorly measured, and are not necessarily back-to-back to each other, and can mimic signal-like displaced dimuons. To suppress such events, we reject all STA-STA dimuons with $N(\text{dimuon segments}) < 5$, with $N(\text{dimuon segments})$ being the sum of the number of segments belonging to each STA muon in the dimuon. This requirement is inherited from the Run 2 search and is 99% efficient for signal events in the STA-STA category.

Another class of cosmic-ray background arises from events involving multiple cosmic-ray muons produced in an atmospheric shower. Such events typically produce a large number of nearly parallel STA muons that are reconstructed in both the upper and lower hemispheres of the detector. Although the $\cos\alpha$ requirement rejects dimuons formed from muons in different hemispheres, dimuons can also be formed from two muons in the same hemisphere. The common vertex fit of two nearly parallel muons in the same hemisphere should normally yield dimuons with $\cos\alpha \approx 1$ and a very small invariant mass, which the invariant mass requirements can suppress. However, common-vertex fit can behave anomalously when fitting two collinear muons with large impact parameters ($d_0 > 100$ cm); the fit changes the ϕ momentum direction of one of the muons by π, producing a dimuon with $\cos\alpha < 1$ and large dimuon invariant mass $m_{\mu\mu}$.

To reject such background events from cosmic muon showers, we apply the following set of requirements (inherited from the Run 2 analysis):

- The event must contain at least one reconstructed primary vertex with at least four associated tracks and with transverse (longitudinal) coordinates within 2 (24) cm of the nominal collision point. This standard 'good primary vertex' requirement rejects events containing only cosmic muons (as opposed to events where pp collisions and cosmic muon showers overlap) and has a negligible impact on the signal efficiency.

- The event must contain fewer than six pairs of parallel STA muons. All pairs of original STA muons (i.e., prior to the common vertex fit) with more than 12 valid CSC or DT hits and with $p_T > 5$ GeV are considered, and the number of pairs with $|\cos\alpha| > 0.99$ is counted, as N(parallel pairs). Most cosmic muon showers have N(parallel pairs) ≥ 6, whereas this criterion is true only for 0.2% of the 2μ and 0.4% of the 4μ signal events.

- Neither of the two muons forming an STA-STA dimuon is back-to-back ($\cos\alpha < -0.9$) to a third STA muon with $p_T > 10$ GeV. For a dimuon to be vetoed, the time difference between the muon in the dimuon and its back-to-back muon, Δt_{b2b}, should be larger than 20 ns. This requirement is only applied in the STA-STA dimuon category, in which the contribution from the cosmic muon background is the largest, and is 98-99% efficient for signal events. In addition to rejecting dimuons from cosmic muon showers, it also helps to suppress dimuon candidates reconstructed from one 'leg' of a cosmic muon, and a muon from an overlapping pp collision.

The above set of requirements suppresses cosmic-ray muons efficiently and renders this background negligible. The residual contamination from cosmic-ray muons in the analysis signal region is estimated in Section 5.6.4 and found to be negligible.

5.5.8 Suppression of pp collision backgrounds

SM events produced in pp collisions are the main source of background in the analysis. Four major types of SM backgrounds are relevant for the analysis of STA-STA dimuons.

The first type of background events are genuinely displaced dimuons. Such dimuons can be produced in decays of non-prompt low-mass resonances such as J/ψ. They can also be formed from the products of B-meson 'cascade' decays not yielding dimuon resonances, such as $b \to c\mu_1 X_1; c \to \mu_2 X_2$. The mass of these dimuons does not exceed a few GeV, and thus, we reject them by requiring that the reconstructed dimuon invariant mass $m_{\mu\mu}$ be larger than 10 GeV. This requirement also rejects dimuons from $\Upsilon(1S)$ decays and low-mass Drell-Yan.

A second type of SM background consists of prompt dimuons that appear displaced due to the effect of instrumental or reconstruction mistakes. Such dimuons mostly arise from Drell-Yan $\mu\mu$ events, with relatively small contributions from processes such as $t\bar{t}$ and diboson production. Drell-Yan $\tau\tau$ events in which both leptons decay to muons result in a background with characteristics similar to those with mismeasured Drell-Yan $\mu\mu$ events. The main discriminating variables between signal and such backgrounds are discussed in Section 5.1. Figure 5.28 shows the distributions of the transverse collinearity angle $|\Delta\Phi|$ in STA-STA dimuons for events in the HAHM and RPV SUSY signal models passing the full selection except for the $|\Delta\Phi|$ requirement, compared to a sample of STA*-STA* dimuons enriched in prompt Drell-Yan events by requiring them to have $L_{xy}/\sigma(L_{xy}) < 1$ and be associated to a TMS-TMS dimuon with an invariant mass around the Z boson mass. The distributions of $|\Delta\Phi|$ for dimuons in the signal show a peak at zero, whereas dimuons originating from mismeasured prompt events, in particular Drell-Yan, are expected to have a $|\Delta\Phi|$ distribution symmetric about $\pi/2$, as the dimuon momentum vector $\vec{p}_{T,\mu\mu}$ is uncorrelated with the $\vec{L}_{xy}$ vector.

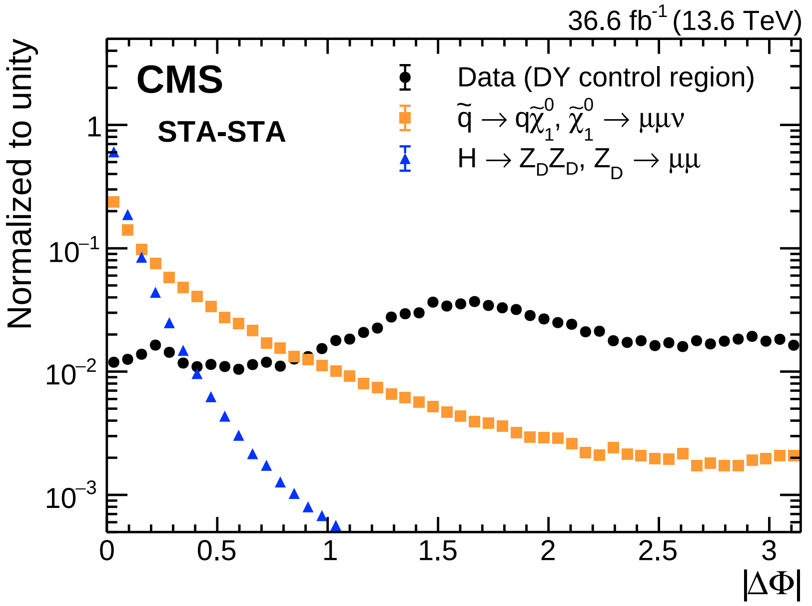

Figure 5.28: Distribution of $|\Delta\Phi|$ for STA-STA dimuons in a data sample enriched in Drell-Yan events (black circles), and for events passing all selection criteria, except for the requirement on $|\Delta\Phi|$, in all HAHM (blue triangles) and RPV SUSY (orange squares) generated signal samples combined. All distributions are normalized to unity. Reproduced from Ref. [23].

The $|\Delta\Phi|$ distribution peaks at zero in both the HAHM model as well as the RPV SUSY model. Therefore, we can define a threshold, $|\Delta\Phi|_C$, and the signal region to be $|\Delta\Phi| < |\Delta\Phi|_C$. However, the $|\Delta\Phi|$ distributions differ in the two types of signal models and, therefore, need two different signal regions in terms of the $|\Delta\Phi|$ selection.

For the HAHM model, the $|\Delta\Phi|$ distribution strongly peaks at zero, with the tail of the distribution determined by the resolution. Thus, we use $|\Delta\Phi|_C = \pi/10$ to define the

HAHM signal region in the STA-STA category. For the RPV SUSY model, the distribution peak is broader due to the undetected neutrino among the decay products, and so we use a looser requirement of $|\Delta\Phi|_C = \pi/4$. Figure 5.29 shows the efficiency for various $|\Delta\Phi|_C$ thresholds in the HAHM model and the RPV SUSY model. For the chosen $|\Delta\Phi|_C$, the signal efficiency ranges from 90-99% in the HAHM model and from 60-99% in the RPV SUSY model.

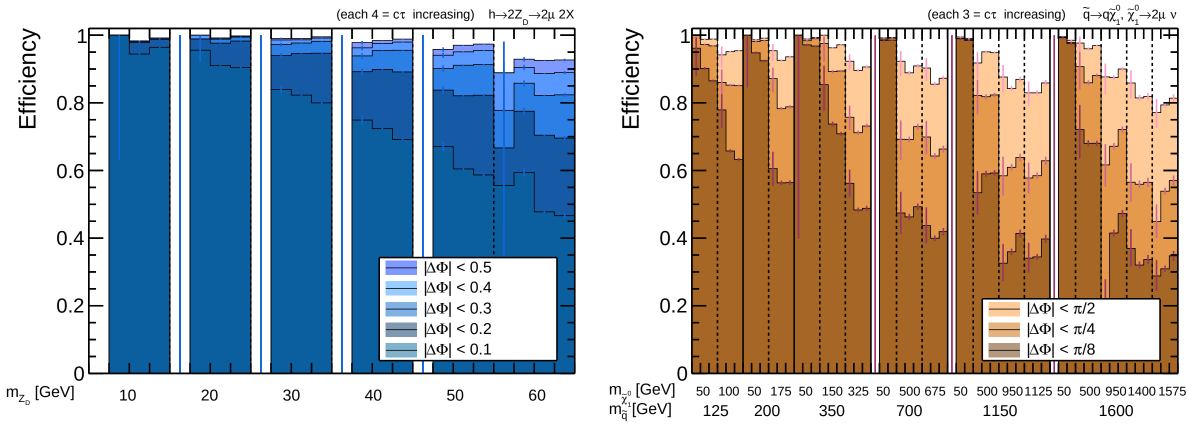

Figure 5.29: Cut efficiency in HAHM (left) and RPV SUSY (right) signal samples for various thresholds of $|\Delta\Phi|$ in the STA-STA category. The dimuons are required to pass the full selection criteria except for the $|\Delta\Phi|$ requirements. The $c\tau$ values for each $m(Z_D)$ and $(m(\tilde{q}), m(\tilde{\chi}_1^0))$ are listed in Tables 4.2 and 4.4, and are shown here in increasing order of magnitude.

Given the $|\Delta\Phi|$-symmetry of the Drell-Yan and other prompt backgrounds, we use regions symmetric about the $|\Delta\Phi| = \pi/2$ value, namely $|\Delta\Phi| > \pi - |\Delta\Phi|_C$, as control regions for the corresponding background predictions. We use central $|\Delta\Phi|$ regions, $\pi/4 < |\Delta\Phi| < \pi/2$ and $\pi/2 < |\Delta\Phi| < 3\pi/4$ for validating the background predictions, as described in Section 5.5.8.

Another important class of pp collision backgrounds is composed of QCD dijet and multijet events yielding dimuons that are formed from genuine or fake muons belonging to different jets. Muon emerging from QCD events such as $J/\psi \to \mu\mu$ may be displaced, and the $|\Delta\Phi|$ distribution in such events may show a signal-like peak at zero.

Figure 5.30 shows the $|\Delta\Phi|$ distributions in 2022 data in control regions enriched in DY-like events (by requiring that the STA*-STA* dimuon is associated to a TMS-TMS dimuon with $L_{xy}/\sigma(L_{xy}) < 1$, an invariant mass $m_{\mu\mu}$ between 80 and 100 GeV, and that both TMS muons are isolated), and in QCD-like events (by requiring that the STA-STA dimuon invariant mass < 10 GeV).

Dimuons in QCD events are expected to be charge-symmetric; in addition to splitting according to $|\Delta\Phi|$, all selected dimuons are classified into OS and SS dimuons, depending on whether two muons forming a dimuon have the same or opposite electric charges. The signal selection requires that dimuons be formed from OS muons, while SS dimuons are

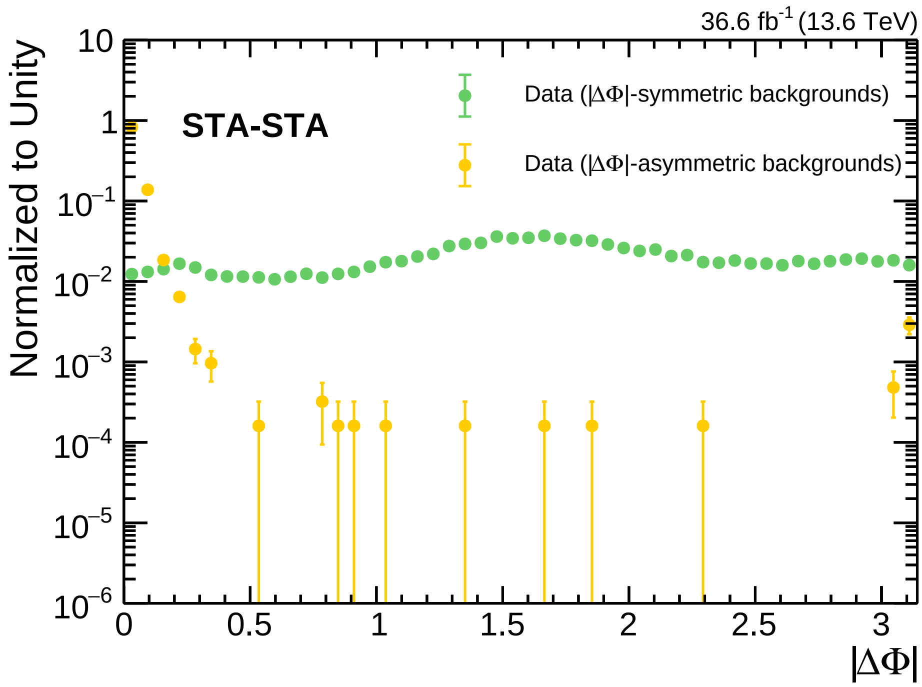

Figure 5.30: Distributions of $|\Delta\Phi|$ in events enriched in DY-like $|\Delta\Phi|$-symmetric backgrounds (left) and QCD-like $|\Delta\Phi|$-symmetric backgrounds (right) that show a signal-like peak at zero. The distributions are normalized to unity to facilitate comparison.

used to define a control region in which we evaluate backgrounds arising from QCD events and B-meson cascade decays.

The dimuons are required to be displaced with respect to the primary vertex by requiring a minimum $L_{xy}/\sigma(L_{xy}) > 6$. The thresholds are optimized separately for each of the dimuon categories by maximizing the expected statistical significance of discovery between the expected number of signal events obtained from simulation and the number of background events estimated from events in control regions in data. This optimization was performed on all signal samples, using a systematic grid of thresholds on several key discriminating variables. We found that a requirement of $L_{xy}/\sigma(L_{xy}) > 6$ maximized sensitivity to the signal in the STA-STA dimuon category. In the STA-STA category, the significance does not increase appreciably for any requirement on the $d_0/\sigma(d_0)$ once a requirement on $L_{xy}/\sigma(L_{xy})$ is imposed, and therefore no cuts are applied on the $d_0/\sigma(d_0)$ in this category. Events rejected by the $L_{xy}/\sigma(L_{xy})$ requirement are used to assess the accuracy of the background prediction and assign the corresponding systematic uncertainty, as discussed in Section 5.6.5.

5.5.9 Suppression of mismeasured low-mass dimuons in the STA-STA category

The looser trigger requirements in Run 3 as compared to the triggers in Run 2, in particular the lower p_T thresholds and the absence of a minimum dimuon invariant mass requirement,

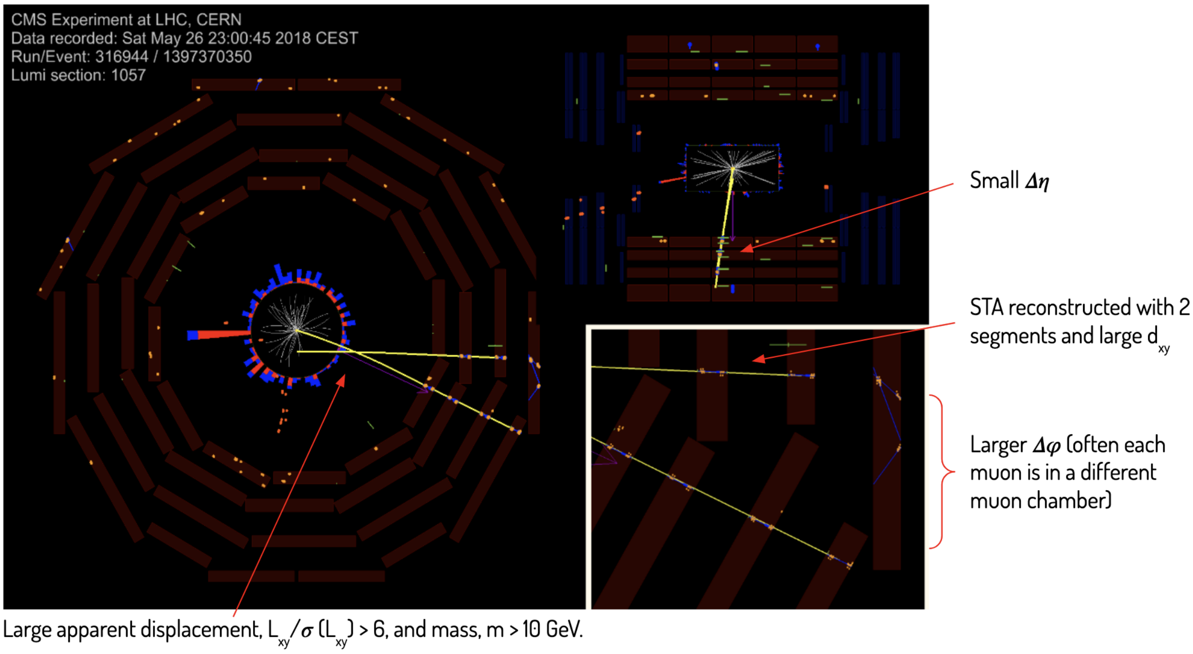

Figure 5.31: Event display of a typical SS event with $L_{xy}/\sigma(L_{xy}) > 6$ in 2018 data. Two STA muons corresponding to two distinct muons produced in a pp collision have very similar η ($R - z$ view, top right), but are well separated in ϕ ($x - y$ view, left and bottom right). The nearly horizontal muon, giving rise to a displaced vertex, contains only two muon segments.

lead to an increase in the number of expected background events at low mass in the STA-STA category before applying the offline analysis selection. This was observed when comparing the number of events with SS STA-STA dimuons passing the full event selection in Run 2 data and in Run 3 data. Such events have some common features:

- Small angular separation between the two muons, $\Delta\eta_{\mu\mu} = |\eta_{\mu_1} - \eta_{\mu_2}| < 0.1$, along with a large opening angle in the azimuthal plane, $\Delta\phi_{\mu\mu} = |\phi_{\mu_1} - \phi_{\mu_2}| > 0.1$.
- One or both STA muons that form a dimuon contain no more than two or three muon segments and a modest number of muon hits. The event display of one such event is shown in Figure 5.31.

Such events are expected to be actually prompt but are misreconstructed as displaced. We can understand events with similar properties in SS STA*-STA* dimuons by looking at the characteristics of the associated TMS-TMS dimuons (as motivated in Section 5.5.2). A study of such events showed that a vast majority of the dimuons reconstructed in the tracker are low-mass ($m_{\mu\mu} < 10$ GeV) dimuons with both $\Delta\eta_{\mu\mu}$ and $\Delta\phi_{\mu\mu}$ smaller than 0.1. These muons are separated due to the magnetic field and appear with a larger $\Delta\phi_{\mu\mu}$ in the muon system. When one or both of them are reconstructed from a small number of measurements in the muon system alone, they can appear as muons with higher p_T, with straighter tracks, leading to an overestimated $m_{\mu\mu}$ above the 10 GeV threshold, and a fake displaced vertex.

To suppress the low-mass background events in the STA-STA dimuon category, we reject all STA-STA dimuons with $\Delta\eta_{\mu\mu} < 0.1$ that have N(dimuon segments) ≤ 5, or barrel-only muons with N(DT hits) ≤ 24. These requirements are 90-99% efficient in signal, depending on the sample. The loss in the signal efficiency is the largest ($\sim 10\%$) for the samples with boosted LLPs, which have the largest fraction of events with $\Delta\eta_{\mu\mu} < 0.1$; for example, $m(H) = 1000$ GeV, $m(X) = 20$ GeV. These requirements reject 22 out of 28 SS events in 2022 data that pass all other selection criteria.

Events rejected by these requirements are used to assess the systematic uncertainty associated with the QCD background evaluation in the STA-STA dimuon category, as described in Section 5.6.5.

As discussed in Section 5.5.4, we require both STA muons to be isolated from hadronic activity in the tracker. This requirement helps suppress this type of background further, since such dimuons often appear close to the hadronic activity in the detector, and helps to compensate for the increased background of this type due to the lowered HLT p_T thresholds in the Run 3 triggers.

5.5.10 Invariant Mass

To test for the existence of LLPs with a given mass in the dark photon and the heavy scalar model, dimuons satisfying the selection criteria are further required to have the invariant mass consistent with the LLP mass. The dimuon invariant mass has to be within a mass interval centered at or around the LLP mass being considered. Therefore, the signal region is subdivided into overlapping mass intervals covering the complete dimuon invariant mass spectrum starting at $m_{\mu\mu} > 10$ GeV, listed in Table 5.2. The boundaries of each mass interval are chosen according to the mass resolution and the expected background, such that the interval contains more than 90-99% of selected events in a signal sample corresponding to the probed LLP mass. For the HAHM model, this choice typically corresponds to the interval width of $\pm$6-8 times the mass resolution (which is about 10-25% for LLP masses between 20 and 350 GeV in the STA-STA category). Since the invariant mass resolution only weakly depends on the LLP lifetime (and $m(H)$ in the heavy scalar model), we use identical intervals for all $c\tau(X)$ and $m(H)$.

In the RPV SUSY model, the presence of a neutrino in the decay vertex, $\tilde{\chi}_1^0 \to \mu\mu\nu$, leads to a non-peaking dimuon invariant mass distribution that falls steeply at $m(\tilde{\chi}_1^0)$. We improve mass resolution and the signal-to-background discrimination in signals featuring a secondary vertex made of a dimuon accompanied by neutral or unobserved particles by using the **'corrected secondary-vertex mass'**, $m_{\mu\mu}^{\text{corr}}$, defined as:

$$m_{\mu\mu}^{\text{corr}} = \sqrt{m_{\mu\mu}^2 + p_{\mu\mu}^2 \sin^2\theta} + p_{\mu\mu}\sin\theta \qquad (5.2)$$

where $p_{\mu\mu}$ is the magnitude of the dimuon three-momentum vector $\vec{p}_{\mu\mu}$, and θ is the 3D angle between $\vec{p}_{\mu\mu}$ and the vector connecting the primary and secondary vertex.

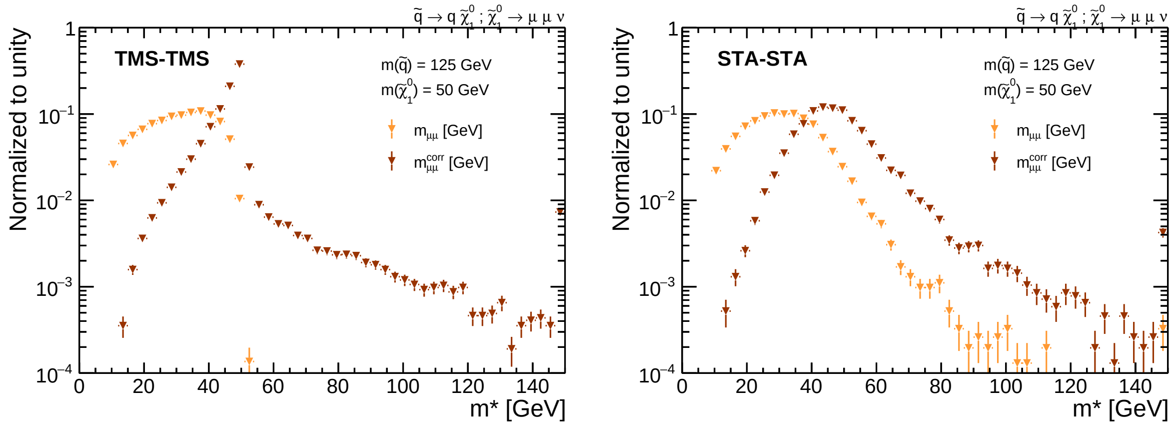

Figure 5.32: Distributions of $m_{\mu\mu}$ and $m_{\mu\mu}^{\text{corr}}$ in the TMS-TMS (left) and STA-STA (right) dimuon categories in the $\tilde{q} \rightarrow q\tilde{\chi}_1^0$ signal samples with $m(\tilde{q}) = 125$ GeV and $m(\tilde{\chi}_1^0) = 50$ GeV, demonstrating the improvement in the location and resolution of the mass peak, upon correcting for the missing neutrino momentum.

This variable was used at LEP and SLC, and recently by the LHCb [108] and CMS [109] collaborations. The motivation behind this variable is to rely on the conservation of momentum and correct for the missing mass of the vertex arising from unaccounted particles. In the case where only a dimuon emerges from the secondary vertex (without any missing products and ignoring the finite detector resolution), the angle θ is zero and $m_{\mu\mu}^{\text{corr}} = m_{\mu\mu}$. In decays where the vertex includes other particles, $\theta > 0$, $m_{\mu\mu}^{\text{corr}} > m_{\mu\mu}$, and the corrected mass is the minimum mass that the LLP, in this case the neutralino $\tilde{\chi}_1^0$, can have while being consistent with the measured direction of flight of the LLP.

Figure 5.32 shows the $m_{\mu\mu}$ and $m_{\mu\mu}^{\text{corr}}$ distributions in a signal with $m(\tilde{q}) = 125$ GeV and $m(\tilde{\chi}_1^0) = 50$ GeV in the TMS-TMS (left) and STA-STA (right) categories. A clear peak at $m(\tilde{\chi}_1^0)$ in the $m_{\mu\mu}^{\text{corr}}$ distribution is visible in the TMS-TMS category after the correction is applied due to the improved resolution. In the STA-STA categories, the peak is not as sharp due to the mass resolution not being as precise, nevertheless significantly improving the discrimination between signal and background. Due to this improvement in mass resolution, we are able to define $m_{\mu\mu}^{\text{corr}}$ intervals to probe different neutralino masses. Here, as well, the width of the intervals is chosen to contain a large fraction of the RPV SUSY signal (typically 80-90%). We additionally require $m_{\mu\mu} > 15$ GeV in the STA-STA category to avoid the leakage of small $m_{\mu\mu}$ background events into large $m_{\mu\mu}^{\text{corr}}$ windows.

For the RPV SUSY model, we define $m_{\mu\mu}^{\text{corr}}$ intervals instead. A study of the optimal $m_{\mu\mu}^{\text{corr}}$ intervals for various $m(\tilde{q})$, $m(\tilde{\chi}_1^0)$ and $c\tau(\tilde{\chi}_1^0)$ has been performed using RPV SUSY signal simulation, optimized using two control regions in background enriched data. Since the background events cluster at low $m_{\mu\mu}^{\text{corr}}$ values, only a lower mass threshold corresponding to approximately 40-50% of the probed LLP mass was found to be useful. We additionally remove events with $10 < m_{\mu\mu} < 15$ GeV to avoid the leakage of background events from small $m_{\mu\mu}$ to large $m_{\mu\mu}^{\text{corr}}$. Since no strong dependence on $m(\tilde{q})$ on $c\tau(\tilde{\chi}_1^0)$ was observed, the chosen value of the $m_{\mu\mu}^{\text{corr}}$ threshold depends only on $m(\tilde{\chi}_1^0)$. The correspondence between the $m_{\mu\mu}$ and $m_{\mu\mu}^{\text{corr}}$ intervals and the probed LLP mass is summarized in Table 5.2.

Mass interval [GeV]	$(m(H), m(X))$ [GeV]	$m(Z_D)$ [GeV]
$10 < m_{\mu\mu} < 15$	-	10
$15 < m_{\mu\mu} < 32$	(125, 20), (200, 20), (400, 20), (1000, 20)	20
$20 < m_{\mu\mu} < 60$	-	30
$25 < m_{\mu\mu} < 80$	-	40
$30 < m_{\mu\mu} < 80$	(125,50), (200, 50), (400, 50), (1000, 50)	50
$40 < m_{\mu\mu} < 120$	-	60
$40 < m_{\mu\mu} < 245$	(400, 150), (100, 150)	-
$60 < m_{\mu\mu}$	(1000, 350)	-
Corrected mass interval [GeV]	$(m(\tilde{q}))$ [GeV]	$m(\tilde{\chi}_1^0)$ [GeV]
$m_{\mu\mu}^{\text{corr}} > 25$	125, 200, 350, 700, 1150, 1600	50
$m_{\mu\mu}^{\text{corr}} > 50$	125	100
$m_{\mu\mu}^{\text{corr}} > 80$	350	150
$m_{\mu\mu}^{\text{corr}} > 90$	200	175
$m_{\mu\mu}^{\text{corr}} > 90$	350	325
$m_{\mu\mu}^{\text{corr}} > 210$	700, 1150, 1600	500
$m_{\mu\mu}^{\text{corr}} > 300$	700	675
$m_{\mu\mu}^{\text{corr}} > 300$	1150, 1600	950
$m_{\mu\mu}^{\text{corr}} > 300$	1150	1125
$m_{\mu\mu}^{\text{corr}} > 300$	1600	1400
$m_{\mu\mu}^{\text{corr}} > 300$	1600	1575

Table 5.2: The mass intervals in $m_{\mu\mu}$ and $m_{\mu\mu}^{\text{corr}}$ in the STA-STA category associated to the parameters of the simulated signal samples.

5.5.11 Summary of Event Selection

The event and object selection criteria used in the analysis are shown in Table 5.3. The updated selections in Run 3 search [23] are highlighted in blue. Only the selection criteria for the STA-STA category was motivated and detailed in the previous section, as it is the focus of this thesis; however, the table also includes the selection criteria for TMS muons and the TMS-TMS category, as described in Section 4.3 of the Run 2 search [22], for completeness.

Figure 5.33 depicts the signal efficiency in the heavy scalar model for the 2022 selections in each of the dimuon categories. The efficiency in each dimuon category is computed as the ratio of the number of dimuons of a given type that passed the full set of the selection criteria (including invariant mass requirements and trigger) to the total number of generated events. As expected, the vast majority of the signal dimuons in the samples with the longest (shortest) lifetimes are found in the STA-STA (TMS-TMS) category. The overall signal efficiency (including acceptance and trigger effects) varies from about 5%

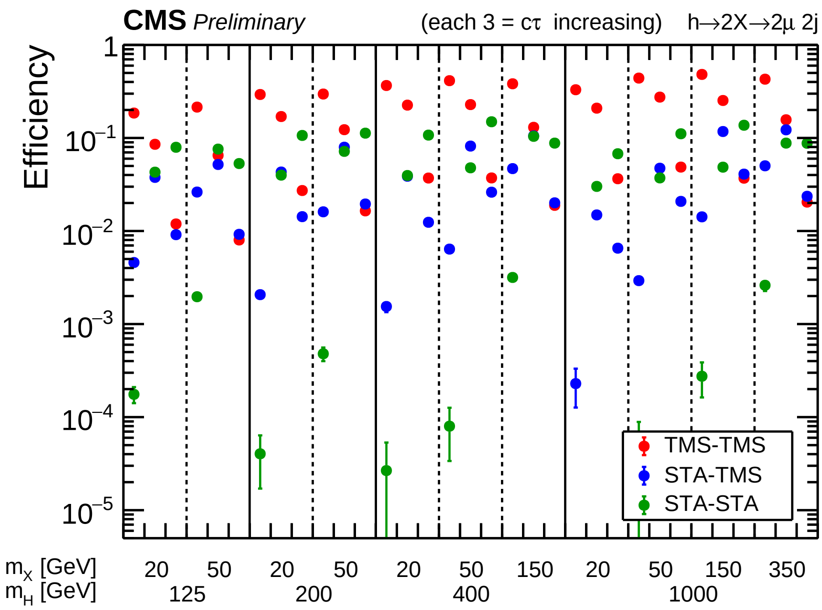

Figure 5.33: Signal Efficiency in the TMS-TMS (red), STA-TMS (blue) and STA-STA (green) dimuon categories in 2022 $H \to 2X \to 2\mu 2j$ signal samples. The efficiency is defined as the fraction of signal events that fired the trigger and passed the full set of the selection criteria. The $c\tau$ values (not labeled) span three orders of magnitude for each combination of $(m(H), m(X))$ are listed in Table 4.3, and are shown in increasing order (from shortest to longest). The signal efficiency for STA-TMS dimuons, included for posterity, is computed by applying the Run 2 analysis selection criteria. The uncertainties shown are statistical.

for the samples with small H and X masses, and long X lifetimes to about 50% for the samples with large masses and short lifetimes.

Figure 5.34 shows the full event selection efficiency as a function of $c\tau$ in the HAHM model, and illustrates the relative improvements in the efficiencies from the new trigger algorithms in the ratio pad at the bottom of each plot. The efficiencies in this figure are *event* efficiencies, defined as the number of events in which at least one dimuon belonging to a given dimuon category passes the full set of selection criteria, divided by the number of generated events with at least one Z_D decaying to a muon pair. The signal samples are generated at 4 $c\tau$ points, and are therefore reweighted for each lifetime point in the graph and combined into one sample, according to the procedure described in Section 4.4.6.2 of Ref. [107]. A large increase in the efficiency is observed in 2022 relative to 2018, as illustrated in the ratio pad, and is apparent for all masses. Comparing this with Figure 5.10, we are able to retain most of the improvement in signal efficiency from the triggers even after the event selection. These efficiencies are used in the statistical interpretation of the results of the search, described in Section 5.8.

Figure 5.35 shows a representation of the offline selection efficiencies in all HAHM signal samples as a function of the generated dimuon vertex position in terms of the gen L_{xy} and the gen $|z|$, overlaid over an approximate schematic of the detector geometry, demonstrating that we retain high efficiencies to reconstruct displaced dimuon vertices up to 500 cm in

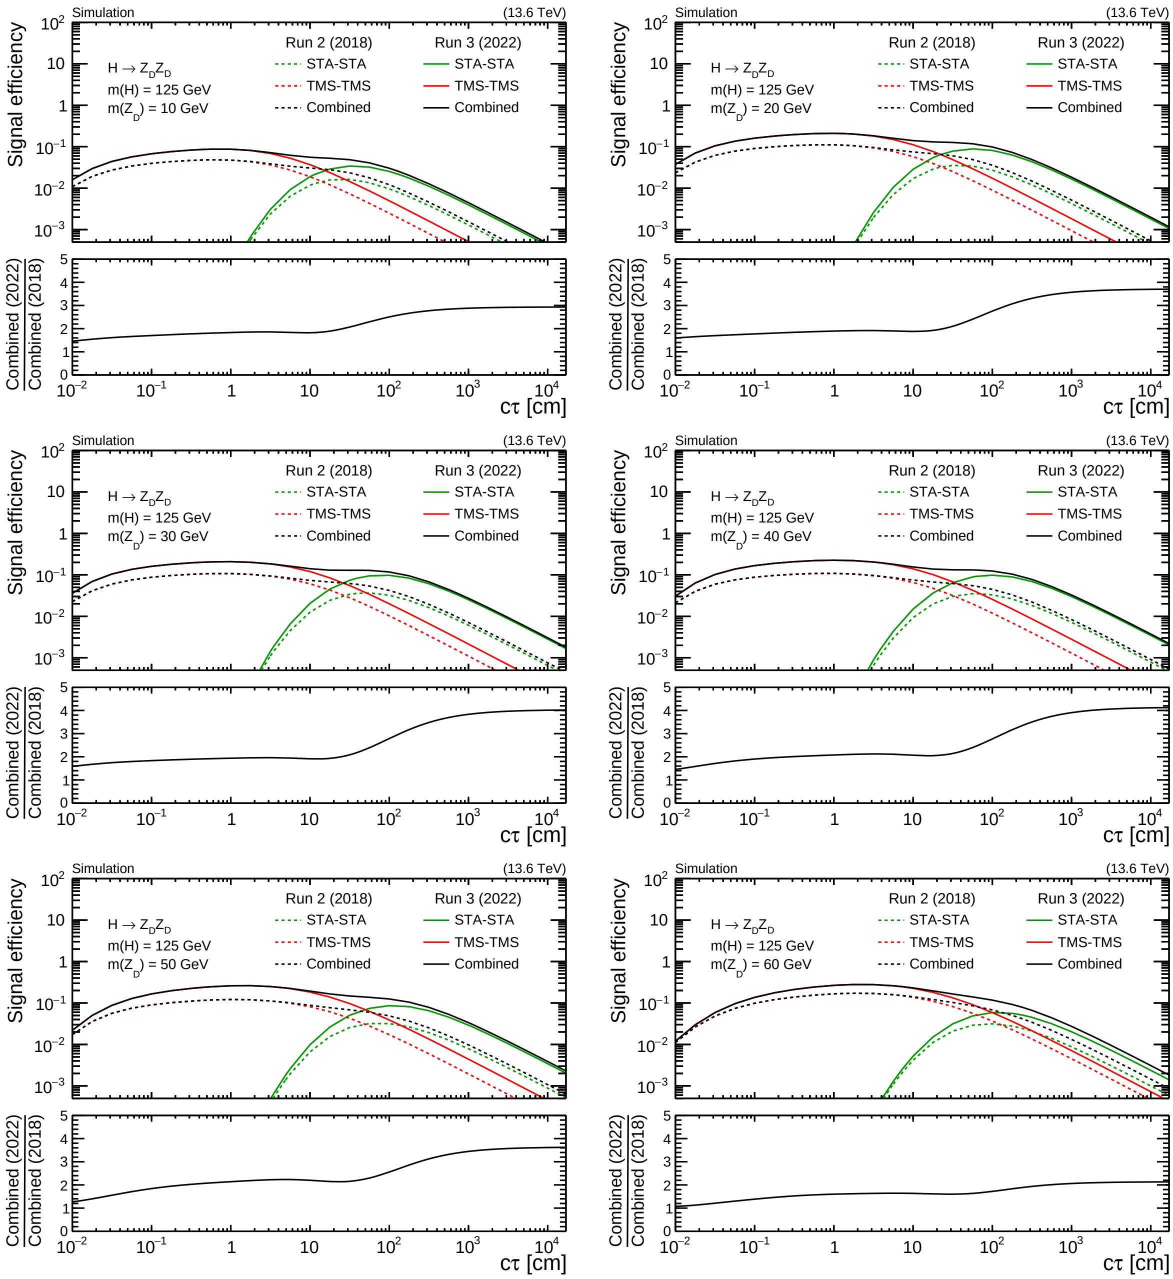

Figure 5.34: Efficiency in the STA-STA (green) and TMS-TMS (red) dimuon categories and their combination (black) for the HAHM signal events with $m(Z_D)$ ranging from 10 GeV (top left) to 60 GeV (bottom right). Solid curves show efficiencies achieved with the OR of all 2022 triggers, whereas dashed curves show efficiencies for the subset of events selected by the triggers used in the 'Run 2 (2018)' configuration. The efficiency is defined as the fraction of signal events that satisfy the criteria of the indicated trigger as well the full set of the offline selection criteria listed in Table 5.3. The ratio plot in the lower panels shows the relative improvement of the overall signal efficiency brought in by the improvements in the trigger.

Event Selection

Requirement	Dimuon Category	
	STA-STA	TMS-TMS
Good primary vertex	yes	yes
HLT-RECO matching	yes	yes
N(parallel pairs)	< 6	< 6

Muon Selection

Requirement	Muon Type	
	STA	TMS
STA to TMS muon association	not matched to TMS μ	matched to STA μ
N(CSC + DT hits)	> 12	> 12 (assoc. STA muon)
N(DT hits) for muon in barrel	> 18	-
Tracker muon	-	yes
N(matched muon segments)	-	> 1
p_T	> 10 GeV	> 10 GeV
$\sigma(p_T)/p_T$	< 1.0	< 1.0
$\chi^2_{\text{track}}/\text{dof}$	< 2.5	-
I^{rel}	< 0.15	< 0.075
$\lvert t_{\text{in-out}}\rvert$	< 12 ns	-
Muon `direction`	+1 ('inside-out')	-
$d_0/\sigma(d_0)$	-	> 6

Dimuon Selection

Requirement	Dimuon Category	
	STA-STA	TMS-TMS
DCA	< 15 cm	< 15 cm
Convergent common vertex fix	yes	yes
Pairing criteria	Select best 1-2 dimuons ranked by $\min(\Sigma\chi^2_{\text{vertex}}/\text{dof})$	
χ^2_{vertex}	< 10	< 10
ΔN(pixel hits)	-	< 3
N(hits before vertex)	-	Not applied
$\sigma(L_{xy})$	< 20 cm	-
$\cos\alpha$	$>$ -0.9	- 0.99
Back-to-back muon timing, $\lvert\Delta t_{\text{b2b}}\rvert$	> 20 ns	-
N(dimuon segments)	> 4	-
If $\lvert\Delta\eta_{\mu\mu}\rvert < 0.1$		
·N(dimuon segments)	> 5	-
·N(DT hits) for muons in barrel	> 24	-
$m_{\mu\mu}$	> 10 GeV	> 10 GeV
Leading muon p_T	-	> 25 GeV
$L_{xy}/\sigma(L_{xy})$	> 6	> 6
Muon sign correlation	opposite-sign (OS)	
$\lvert\Delta\Phi\rvert$ for $H \to Z_D Z_D$ (HAHM)*	$< \pi/10$	$< \pi/30$
$\lvert\Delta\Phi\rvert$ for $\tilde{q} \to q\tilde{\chi}^0_1$ (RPV SUSY)	$< \pi/4$	$< \pi/4$
$m_{\mu\mu}, m^{\text{corr}}_{\mu\mu}$ intervals	See Section 5.8	

Table 5.3: Summary of the selection criteria used in the analysis, grouped into event, muon, and dimuon requirements. The updated selections in the Run 3 search [22] are highlighted in blue.

L_{xy} and about 600 cm in $|z|$. While it is possible to reconstruct dimuons beyond this point, the muons are reconstructed with lower quality and have poorer resolution due to a smaller number of hits in the muon system. A similar map is provided for TMS-TMS dimuons, overlaid on top of the CMS tracker geometry, in Figure B.3 in Appendix B.

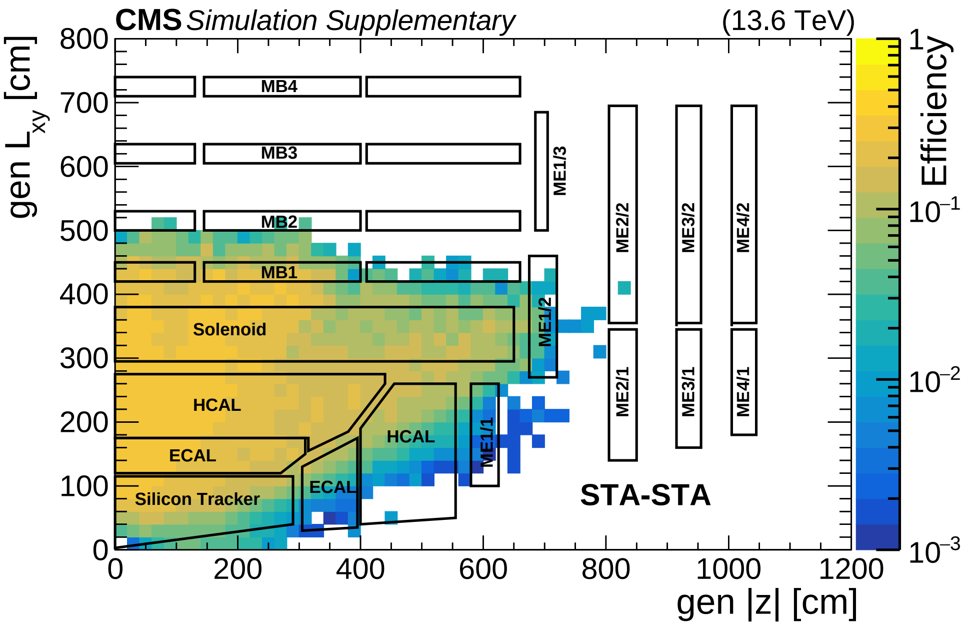

Figure 5.35: Offline selection efficiency of STA-STA dimuons as a function of the simulated L_{xy} and $|z|$ decay positions of the dark photon Z_D, in all generated HAHM signal samples combined. Approximate locations of various CMS subdetectors are drawn superimposed as black boxes. MB1-MB4 denote the four DT stations in the barrel, while the ME stations show the CSC modules, with the first index denoting the CSC stations and the second index denoting the rings in the endcap.

5.6 Backgrounds in the STA-STA Category

Background yields in CMS analyses are generally estimated via MC simulation and data driven estimates. In this analysis, we cannot rely on simulated samples for an accurate evaluation of background yields due to a variety of reasons, the most important of which are mentioned here:

- The backgrounds from pp collisions, mainly from Drell-Yan and QCD processes, are mostly instrumental due to misreconstruction. MC simulations may not faithfully reproduce such mistakes and their rate.
- The MC samples have limited statistical power, as the number of generated events is much smaller than the number of events expected in data.

- The available MC simulation of cosmic muons does not include the simulation of multiple cosmic muons from an atmospheric shower or of cosmic muons occurring with (and therefore overlapping with) pp collision data taking.

For these reasons, we use a data-driven approach to evaluate the expected background. The control regions used to evaluate various types of background processes are defined by inverting one or more selection criteria, giving regions mostly containing events of a given type of background while only containing a negligible signal contribution. The control regions and the details of the background evaluation procedure are defined differently for each dimuon category. The procedure for the STA-STA category is described in the rest of this section. The events passing the full selection (i.e., events in the signal region) were blinded until the last step of the analysis to avoid potential bias in the event selection.

The methods described in this section closely follow the strategy employed in the Run 2 search. The evaluation of the QCD background in the STA-STA category was updated and improved to take advantage of the STA muon isolation in the Run 3 analysis. In general terms, we expect background in two broad categories, as illustrated in Figure 5.30:

- $|\Delta\Phi|$-symmetric, such as Drell-Yan.
- $|\Delta\Phi|$-asymmetric, such as from QCD processes.

5.6.1 Estimation of backgrounds with ‘symmetric’ collinearity

As discussed in Section 5.5.8, the signal processes have a $|\Delta\Phi|$ distribution that is significantly different from what is expected in many types of background events. In the signal, both the LLP transverse momentum vector (approximated by the dimuon $\vec{p}_T$ vector) and the $\vec{L}_{xy}$ vector point away from the primary vertex, leading to a $\Delta\Phi$ distribution that strongly peaks at zero. In background events with no genuine secondary vertex (such as in Drell-Yan events), or with an apparent secondary vertex not related to the particle decay (such as in ‘split’ cosmic muons), the $\vec{p}_{T,\mu\mu}$ and the $\vec{L}_{xy}$ vectors are uncorrelated, leading to a $|\Delta\Phi|$ distribution that is expected to be symmetric about $\pi/2$. We refer to such $|\Delta\Phi|$-symmetric backgrounds as Drell-Yan (DY) type backgrounds.

To understand the shape of the $|\Delta\Phi|$ distribution in background events such as Drell-Yan, we use the events in the control region obtained by reversing the STA to TMS muon association for both muons. Such events consist of STA-STA dimuons passing all selection criteria, but in which each of the two constituent STA muons is associated with a TMS muon. Such dimuons are referred to as STA*-STA* dimuons here. Ordinarily, such events would be part of the TMS-TMS category, but we ensure that these are promptly produced (and therefore are not in any signal region) by requiring that the $L_{xy}/\sigma(L_{xy})$ of the resulting associated TMS-TMS dimuons is small, $L_{xy}/\sigma(L_{xy}) < 1$. This sample of events should have characteristics expected from prompt dimuons that are reconstructed as displaced STA-STA dimuons due to reconstruction mistakes or vertex fit anomalies —

precisely the type of events we wish to study. Failures in the STA muon reconstruction are largely decoupled from failures of the TMS muon reconstruction: when a prompt muon gives rise to an STA muon reconstructed as displaced, the same muon is usually accurately reconstructed as prompt by the TMS muon reconstruction. The STA muon reconstruction is very different from the TMS muon reconstruction, as the former is based on information in the muon system, whereas the latter is based predominantly on information from the tracker, with the muon system used mostly to identify muons and not to reconstruct the tracks. Therefore, the $L_{xy}/\sigma(L_{xy}) < 1$ requirement applied to TMS-TMS dimuons only ensures that the dimuons are prompt, but does not affect the $L_{xy}/\sigma(L_{xy})$ distribution of the associated STA-STA dimuons. At the same time, this sample of events allows us to accurately study STA-STA dimuons with large $L_{xy}/\sigma(L_{xy})$, because a vast majority of Drell-Yan events are reconstructed as both STA-STA dimuons as well as prompt TMS-TMS dimuons. This is illustrated in Figure 5.36, which shows a comparison between the $|\Delta\Phi|$ distributions in STA-STA (red) and STA*-STA* (blue) dimuons in a control region enriched in Drell-Yan background events by requiring $L_{xy}/\sigma(L_{xy}) < 6$. The distributions are scaled to unit area to facilitate a comparison of the shapes, which are similar, confirming that a control region with the STA to TMS muon association reversed is well suited for studying backgrounds in STA-STA dimuons without associated TMS muons. We also benefit from the fact that STA*-STA* dimuons are far more numerous than STA-STA dimuons (11.2 million STA*-STA* dimuon events vs 96 STA-STA dimuon events in this control region), corresponding to much smaller statistical uncertainties in distributions with STA*-STA* dimuons, allowing an accurate measurement of R_{DY}.

In order to minimize contamination from QCD events, which are studied separately, we additionally require that the dimuon invariant mass calculated using TMS muon quantities is larger than 15 GeV and that each TMS muon is isolated. For this purpose, the relative isolation of the associated TMS muon track, defined as the scalar sum of the p_T of tracks in a cone of $\Delta R < 0.3$ around the muon divided by the muon p_T, is required to be smaller than 0.05.

Our estimate of backgrounds other than QCD is based on the number of events in the control regions symmetric to the signal regions, $|\Delta\Phi| > \pi - |\Delta\Phi|_{\mathrm{C}}$, which are

- $|\Delta\Phi| > 3\pi/4$ in the case of the RPV SUSY model, and
- $|\Delta\Phi| > 9\pi/10$ in the case of the HAHM model.

We estimate the number of expected $|\Delta\Phi|$-symmetric background events, N_{DY}, by:

$$N_{\mathrm{DY}}(\mathrm{OS}; |\Delta\Phi| < |\Delta\Phi|_{\mathrm{C}}) = N(\mathrm{OS}; |\Delta\Phi| > \pi - |\Delta\Phi|_{\mathrm{C}}) \cdot R_{\mathrm{DY}} \tag{5.3}$$

where $N_{\mathrm{DY}}(\mathrm{OS}; |\Delta\Phi| < |\Delta\Phi|_{\mathrm{C}})$ and $N_{\mathrm{DY}}(\mathrm{OS}; |\Delta\Phi| > \pi - |\Delta\Phi|_{\mathrm{C}})$ are, respectively, the numbers of DY background events and its $|\Delta\Phi|$ symmetric CR. This CR is defined by $|\Delta\Phi|_{\mathrm{C}} = \pi/4$ for the RPV SUSY model, and $|\Delta\Phi|_{\mathrm{C}} = \pi/10$ for the HAHM model. R_{DY} is

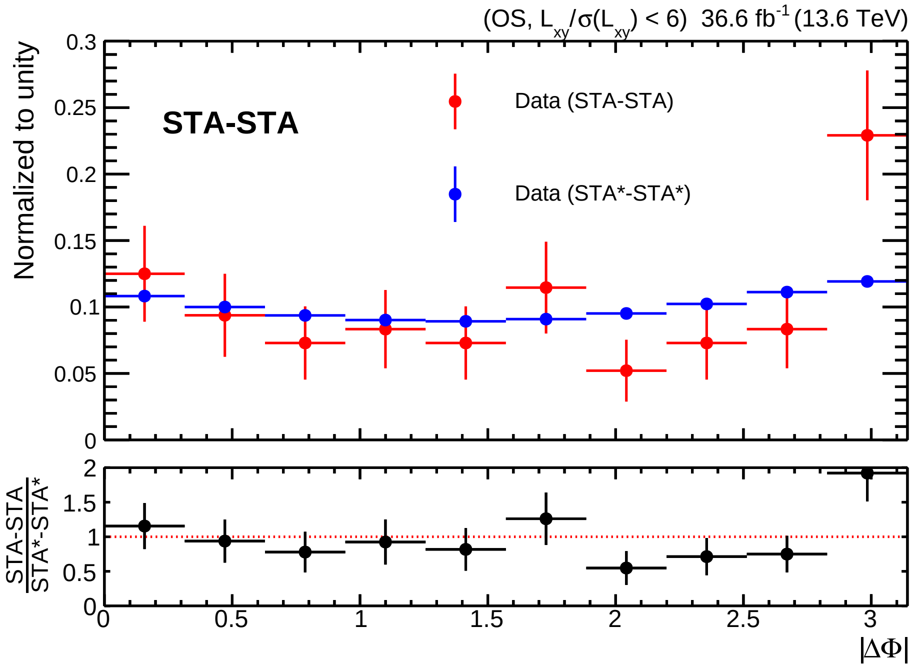

Figure 5.36: Comparison of the $|\Delta\Phi|$ distributions of OS STA-STA dimuons with neither of the two (red) and both (blue) STA muons associated to TMS muons, denoted as STA-STA and STA*-STA* dimuons respectively, in a control region enriched in Drell-Yan events by requiring $L_{xy}/\sigma(L_{xy}) < 6$. Both distributions are normalized to unity. The lower panel shows the ratio of the two distributions. The displayed errors are statistical.

the transfer factor, defined as

$$R_{\mathrm{DY}} = \frac{N_{\mathrm{DY}}^{\text{STA*-STA*}}(\mathrm{OS}; |\Delta\Phi| < |\Delta\Phi|_{\mathrm{C}})}{N_{\mathrm{DY}}^{\text{STA*-STA*}}(\mathrm{OS}; |\Delta\Phi| > \pi - |\Delta\Phi|_{\mathrm{C}})} \tag{5.4}$$

is measured in the control region with the STA to TMS muon association reversed for both muons (denoted by the superscript ‘STA*-STA*’). Figure 5.37 illustrates the signal and control regions that are used in this background evaluation method. It accounts for the deviation from unity of the ratio of such events in the signal region to that in the $|\Delta\Phi|$-symmetric control region. From this sample, we measure $R_{\mathrm{DY}} = 0.75 \pm 0.02$ (stat.) for $|\Delta\Phi|_{\mathrm{C}} = \pi/10$, and $R_{\mathrm{DY}} = 0.73 \pm 0.02$ (stat.) for $|\Delta\Phi|_{\mathrm{C}} = \pi/4$.

Figure 5.38 shows the $|\Delta\Phi|$ distributions for STA-STA dimuons with $|\Delta\Phi| < \pi/2$ (blue) and $|\Delta\Phi| > \pi/2$ (red) with the STA to TMS association reversed in 2022 data. The plot is constructed by superimposing a mirrored $|\Delta\Phi| > \pi/2$ distribution upon the $|\Delta\Phi| < \pi/2$ distribution, with the ratio of the two distributions, corresponding to the bin-wise transfer-factor R_{DY}, on the bottom pad. The ratio is nearly flat, with an excess of about 20% in events with $|\Delta\Phi| > \pi/2$ over events with $|\Delta\Phi| < \pi/2$.

The enhancement of events in both halves of the $|\Delta\Phi|$ spectrum around $|\Delta\Phi| = \pi/2$ is introduced by a correlation between $\sigma(L_{xy})$ and the dimuon orientation, illustrated in Figure 5.39. When a dimuon is formed from two muons that are back-to-back in the

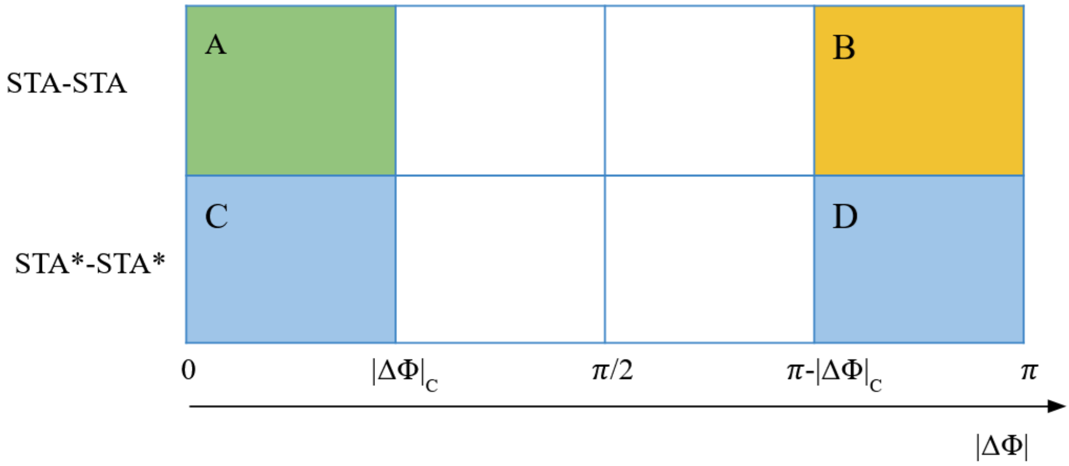

Figure 5.37: Illustration of the background estimation strategy in terms of the signal and control regions used. The signal region, defined as the region that passes all the selection criteria, is denoted by A, and we aim to estimate the DY-like background yield in this region. The control region B consists of events that fall in the $|\Delta\Phi|$-symmetric region, i.e. events that have $|\Delta\Phi| > \pi - |\Delta\Phi|_{\mathrm{C}}$, but pass all the other selection criteria. The measurements regions C and D consist of events in which the STA-STA dimuons are associated to a TMS-TMS dimuon (denoted by STA*-STA*) with $L_{xy}/\sigma(L_{xy}) < 1$, an invariant mass $80 < m_{\mu\mu} < 100$ GeV, and both TMS muons are required to be isolated. The STA*-STA* dimuons have $|\Delta\Phi| < |\Delta\Phi|_{\mathrm{C}}$ in region C, and $|\Delta\Phi| > \pi - |\Delta\Phi|_{\mathrm{C}}$ in region D. Thus, the expected background yield can be computed as $N_{\mathrm{A}} = N_{\mathrm{C}}/N_{\mathrm{D}} \cdot N_{\mathrm{B}}$. These quantities are explicitly described in Equations 5.3 and 5.4.

transverse place, as in the majority of Drell-Yan events, its $\vec{p}_T$ vector is usually pointing in the direction of one of the muons. This is because the directions of the muons are usually well reconstructed, whereas their p_T may not be. The uncertainty in the position of the fitted dimuon vertex is generally larger along the direction of the dimuon $\vec{p}_{T,\mu\mu}$ vector than orthogonal to it. In the extreme case of exactly back-to-back muons, the position along the $\vec{p}_{T,\mu\mu}$ vector is, in fact, undefined. Therefore, the selection of $L_{xy}/\sigma(L_{xy}) > 6$ enhances the contribution at $|\Delta\Phi| = \pi/2$ – namely, where the dimuon $\vec{p}_{T,\mu\mu}$ is perpendicular to the $\vec{L}_{xy}$ vector, and $\sigma(L_{xy})$ is the smallest. However, it does not break the symmetry of the $|\Delta\Phi|$ distribution around $\pi/2$.

5.6.2 Validation of Backgrounds with ‘symmetric’ collinearity

The $|\Delta\Phi|$-symmetric background evaluation method outlined in Section 5.6.1 is tested in a validation region obtained by requiring $L_{xy}/\sigma(L_{xy}) \leq 6$. Dimuon events in this sample are expected to originate predominantly in prompt processes such as Drell-Yan, which have an approximately symmetric $|\Delta\Phi|$ distribution. To remove any residual contamination from asymmetric $|\Delta\Phi|$ backgrounds such as QCD (discussed separately in Section 5.6.3), we also require the dimuon invariant mass $m_{\mu\mu} > 15$ GeV. The signal contribution to this

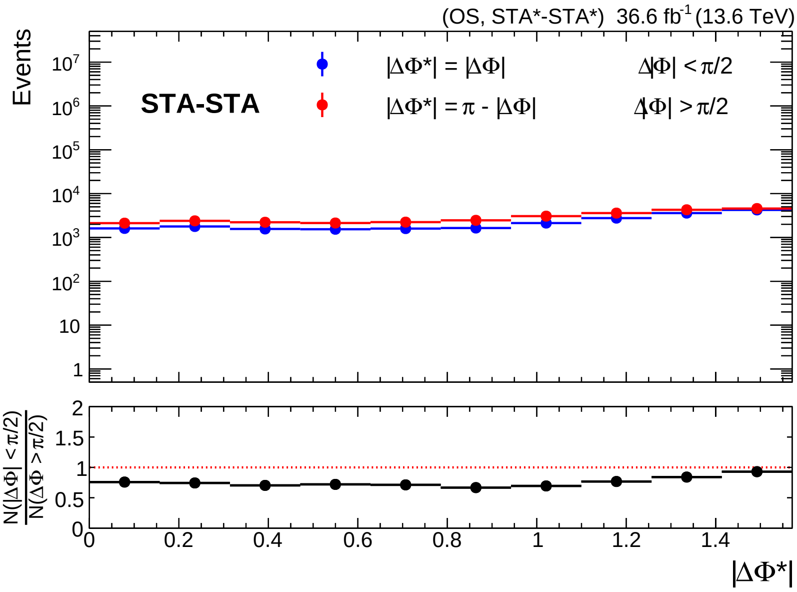

Figure 5.38: Distributions of the $|\Delta\Phi^*|$, defined in the legend, for STA*-STA* dimuons in 2022 data in events in the control region with the STA to TMS muon association reversed. The bottom pad shows the ratio of the distributions in the $|\Delta\Phi|_\mathrm{C}$ region to the symmetric $\pi - |\Delta\Phi|_\mathrm{C}$ region, which corresponds to the transfer factor R_DY.

region is expected to be negligible, given the limits obtained by the previous searches [22, 57, 58].

To validate the method and its dependence on the dimuon displacement and collinearity, we perform the background evaluation in the validation region, defined as above, in bins of $L_{xy}/\sigma(L_{xy})$ and $|\Delta\Phi|$. We compare the expected number of background events in this validation region, computed according to Equations 5.3 and 5.4 separately in each $L_{xy}/\sigma(L_{xy})$ and $|\Delta\Phi|$ bin, to the observed data in this region and bin. Figure 5.40 shows the results of the comparison in 2022 data.

The background prediction is consistent with the data. The statistical uncertainties in the background prediction are dominated by the limited number of events in the $|\Delta\Phi|$-symmetric control region, $\pi - |\Delta\Phi|_\mathrm{C}$. The measured values of R_DY are found to be similar between $L_{xy}/\sigma(L_{xy}) \leq 6$ and $L_{xy}/\sigma(L_{xy}) > 6$, and for data from all data-taking years analyzed in Run 3 and Run 2, namely 2022, 2018 and 2016.

5.6.3 Estimation of backgrounds with 'asymmetric' collinearity

The background evaluation method described in Section 5.6.1 relies on the symmetry of the $|\Delta\Phi|$ distributions and does not account for the contribution from background sources that predominantly yield dimuons with small $|\Delta\Phi|$, such as QCD processes that feature:

- dimuon decays of J/ψ and other low-mass resonances,
- cascade decays of B mesons, $b \to c\mu_1 X_1; c \to \mu_2 X_2$,

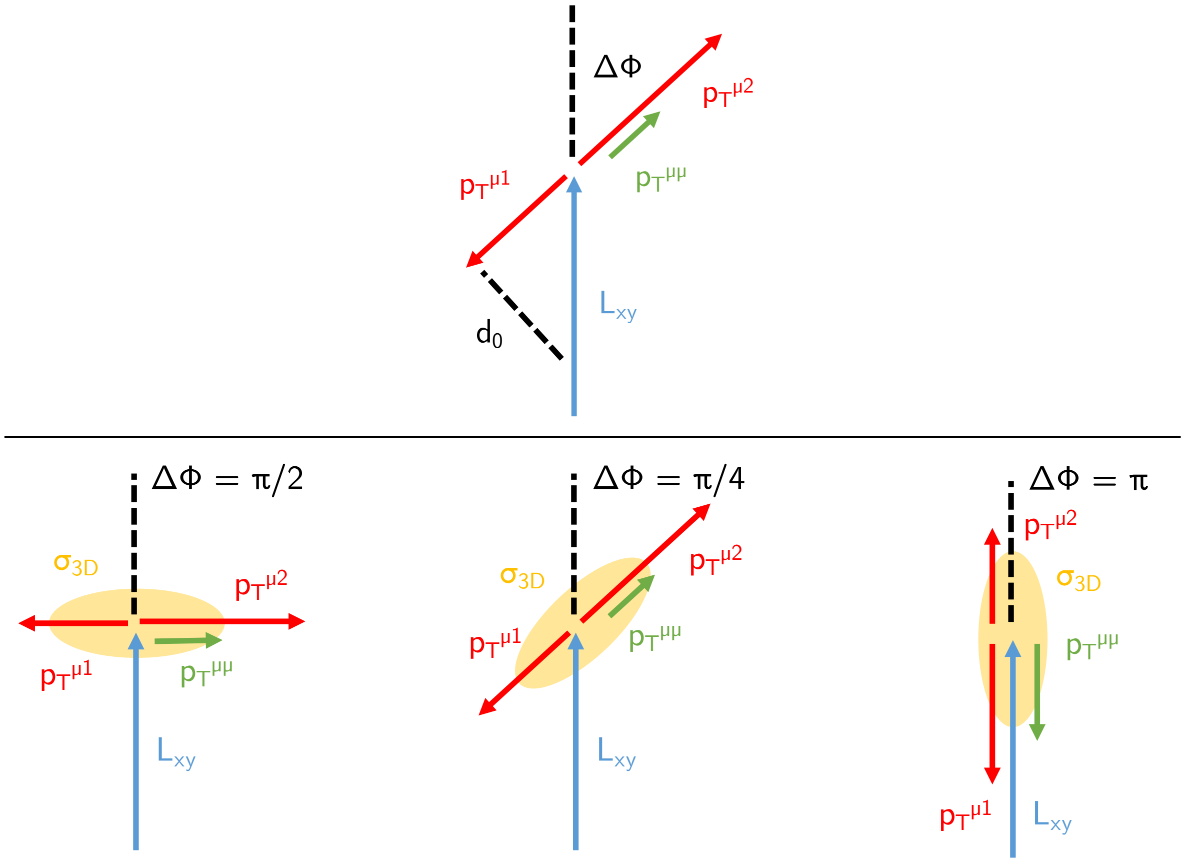

Figure 5.39: Diagrams depicting the effect of dimuon orientation on the shape of the $|\Delta\Phi|$ distribution. The dimuon p_T is assumed to point along one of the muons, which are back-to-back in the transverse plane. For a given L_{xy}, the uncertainty in the position of the fitted dimuon vertex (represented by a yellow ellipse around the vertex) is larger when the muons are parallel to the $\vec{L}_{xy}$ vector (i.e. $|\Delta\Phi| = 0$ or π), and smaller when the muons are perpendicular to it (i.e. $|\Delta\Phi| = \pi/2$), taking an intermediate value at $|\Delta\Phi| = \pi/4$ or $3\pi/4$. Reproduced from Ref. [96].

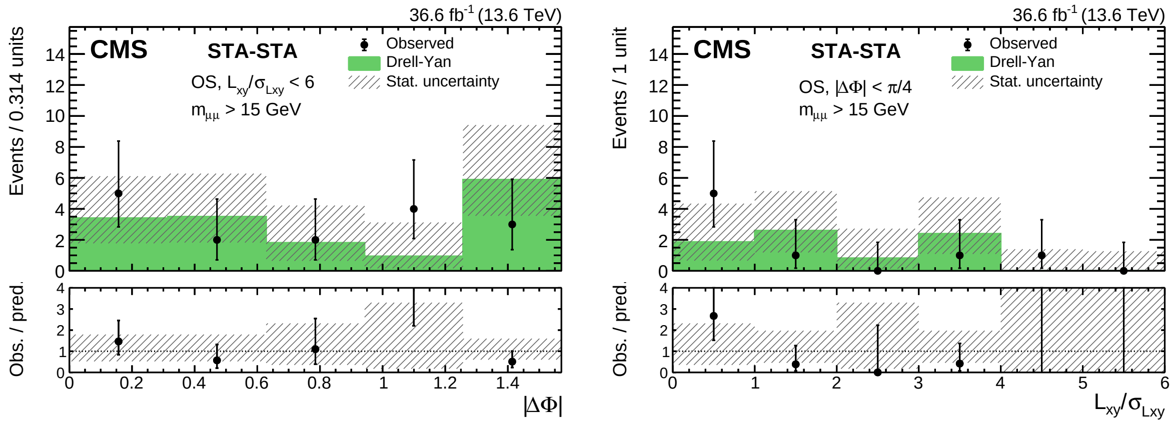

Figure 5.40: Background prediction checks in the STA-STA category, showing distributions of $|\Delta\Phi|$ (left) and $L_{xy}/\sigma(L_{xy})$ (right) for events with $m_{\mu\mu} > 15$ GeV in the $L_{xy}/\sigma(L_{xy}) < 6$ validation region in data (black circles) compared to the background predictions (histogram). The lower panels show the ratio of the observed to the predicted number of events. Hatched histograms show the statistical uncertainty in the background prediction. Reproduced from Ref. [23].

- dimuons formed from a pair of unrelated non-prompt muons in the same jet, for example, one muon produced in a hadron decay in flight and the other one arising from a b or c quark decay.

We collectively refer to these '$|\Delta\Phi|$-asymmetric' background sources as QCD backgrounds. If well reconstructed, most of the dimuons from this type of background have an invariant mass up to a few GeV and are mostly rejected by the $m_{\mu\mu} > 10$ GeV requirement. However, a small fraction of them could satisfy the invariant mass requirement due to a mismeasured invariant mass, and pass the event selection. These dimuons usually have small $|\Delta\Phi|$ values, as the $\vec{p}_{T,\mu\mu}$ and their $\vec{L}_{xy}$ vectors are collinear, and may have large signal-like $L_{xy}/\sigma(L_{xy})$ values. They may also feature an invariant mass close to the 10 GeV threshold.

Muons from such backgrounds are expected to be a part of jets, surrounded by tracks from other particles, and are hence not isolated. Additionally, while well-reconstructed signal muons are almost exclusively OS, muons from QCD backgrounds may also be SS. Exploiting this nature of muons from QCD background processes, we can study two small-$|\Delta\Phi|$ control regions with SS and OS dimuons where at least one of the STA muons fails the isolation requirement described in Section 5.5.4. Figure 5.41 shows the $|\Delta\Phi|$ (left) and invariant mass $m_{\mu\mu}$ (right) distributions for SS STA-STA dimuons passing (blue) and failing (red) the STA muon isolation requirements. In contrast to the $|\Delta\Phi|$-symmetric (DY) backgrounds described in Section 5.6.1, the $|\Delta\Phi|$-asymmetric (QCD) background distributions feature a distribution that falls off as $|\Delta\Phi|$ increases. From both the $|\Delta\Phi|$ and the $m_{\mu\mu}$ distributions, we see that the behavior of dimuons with both muons isolated is similar to dimuons in which at least one of the muons is not isolated.

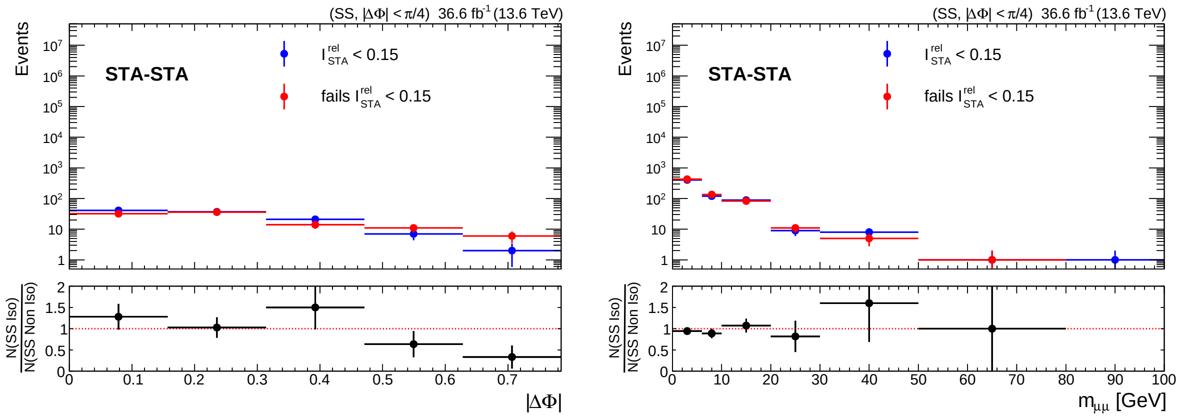

Figure 5.41: Distributions of $N_{\text{non-iso}}(\text{SS}; |\Delta\Phi| < \pi/4)$ in red, and $N_{\text{iso}}(\text{SS}; |\Delta\Phi| < \pi/4)$ in blue, as well as their ratio R_{QCD} in the lower panels, as a function of $|\Delta\Phi|$ (left) and the dimuon invariant mass (right) in 2022 data. The values of R_{QCD} in the mass bins below 10 GeV are used in the QCD-enriched validation region discussed in Section 5.6.3.

Therefore, we base our estimate of the QCD background on the number of events with OS non-isolated dimuons for the HAHM model (with $|\Delta\Phi|_{\text{C}} = \pi/10$) and for the RPV SUSY model (with $|\Delta\Phi|_{\text{C}} = \pi/4$). The number of estimated background events, N_{QCD} is obtained as follows:

$$N_{\text{QCD}}(\text{OS}; |\Delta\Phi| < |\Delta\Phi|_{\text{C}}) = N_{\text{non-iso}}(\text{OS}; |\Delta\Phi| < |\Delta\Phi|_{\text{C}}) \cdot R_{\text{QCD}} \tag{5.5}$$

where $N_{\text{non-iso}}(\text{OS}; |\Delta\Phi| < |\Delta\Phi|_{\text{C}})$ is the number of small-$|\Delta\Phi|$ OS non-isolated dimuons, and R_{QCD} is the transfer factor between isolated to non-isolated dimuons that is evaluated in SS dimuons in the same small-$|\Delta\Phi|$ region as follows:

$$R_{\text{QCD}} = \frac{N_{\text{iso}}(SS; |\Delta\Phi| < |\Delta\Phi|_{\text{C}})}{N_{\text{non-iso}}(SS; |\Delta\Phi| < |\Delta\Phi|_{\text{C}})} \tag{5.6}$$

where $N_{\text{iso}}(SS; |\Delta\Phi| < |\Delta\Phi|_{\text{C}})$ and $N_{\text{non-iso}}(SS; |\Delta\Phi| < |\Delta\Phi|_{\text{C}})$ are the numbers of isolated and non-isolated SS STA-STA dimuons with $|\Delta\Phi| < |\Delta\Phi|_{\text{C}}$. To improve the statistical precision on R_{QCD}, we increased the sample of SS dimuons by relaxing the tight requirements on the number of DT hits and muon segments applied to dimuons with $|\eta| < 0.1$ as well as the generic requirement on the minimum number of DT hits. Figure 5.42 illustrates the signal and control regions that are used in this background evaluation method.

The ratio pad in Figure 5.41 shows distributions of R_{QCD} in 2022 data as a function of the $|\Delta\Phi|$ (left) and invariant mass (right). R_{QCD} is compatible with unity, within statistical uncertainties. We don't observe a strong dependence of R_{QCD} on $|\Delta\Phi|$ or mass; however, due to the large statistical uncertainties of these measurements, we cannot exclude residual dependence on these variables. Therefore, we evaluate R_{QCD} separately for each $|\Delta\Phi|_{\text{C}}$ and mass window in the signal region.

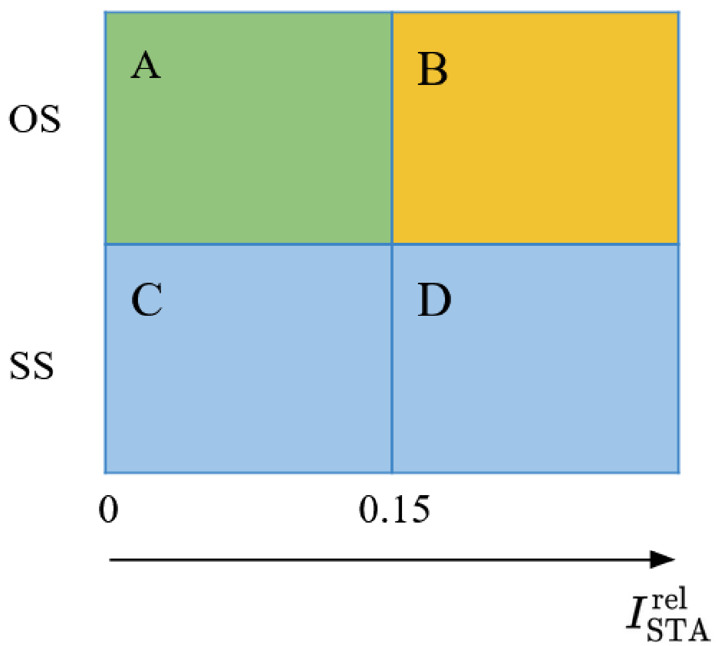

Figure 5.42: Illustration of the $|\Delta\Phi|$-asymmetric background estimation strategy in terms of the signal and control regions used. The signal region, defined as the region that passes all the selection criteria, is denoted by A, and we aim to estimate the QCD-like background yield in this region. The control region B consists of events in which at least one STA muon fails the isolation requirement I^{rel}_{STA}, but which pass all the other selection criteria. The measurements regions C and D consist of SS events, in which both STA muons are isolated (C), or in which at least one STA muon fails this requirement (D). Thus, the expected background yield can be computed as $N_{\mathrm{A}} = N_{\mathrm{C}}/N_{\mathrm{D}} \cdot N_{\mathrm{B}}$. These quantities are described explicitly in Equations 5.5 and 5.6.

Validation of Backgrounds with 'asymmetric' collinearity

The background evaluation method for '$|\Delta\Phi|$-asymmetric' backgrounds is tested in QCD-enriched validation regions obtained by inverting the invariant mass requirement and selecting dimuons with low mass ($6 < m_{\mu\mu} < 10$ GeV). The signal contribution to this region is expected to be negligible given the limits obtained by the previous Run 2 searches [22, 54].

The expected number of opposite-sign background events in this control region with $|\Delta\Phi| < \pi/4$ is calculated according to Equations 5.5 and 5.6 and compared to the corresponding number of observed events. The results are shown in Figure 5.43 as a function of the $|\Delta\Phi|$ (left), showing the method is able to reproduce the $|\Delta\Phi|$-asymmetric shape of the QCD background, and as a function of the dimuon invariant mass (right). There is a good agreement between the predicted and observed backgrounds.

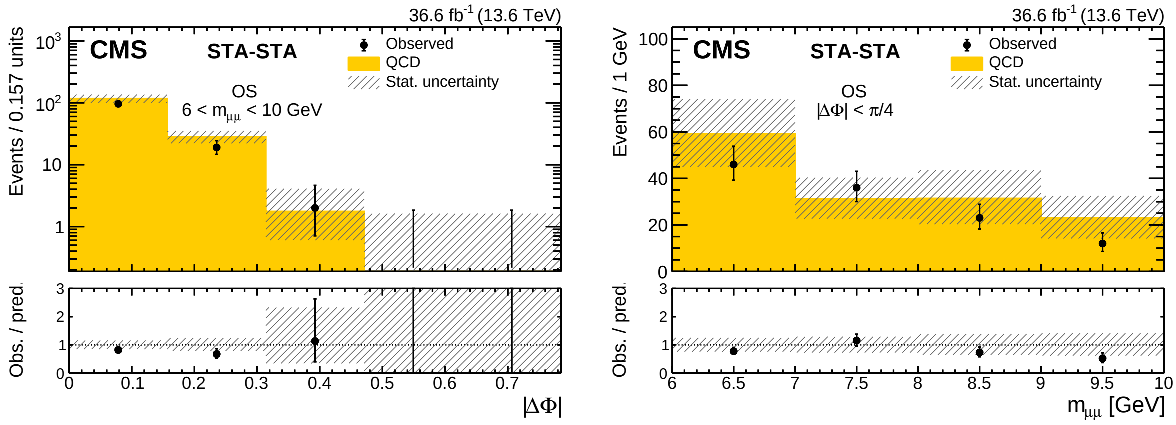

Figure 5.43: Background prediction checks in the STA-STA category, showing distributions of $|\Delta\Phi|$ (left) and $m_{\mu\mu}$ (right) for dimuons in the low-mass ($6 < m_{\mu\mu} < 10$ GeV) validation region in data (black circles) compared to background predictions (histograms). The lower panels show the ratio of observed to predicted number of events. Hatched histograms show the statistical uncertainty in the background prediction. Reproduced from Ref. [23].

5.6.4 Cosmic ray backgrounds

We computed the residual contamination from cosmic-ray muons in the analysis signal region with an auxiliary measurement that provides an upper bound for the number of dimuons passing the $\cos\alpha$ cut, N_p, via

$$N_p = N_f \frac{\varepsilon}{1-\varepsilon} \tag{5.7}$$

where N_f is the number of events that fail the nominal $\cos\alpha$ cut, and ε is the efficiency to pass the $\cos\alpha$ cut obtained from a sample enriched in cosmic-ray muons. This estimate is conservative in that it assumes all dimuons failing the $\cos\alpha$ requirement are of cosmic

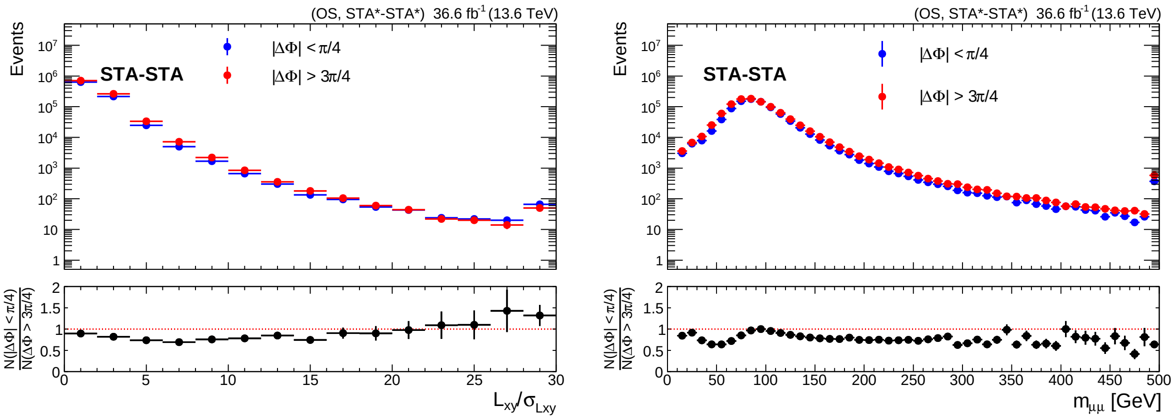

Figure 5.44: Distributions of $L_{xy}/\sigma(L_{xy})$(left) and $m_{\mu\mu}$ (right) for STA*-STA* dimuons in 2022 data in events in the control region defined by reversing the STA to TMS muon association. The ratio pad shows the differential DY transfer factor R_{DY}.

origin. The measurement results in an upper limit of 0.01 residual cosmic-ray events in 2022.

5.6.5 Systematic Uncertainties in background prediction

The $|\Delta\Phi|$-symmetric (DY) background estimation uncertainty is dominated by statistical uncertainties as a result of the small number of events in the STA-STA Drell-Yan enriched control regions, with $|\Delta\Phi| > \pi - |\Delta\Phi|_{\mathrm{C}}$. In contrast, the measurement of the transfer factor R_{DY} is performed in control regions with large statistical power due to the reversal of the STA to TMS muon association. Therefore, while the statistical uncertainties of R_{DY} are small, its measurement is still prone to systematic uncertainties.

We studied the dependence of R_{DY} on several key variables. The ratio pad of Figure 5.44 (left) shows R_{DY} as a function of $L_{xy}/\sigma(L_{xy})$. Its dependence on the STA-STA invariant mass is shown in Figure 5.44 (right). We measure that R_{DY} varies between 0.5 and 1.0, with a mean value of 0.75 (0.72) in the $|\Delta\Phi| < \pi/4$ ($|\Delta\Phi| < \pi/10$) signal regions. We also measured R_{DY} in events where only one of the STA muons is associated to a TMS muon, to verify the robustness of the method against the reversal of the TMS muon association and replacement, and found compatible results within statistical uncertainties, as shown in Table 5.4. Based on these studies, we assign a systematic uncertainty of 30% in R_{DY} for all signal regions considered. This uncertainty has a negligible effect on the DY background prediction.

Similarly, the $|\Delta\Phi|$-asymmetric (QCD) background prediction in the low-mass validation region, shown in Figure 5.43 is compatible with the observed yields, and no systematic uncertainty in R_{QCD} is derived from these checks. However, Equation 5.5 assumes that R_{QCD} can be measured with SS dimuons and applied to OS dimuons. To verify this assumption, we compute R_{QCD} using OS events failing the N(dimuon segments) or the N(DT hits) requirements applied to STA-STA dimuons with $|\eta| < 0.1$, described in Sec-

$R_{\rm DY}^{\rm STA\text{-}STA}$ definition	$\lvert\Delta\Phi\rvert_{\rm C} = \pi/4$	$\lvert\Delta\Phi\rvert_{\rm C} = \pi/10$
Both STA muons associated to TMS muons	0.75 ± 0.02	0.72 ± 0.03
One STA muon associated to TMS muon	0.87 ± 0.33	0.50 ± 0.31
(One STA muon associated) / (Both STA muons associated)	1.16 ± 0.44	0.69 ± 0.43

Table 5.4: $R_{\rm DY}$ obtained in the samples of dimuons with one or both STA muons associated to TMS muons, and their ratio.

tion 5.5.9. The results, shown in Table 5.5, confirm that both methods yield compatible transfer factors within their statistical uncertainties. Based on this study, we conclude that $R_{\rm QCD}$ is dominated by statistical uncertainties in all considered $m_{\mu\mu}$ intervals, and we do not assign any systematic uncertainty to this prediction.

Mass interval [GeV]	$\lvert\Delta\Phi\rvert_{\rm C}$	$R_{\rm QCD}$ in SS	$R_{\rm QCD}$ in OS small-$\lvert\Delta\eta_{\mu\mu}\rvert$
$6 < m_{\mu\mu} < 10$	$\pi/4$	0.89 ± 0.12	0.74 ± 0.06
$10 < m_{\mu\mu} < 15$	$\pi/10$	1.09 ± 0.22	1.13 ± 0.20
$15 < m_{\mu\mu} < 32$	$\pi/10$	1.21 ± 0.38	0.88 ± 0.32
$20 < m_{\mu\mu} < 60$	$\pi/10$	1.50 ± 0.68	1.20 ± 0.73
$m_{\mu\mu} > 15$	$\pi/4$	1.02 ± 0.22	0.81 ± 0.21

Table 5.5: $R_{\rm QCD}$ obtained in with SS or OS small-$\lvert\Delta\eta_{\mu\mu}\rvert$ STA-STA dimuons.

5.7 Systematics uncertainties affecting signal

The modelling of signal efficiencies in the simulation was extensively studied in the Run 2 search. As the analysis does not make changes to the CMS detector or in the muon reconstruction algorithms between Run 2 and Run 3, we take many of the systematic uncertainties and data-to-simulation corrections, called scale factors (SFs), directly from the preceding search [22]. We performed new studies to address the modelling of signal efficiencies related to the new trigger algorithms and modified event selection criteria, such as the STA muon isolation. We performed these studies for each dimuon category separately, using dedicated data samples. In both categories, the muon identification, muon reconstruction, and trigger efficiencies are the sources of the dominant systematic uncertainties and the largest data-to-simulation corrections.

All the sources of uncertainty are listed and described below:

- **Luminosity**: The uncertainty in the integrated luminosity is 2.3% [110], and affects the signal rate normalization.

- **Pileup**: To obtain the scale factors and systematic uncertainties due to pileup, we first obtained normalized distributions of the number of primary vertices per event in data and in each simulated signal sample, as shown in Figure 5.45. The ratio of these distributions in data to simulation yielded weights, $\omega = N_{\rm Data}({\rm PV})/N_{\rm MC}({\rm PV})$

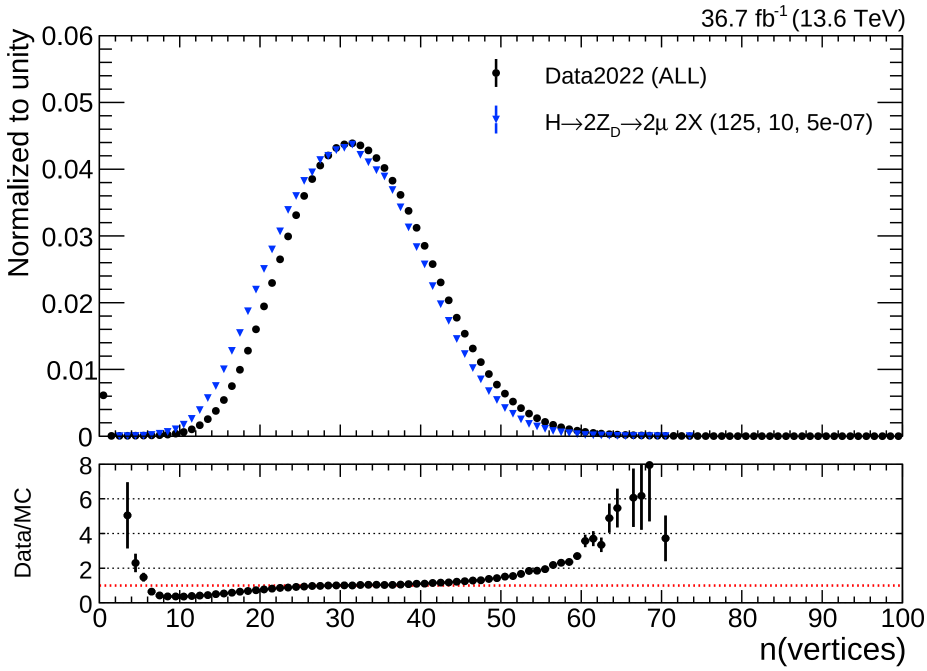

Figure 5.45: Normalised distribution of the number of primary vertices in data (black dots) and in one HAHM signal sample (blue triangles) with $m(H) = 125$ GeV, $m(Z_D) = 20$ GeV, and kinetic mixing constant $\epsilon = 5 \times 10^{-7}$. The ratio pad at the bottom shows the ratio of these two distributions.

by which each simulated event was scaled. This was done separately for each signal sample. We then computed signal efficiencies after the full event selection from these reweighted samples, and compared them to the signal efficiencies obtained without this reweighting procedure. The ratio of signal efficiencies after reweighting to before is between a factor 0.94 and 0.97 depending on $c\tau$, as shown in Figure 5.46 for all signal samples. Therefore, we correct the signal efficiencies by a factor of 0.95 and assign a systematic uncertainty of 2%.

- **STA muon identification and reconstruction**: As stated above, we rely on the studies done for Run 2 in 2018 data for the STA muon identification and reconstruction, as these aspects of the analysis were inherited without changes, and use the same SFs and uncertainties. The STA muon efficiencies given by the signal simulation are corrected by reconstruction and identification SFs for prompt muons extrapolated to large d_0. The p_T and η dependent SFs for prompt STA muons are obtained by applying the 'tag-and-probe' method [111] to muons from J/ψ and Z boson decays. The extrapolation to non-zero d_0 values is performed using d_0 dependent correction functions derived in the analysis of cosmic-ray muon data and simulation. The resulting efficiency corrections for STA-STA dimuons in the signal samples range from 0.75 to 1, depending on the sample, and represent the largest data-driven correction in the STA-STA category. The overall uncertainty increases continuously with displacement, from $\approx 5\%$ for prompt-like muons to $\approx 15\%$ for

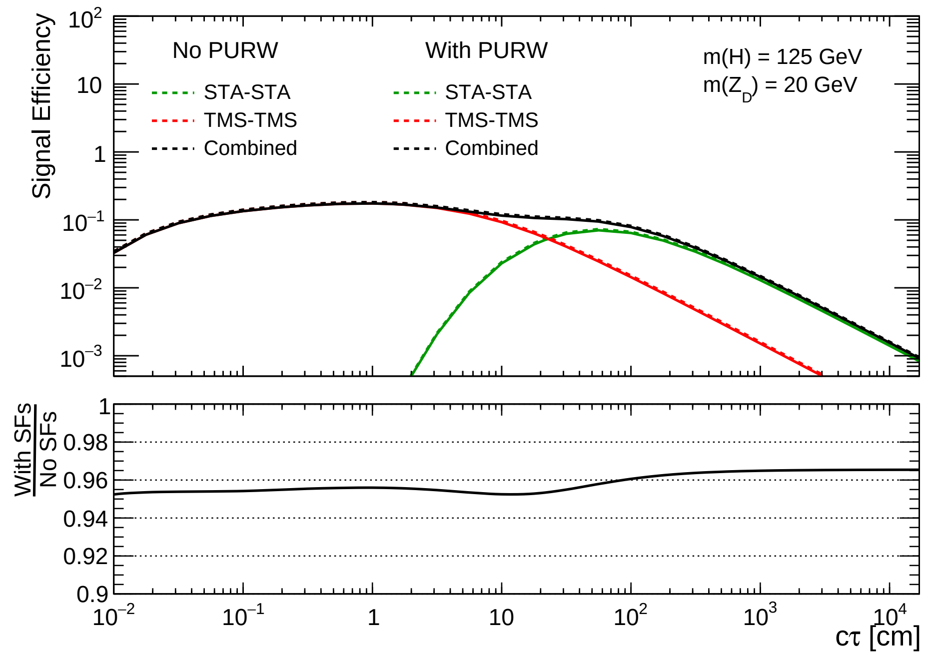

Figure 5.46: Signal efficiencies showing the effect of the pileup reweighting procedure in the TMS-TMS and STA-STA dimuon categories, compared to the signal efficiencies obtained without pileup reweighted, for the HAHM signal samples with $m(Z_D) = 20$ GeV. The bottom pad shows the ratio of the reweighted to un-reweighted efficiencies.

muons with $d_0 \approx 100$ cm, corresponding to the precision of studies using cosmic-ray muons. Systematic uncertainties related to the muon p_T resolution are less than 5%.

- **Trigger**: As the triggers have changed significantly in Run 3, the trigger SFs and associated systematic uncertainties have been re-evaluated. We performed dedicated studies to examine the effect of the d_0 thresholds and looser muon p_T requirements of the new displaced dimuon triggers by comparing the efficiencies evaluated with $J/\psi \to \mu\mu$ data with the efficiencies using the HAHM signal samples with $m(Z_D) = 10$ GeV. These results are shown in Figure 5.47, as a function of $\min(d_0)$ (left) and $\min(d_0/\sigma(d_0))$ (right), in TMS-TMS dimuons. A more robust measurement was also obtained using the JetMET dataset and simulated $J/\psi \to \mu\mu$ samples, as shown previously in Figure 5.15. From the ratio of these distributions, we obtain a scale factor of 0.95, by which we correct the simulated yields of signal events selected by the 'Run 3 (2022, L3)' set of triggers and assign an uncertainty of 5%.

  Similarly, to verify the modelling of displacement-related effects in the new 'Run 3 (2022, L2)' set of triggers, we compare the efficiencies in Figure 5.16, obtained using cosmic-ray muon data, with the efficiencies in the HAHM signal samples. The results are shown in Figure 5.48 as a function of $\min(d_0)$ in STA-STA dimuons. While no systematic biases are observed, we see that the efficiencies in signal simulation agree

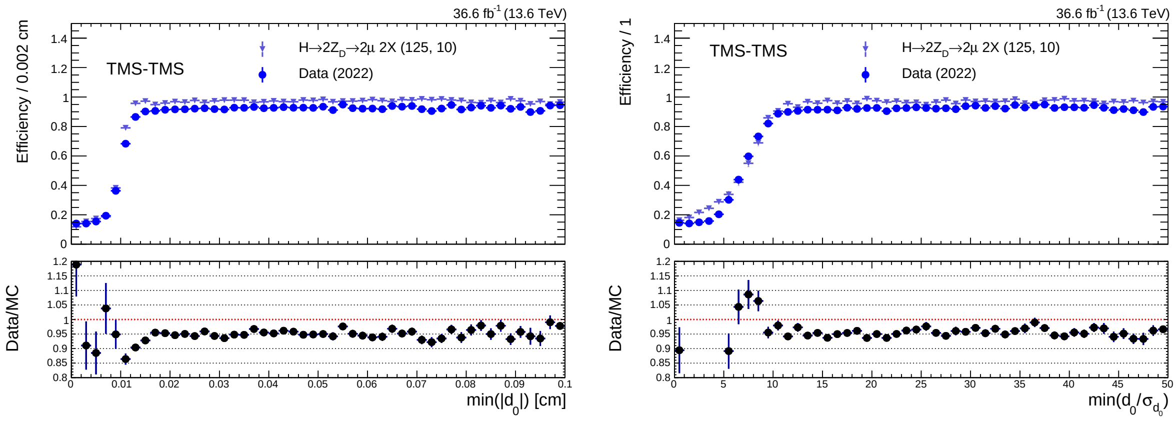

Figure 5.47: Efficiency of the 'Run 3 (2022, L3)' trigger as a function of the $\min(d_0)$ (left) and $\min(d_0/\sigma(d_0))$ (right) of TMS-TMS dimuons in $J/\psi \to \mu\mu$ data (blue) and the HAHM signal samples with $m(Z_D) = 10$ GeV (light blue), and their ratios.

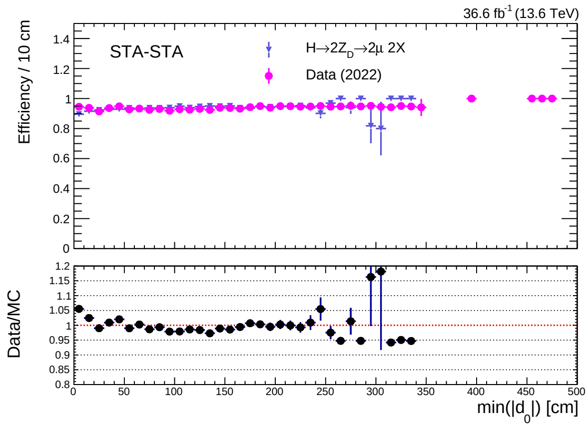

Figure 5.48: Efficiency of the 'Run 3 (2022, L2)' trigger as a function of $\min(d_0)$ of STA-STA dimuons in cosmic ray data (pink) and the HAHM signal samples (light blue), and their ratios.

to within 3% with those in data, up to $\min(d_0)\approx 300$ cm. The measurement in the first bin, $\min(d_0) < 10$ cm, is not expected to yield an accurate SF as it includes the trigger turn-on and is sensitive to the underlying $\min(d_0)$ distribution in each of the samples. Thus, we conclude that the dependence of the efficiencies on d_0 is well reproduced in simulation, and assign the largest data-to-simulation difference of 3%, in bins of d_0, as a systematic uncertainty in the efficiency of these triggers.

- **STA muon isolation**: The STA muon isolation is a new handle in the Run 3 search, and thus we performed dedicated studies to evaluate the accuracy of its modelling in signal simulation. As shown in Figure 5.49, we compare the isolation efficiency for muons originating in Drell-Yan processes (left) and HAHM signal samples (right) as a function of the subleading STA muon p_T in events with the leading STA muon passing the $I^{rel}_{\mathrm{STA}} < 0.15$ requirement. The Drell-Yan sample is obtained from STA-STA dimuons with $|\Delta\Phi| > 3\pi/4$, associated to TMS-TMS dimuons with $80 < m_{\mu\mu} < 100$ GeV. For the HAHM signal, all available samples have been considered. The STA isolation efficiency in data is close to 100% and drops to $\approx$ 90% at 10 GeV. This behavior is well reproduced in signal simulation, as shown by the data-to-simulation SFs in Figure 5.49 (bottom), except for STA muon with $p_T < 20$ GeV, where the SFs are larger than unity and simulation underestimates the efficiencies. For the analysis, we take a conservative approach and do not correct this effect in simulation. We assume an SF equal to unity and apply a 6% systematic uncertainty on the signal efficiency to account for these differences.

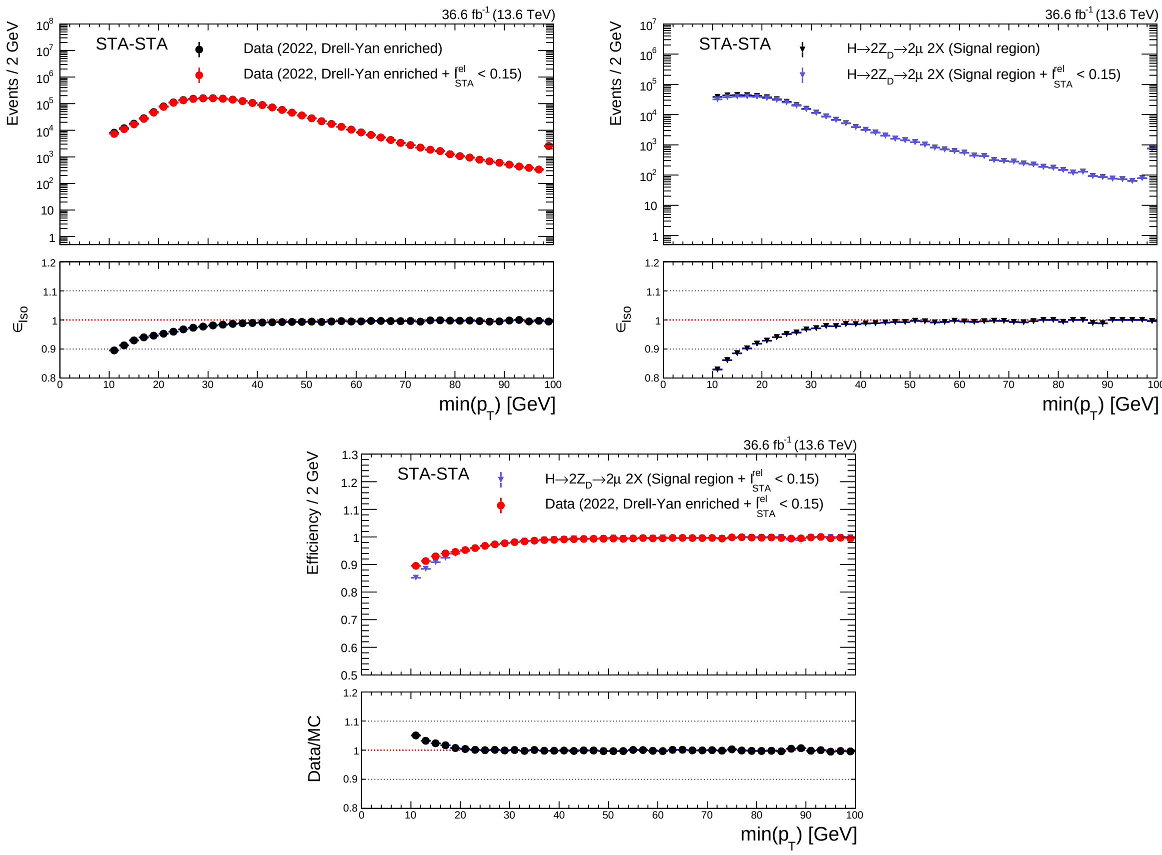

Figure 5.49: STA isolation efficiency for the subleading STA muon in Drell-Yan (left) and HAHM signal samples (right) as a function of the minimum STA muon p_T, computed in events passing the full event selection except the requirement on I^{rel}_{STA}. Data-to-simulation scale factors (bottom) are obtained from the ratio of these two efficiencies.

We studied the dependence of the SFs for the STA muon isolation on displacement by comparing the isolation efficiency for cosmic-ray muons selected in pp collision events with those in signal simulation. The results, shown as a function of min(d_0) of STA muons in Figure 5.50, demonstrate that the efficiencies and SFs are nearly independent of d_0, with variations as large as 3%, derived at small d_0. Thus we use this value as the systematic uncertainty.

- **STA-STA dimuon selection (vertex χ^2, DCA)**: The accuracy of the modelling of the vertex $\chi^2 < 10$ and DCA < 15 cm requirements in the signal simulation was

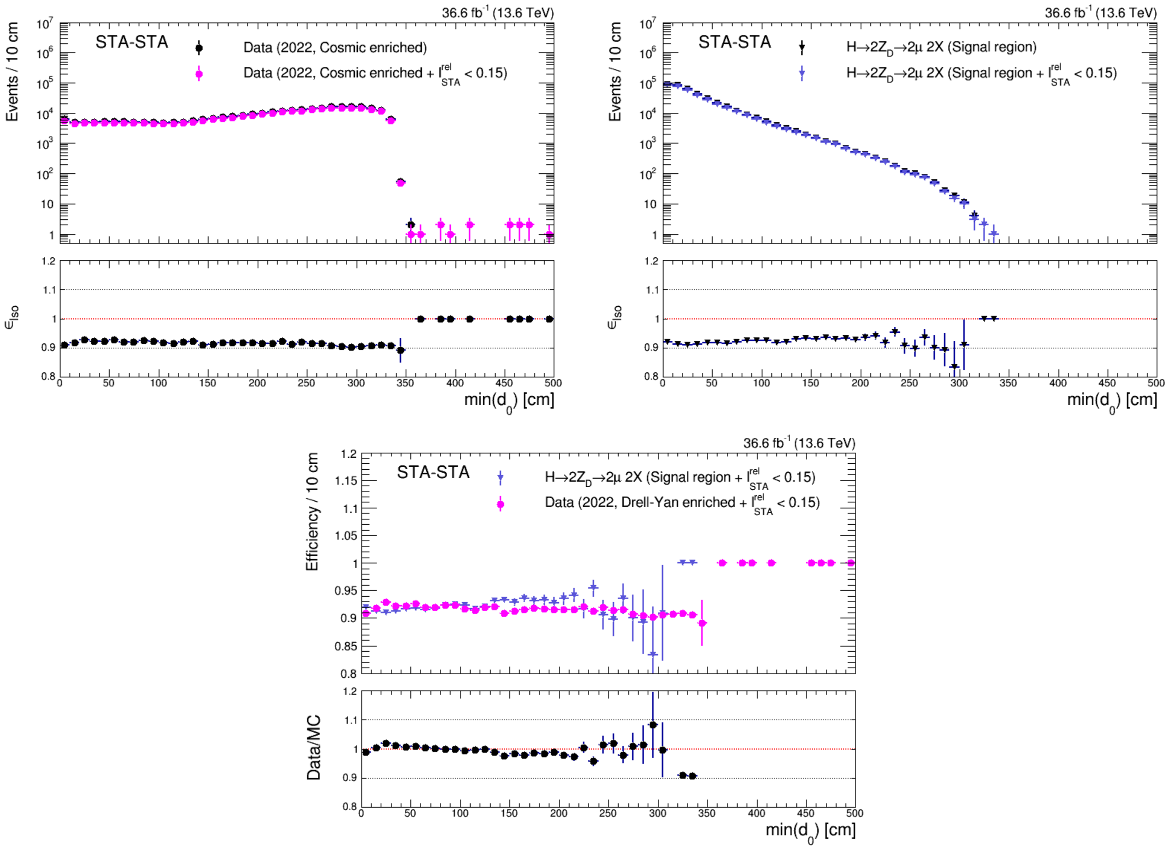

Figure 5.50: STA isolation efficiency for the subleading STA muon in cosmic ray data (left) and HAHM signal samples (right) as a function of the STA muon min(d_0), computed in events passing the full event selection except the requirement on I^{rel}_{STA}. Data-to-simulation scale factors (bottom) are obtained from the ratio of these two efficiencies.

checked in Run 2 and found to be similar between data and simulation. At the time, the dependence with d_0 was studied in cosmic muon data and signal simulation, and a very mild dependence on d_0 was found. The vertex χ^2 efficiency is flat up to $d_0 = 200$ cm, while the DCA efficiency increases by 1% (2%) at $d_0 = 100(200)$ cm. Thus, we assign a 2% systematic uncertainty in the vertex χ^2 and DCA selection efficiencies.

- **Higgs cross section**: The uncertainties in the Higgs boson production cross section at 13.6 TeV are $^{+5\%}_{-7\%}$ [112].

5.8 Statistical Analysis and Results

The analysis SR contains all OS dimuons passing the selection criteria outlined in Section 5.5 with the transverse collinearity $|\Delta\Phi| < \pi/10$ for the HAHM model, and $|\Delta\Phi| < \pi/4$ for the RPV SUSY model. This region was blinded until the final stages of the analysis, and these events were only observed after finalizing the full event selection and the background predictions.

The SR data was unblinded in two steps, with a randomly selected sample consisting of 30% of the data being unblinded first, and then the full dataset. Following the first step, we scrutinized the observed events for any anomalies originating from SM backgrounds, in case we missed some critical aspect of the background evaluation and required a modification of the analysis strategy. No such anomalies were found, and the full unblinding process concluded smoothly.

The background prediction in each mass interval, listed in Table 5.2, is computed as the sum of the $|\Delta\Phi|$-symmetric backgrounds composed of Drell-Yan and other non-QCD backgrounds (Section 5.6.1), and the $|\Delta\Phi|$-asymmetric backgrounds composed of the non-prompt QCD background (Section 5.6.3). The results are shown in Tables 5.6 and 5.7, in the intervals of $m_{\mu\mu}$ and $m_{\mu\mu}^{\mathrm{corr}}$, respectively. For each mass interval, the first four columns list the numbers of events observed in the $\pi - |\Delta\Phi|_{\mathrm{C}}$ and OS non-isolated small-$|\Delta\Phi|$ control regions, and the corresponding transfer factors. The penultimate column, $N_{\mathrm{bkg}}^{\mathrm{est}}$, shows the total predicted number of background events as well as its statistical uncertainty. This number illustrates the pre-fit background predictions.

Due to the small number of events in the control regions, the dominant uncertainty in the background predictions is statistical. According to the predictions, the residual background in the STA-STA category is expected to be mainly composed of QCD events. The QCD contribution is expected to be particularly large in the $10 < m_{\mu\mu} < 15$ GeV mass interval. Compared to Run 2, the expected background in 2022 is higher for $10 < m_{\mu\mu} < 15$ GeV, and similar or lower for $m_{\mu\mu} > 15$ GeV.

The number of observed events is in agreement with the expected background. We observe:

- 26 events with $m_{\mu\mu} > 10$ GeV and $|\Delta\Phi| < \pi/10$, compared to $36.3^{+9.8}_{-8.7}$ (stat) expected background events.

- 14 events with $m_{\mu\mu} > 15$ GeV and $|\Delta\Phi| < \pi/4$, compared to $8.9^{+4.7}_{-3.4}$ (stat) expected background events.

- Overall, in the region inclusive in mass and $|\Delta\Phi|$, with $m_{\mu\mu} > 10$ GeV and $|\Delta\Phi| < \pi/4$, we observe 37 events, compared to $45.4^{+10.4}_{-9.3}$ (stat.) expected background events.

To facilitate interpretation in terms of the HAHM and RPV-SUSY signal models, we further bin these events in terms of the mass windows described in 5.5.10. The number of observed events in each bin are in agreement with the expected background, and are listed in Tables 5.6 and 5.7.

Given that the expected and observed events are compatible in all search regions, we conclude that there is no statistically significant excess of signal-like events beyond the expected background from SM processes. Figure 5.51 shows the number of observed events overlaid over the predicted background yields in non-overlapping $m_{\mu\mu}$ (left) and $m_{\mu\mu}^{\mathrm{corr}}$ (right) intervals for the HAHM and RPV SUSY models respectively. Most of the

Mass interval [GeV]	N(OS; $\|\Delta\Phi\| > 9\pi/10$)	R_{DY}	$N_{\text{non-iso}}$(OS; $\|\Delta\Phi\| < \pi/10$)	R_{QCD}	$N_{\mathrm{bkg}}^{\mathrm{est}}$	N^{obs}
$10 < m_{\mu\mu} < 15$	0	0.75 ± 0.02	28	1.09 ± 0.22	$30.5^{+9.5}_{-8.5}$	18
$15 < m_{\mu\mu} < 32$	0	0.75 ± 0.02	2	1.21 ± 0.38	$2.4^{+3.6}_{-1.7}$	8
$20 < m_{\mu\mu} < 60$	1	0.75 ± 0.02	2	1.50 ± 0.68	$3.7^{+4.5}_{-2.4}$	1
$25 < m_{\mu\mu} < 80$	1	0.75 ± 0.02	2	1.50 ± 0.68	$3.7^{+4.5}_{-2.4}$	1
$30 < m_{\mu\mu} < 80$	1	0.75 ± 0.02	1	1.50 ± 0.68	$2.2^{+3.9}_{-1.5}$	0
$40 < m_{\mu\mu} < 120$	1	0.75 ± 0.02	1	1.50 ± 0.68	$2.2^{+3.9}_{-1.5}$	0
$40 < m_{\mu\mu} < 245$	1	0.75 ± 0.02	1	1.50 ± 0.68	$2.2^{+3.9}_{-1.5}$	0
$m_{\mu\mu} > 60$	0	0.75 ± 0.02	0	1.50 ± 0.68	$0.0^{+3.1}_{-0.0}$	0
$m_{\mu\mu} > 10$	1	0.75 ± 0.02	31	1.15 ± 0.19	$36.3^{+9.8}_{-8.7}$	26

Table 5.6: Background estimation and observed number of events in the STA-STA dimuon category in 2022 data in the $m_{\mu\mu}$ mass intervals used for the HAHM model. For each $m_{\mu\mu}$ interval, the first four columns show the number of events observed in the $|\Delta\Phi| > 9\pi/10$ and OS non-isolated control regions, and the measured transfer factors for each of the regions. The two last columns show the total number of predicted background events and the number of observed events, respectively. The quoted uncertainties are statistical only.

observed events have low $m_{\mu\mu}$, as expected from the background predictions. Post-fit signal distributions at the level of the median expected exclusion limits at 95% confidence level (CL) obtained from the ensemble of both dimuon categories in the absence of signal are also shown, enhanced for visibility.

The expected and observed numbers of events in each mass interval in the STA-STA category are interpreted as a single-bin counting experiment and combined with the results obtained in the TMS-TMS category, as explained in Section 5.8.1

5.8.1 Limit setting procedure

We use these results to set upper limits, as a function of the mass and lifetime, on the branching fraction of the Higgs boson to dark photons in the HAHM model, $\mathcal{B}(H \to Z_D Z_D)$, and on the product of the squark-antisquark production cross section and the branching ratio of the squark decay to quark plus neutralino in the RPV SUSY model, $\sigma(\mathrm{pp} \to \tilde{q}\bar{\tilde{q}}) \cdot \mathcal{B}(\tilde{q} \to q\tilde{\chi}^0_1)$. We extract the limits based on a modified frequentist approach [113, 114], using the Combine package [115] developed for statistically combining the results of the Higgs boson searches [116]. The calculation is performed using a likelihood, which is the product of Poisson distributions of the number of observed events in the TMS-TMS or STA-STA signal regions, the expected number of background events, the expected number of signal events, and the systematic uncertainties described in Section 5.7. The systematic uncertainties are incorporated as nuisance parameters in the likelihood using log-normal distributions in Combine, except the uncertainties in the background evaluation from control regions, which is discussed later in this section. The signal efficiencies are computed separately for each dimuon category, mass interval, LLP lifetime, and signal model. The

Mass interval [GeV]	N(OS; $\|\Delta\Phi\| > 3\pi/4$)	R_{DY}	$N_{\text{non-iso}}$(OS; $\|\Delta\Phi\| < \pi/4$)	R_{QCD}	$N_{\mathrm{bkg}}^{\mathrm{est}}$	N^{obs}
$m_{\mu\mu}^{\mathrm{corr}} > 25$	1	0.73 ± 0.02	8	1.05 ± 0.24	$9.2^{+4.9}_{-3.6}$	14
$m_{\mu\mu}^{\mathrm{corr}} > 50$	1	0.73 ± 0.02	5	0.95 ± 0.30	$5.5^{+3.9}_{-2.6}$	9
$m_{\mu\mu}^{\mathrm{corr}} > 80$	1	0.73 ± 0.02	4	1.50 ± 0.68	$6.7^{+5.7}_{-4.0}$	5
$m_{\mu\mu}^{\mathrm{corr}} > 90$	1	0.73 ± 0.02	4	1.50 ± 0.68	$6.7^{+5.7}_{-4.0}$	5
$m_{\mu\mu}^{\mathrm{corr}} > 210$	0	0.73 ± 0.02	1	1.50 ± 0.68	$1.5^{+3.8}_{-1.4}$	3
$m_{\mu\mu}^{\mathrm{corr}} > 300$	0	0.73 ± 0.02	0	1.50 ± 0.68	$0.0^{+3.1}_{-0.0}$	3
$m_{\mu\mu}^{\mathrm{corr}} > 15$	1	0.73 ± 0.02	8	1.02 ± 0.22	$8.9^{+4.7}_{-3.4}$	14

Table 5.7: Background estimation and observed number of events in the STA-STA dimuon category in 2022 data in the $m_{\mu\mu}^{\mathrm{corr}}$ mass intervals used for the RPV SUSY model. For each $m_{\mu\mu}^{\mathrm{corr}}$ interval, the first four columns show the number of events observed in the $|\Delta\Phi| > 3\pi/4$ and OS non-isolated control regions and the measured transfer factors for each of the regions. The two last columns show the total number of predicted background events and the number of observed events, respectively. The quoted uncertainties are statistical only.

individual likelihoods are then combined to obtain the combined results in the ensemble of both categories. To achieve even stronger bounds on the model parameters of the HAHM model, the results obtained in this analysis are also combined with the results of the Run 2 analysis [22]. The combination of the two results accounts for the different cross sections of the Higgs boson production in collisions at 13 and 13.6 TeV [101, 112].

Given the smallness of the expected background for events with one dimuon, and taking into account the selection efficiencies discussed in Section 5.5.11, introducing a separate category for events with two dimuons does not increase the sensitivity of the analysis significantly even in the most favorable case for the 4μ signal events, where $\mathcal{B}(X \to \mu\mu) = 1$. The gain is negligible for smaller $\mathcal{B}(X \to \mu\mu)$ values, such as those predicted by the HAHM model. One could potentially increase the selection efficiency for events with two dimuons by designing a dedicated set of selection criteria for this particular category of events, but this is beyond the scope of this thesis. Thus, no distinction is made between events with one or two reconstructed dimuons of the same type.

If a genuine signal were present, it would give rise to the number of signal events with at least one selected dimuon, N_S, given by:

$$N_S = \mathcal{L}\sigma \left[2\mathcal{B}(1-\mathcal{B})\epsilon_1 + \mathcal{B}^2\epsilon_2\right] \tag{5.8}$$

where $\mathcal{L}$ is the integrated luminosity of the data-taking period (36.6 fb^{-1}), σ is the signal production cross section, $\mathcal{B}$ is the branching fraction for the decay of the long-lived particle to a dimuon pair, and ϵ_1 and ϵ_2 are the fraction of 2μ and 4μ signal events, respectively, in which at least one dimuon candidate passes the full set of selection criteria. The signal

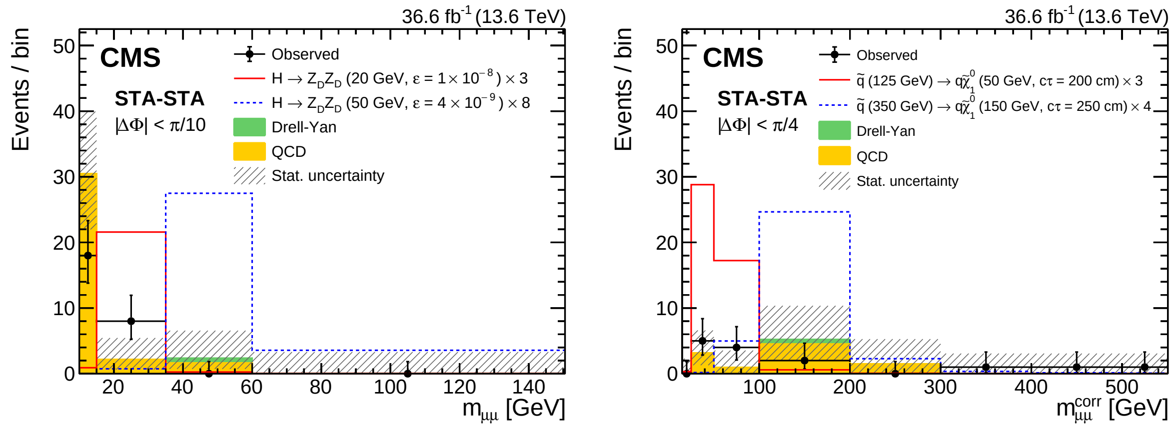

Figure 5.51: Comparison of the observed (black point) and expected (histograms) numbers of events in non-overlapping $m_{\mu\mu}$ (left) and $m_{\mu\mu}^{\text{corr}}$ (right) intervals in the STA-STA dimuon category, in the signal regions optimized for the HAHM (left) and RPV SUSY (right) models. The stacked yellow and green filled histograms represent mean expected background contributions from QCD and DY, respectively, while statistical uncertainties in the total expected background are shown as hatched histograms. Signal contributions expected from simulated signals indicated in the legends are shown in red and blue. Their yields are set to the corresponding median expected 95% CL exclusion limits obtained from the ensemble of both dimuon categories, scaled up as indicated in the legend to improve visibility. The last bin includes events in the histogram overflow. Reproduced from Ref. [23].

efficiencies are obtained from simulation and further corrected by the data-to-simulation scale factors described in Section 5.7, computed separately for each signal model, dimuon category, and mass window.

The predicted branching fraction of the SM Higgs boson decay into a pair of long-lived dark photons, $\mathcal{B}(H \to Z_D Z_D)$ depends on the parameters of the hidden sector, as shown in Equation 2.10. The dark photon branching ratio $\mathcal{B}(Z_D \to \mu\mu)$ depends only on $m(Z_D)$ and ranges between 10% and 16% in the mass range of interest, as summarized in Table 4.2. The signal efficiencies ε_{Z_D} are computed as the ratio of the number of signal events in which at least one dimuon candidate passes the full set of selection criteria to the number of generated events with at least one Z_D decaying into a muon pair (Section 5.5.11), evaluated using dark photon samples generated with the predicted branching ratio, and therefore containing the appropriate mix of 2μ and 4μ events. The signal yields in the HAHM model can then be expressed as:

$$N_S = \mathcal{L}\sigma(H)\mathcal{B}(H \to Z_D Z_D)\mathcal{B}(Z_D \to \mu\mu)\left[2 - \mathcal{B}(Z_D \to \mu\mu)\right] \cdot \varepsilon_{Z_D} \tag{5.9}$$

In each dimuon category, the backgrounds are implemented separately in the COMBINE datacards. The ABCD method yielding background predictions in the signal region is implemented using a multi-bin likelihood, which is the product of Poisson distributions corresponding to the signal regions and the control regions. Since the expected number of background events is small in all categories, the limits are calculated by generating toy experiments to compute the distributions of the test statistic.

To calculate an estimated number of signal events for lifetimes other than those for the MC signal samples listed in Tables 4.2 and 4.4, we use the lifetime reweighting procedure outlined in Appendix C. The procedure was refined to combine events in several MC samples produced with the same set of signal masses and different lifetimes and to use the combination of these events in the reweighting [107]. The validity of this method was verified by excluding one of the available samples from the reweighting procedure and comparing the computed number of events to the 'true' number of events in that sample. For the purpose of calculating upper limits, the relative statistical uncertainty is computed for each reweighted signal sample and taken as a systematic uncertainty in the estimated number of signal events. Only $c\tau$ points corresponding to signal rate estimates with statistical uncertainty smaller than 50% are included in the limit plots.

5.8.2 Upper limits in the HAHM model

Figures 5.52 and 5.53 show the 95% CL upper limits obtained in the framework of the HAHM model under the assumption that $m(H_D) > m(H)/2$, set on $\mathcal{B}(H \to Z_D Z_D)$ as functions of the mean proper decay length of Z_D for $m(Z_D)$ in the range 10-60 GeV. Figure 5.52 shows the expected limits in the individual dimuon categories as well as the expected and observed limits obtained in the ensemble of both categories. Figure 5.53 shows a comparison of the observed limits obtained in this analysis with the corresponding Run 2 limits [22], along with the expected and observed limits obtained in a combination of the two results.

The constraints on the parameters of the HAHM model are comparable or tighter in a significant fraction of the parameter space despite a dataset that is a factor of 2.5 smaller than that used by the Run 2 analysis, mainly due to the improvements in the trigger algorithms. The combined limits on $\mathcal{B}(H \to Z_D Z_D)$ are approximately a factor of 2 more stringent than the limits obtained in the Run 2 analysis, and exclude $\mathcal{B}(H \to Z_D Z_D)$ of 1% in the range of $c\tau(Z_D)$ from a few tens of μm to 30 m (700 m) for $m(Z_D) = 10$ GeV (60 GeV) at 95% CL. In the $m(Z_D)$ interval 20-60 GeV, $\mathcal{B}(H \to Z_D Z_D)$ of 0.01% is excluded at 95% CL in the $c\tau(Z_D)$ range of about 0.3 mm to 0.5 m, corresponding to a wide range of ϵ values from $\approx 10^{-8}$ to $\approx 10^{-6}$. These constraints on rare SM Higgs boson decays are tighter than those derived from searches for invisible Higgs boson decays [117, 118], and from measurements of the SM Higgs boson couplings [119]. At $m(Z_D) > 15$ GeV, the limits obtained are the most stringent limits to date for all $c\tau(Z_D)$ values except those between ≈ 0.1 cm and ≈ 10 cm, depending on $m(Z_D)$, where the best limits are provided by a previous CMS search [54] using data collected with a dedicated high-rate data stream, as shown in Figure 5.54.

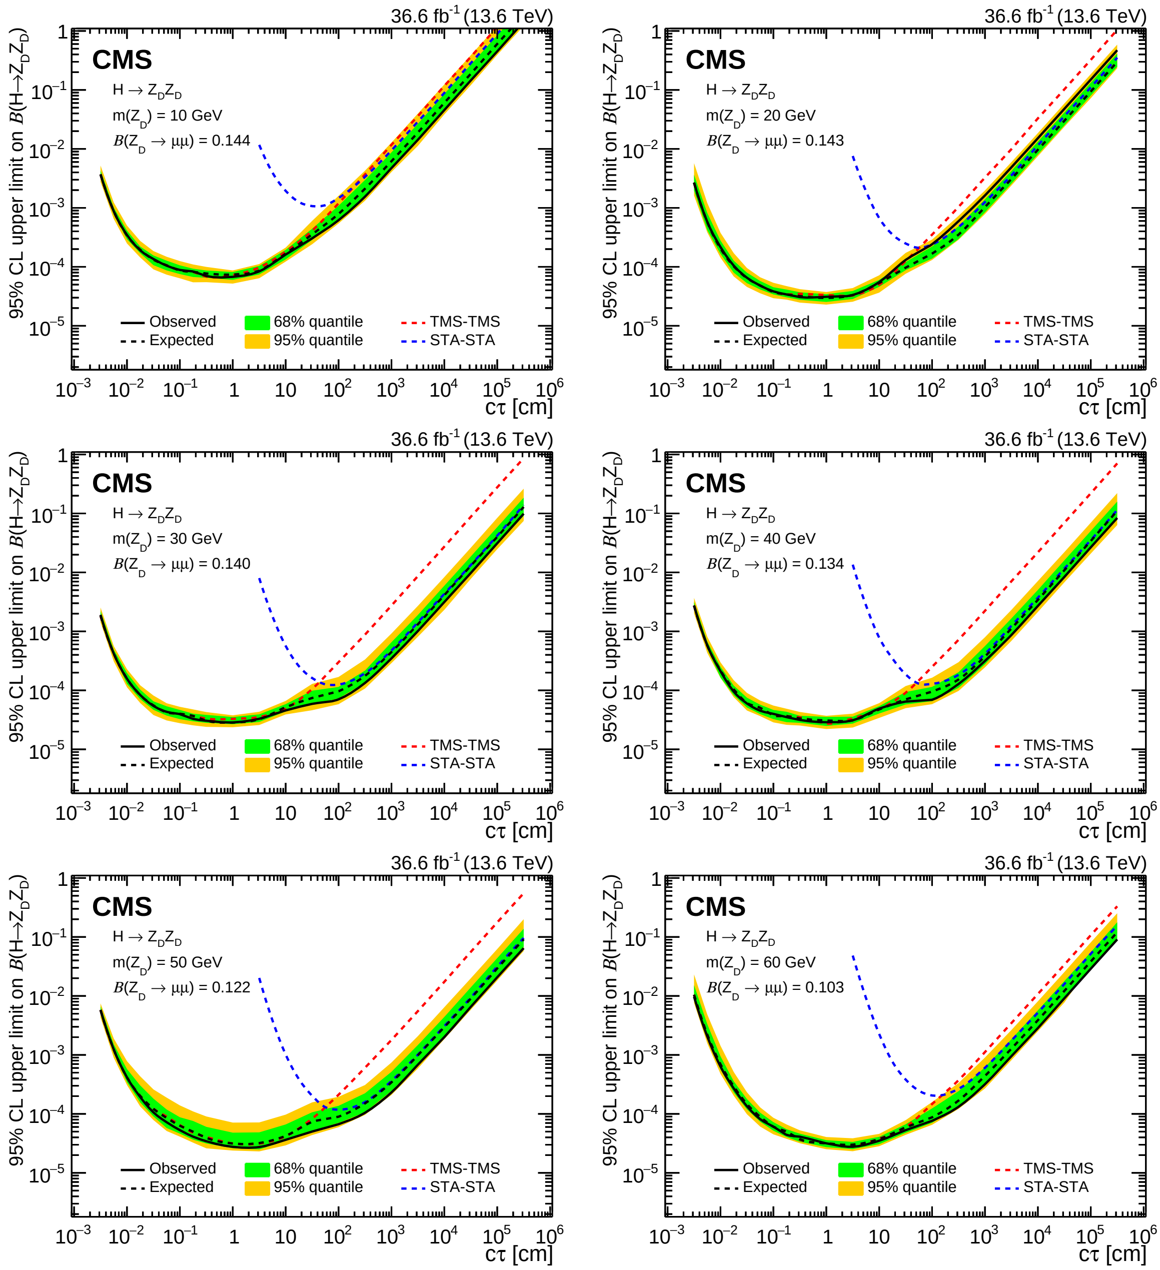

Figure 5.52: The 95% CL upper limits on $\mathcal{B}(H \to Z_D Z_D)$ as a function of the $c\tau(Z_D)$ in the HAHM model, for $m(Z_D)$ ranging from 10 GeV (top left) to 60 GeV (bottom right), in the STA-STA and TMS-TMS dimuon categories in 2022 data and their combination. The median expected limits obtained from the STA-STA and TMS-TMS dimuon categories are shown as dashed blue and red curves, respectively. The combined median expected limits are shown as dashed black curves, and the combined observed limits are shown as solid black curves. The green and yellow bands correspond, respectively, to the 68% and 95% quantiles for the combined expected limits. Reproduced from Ref. [23].

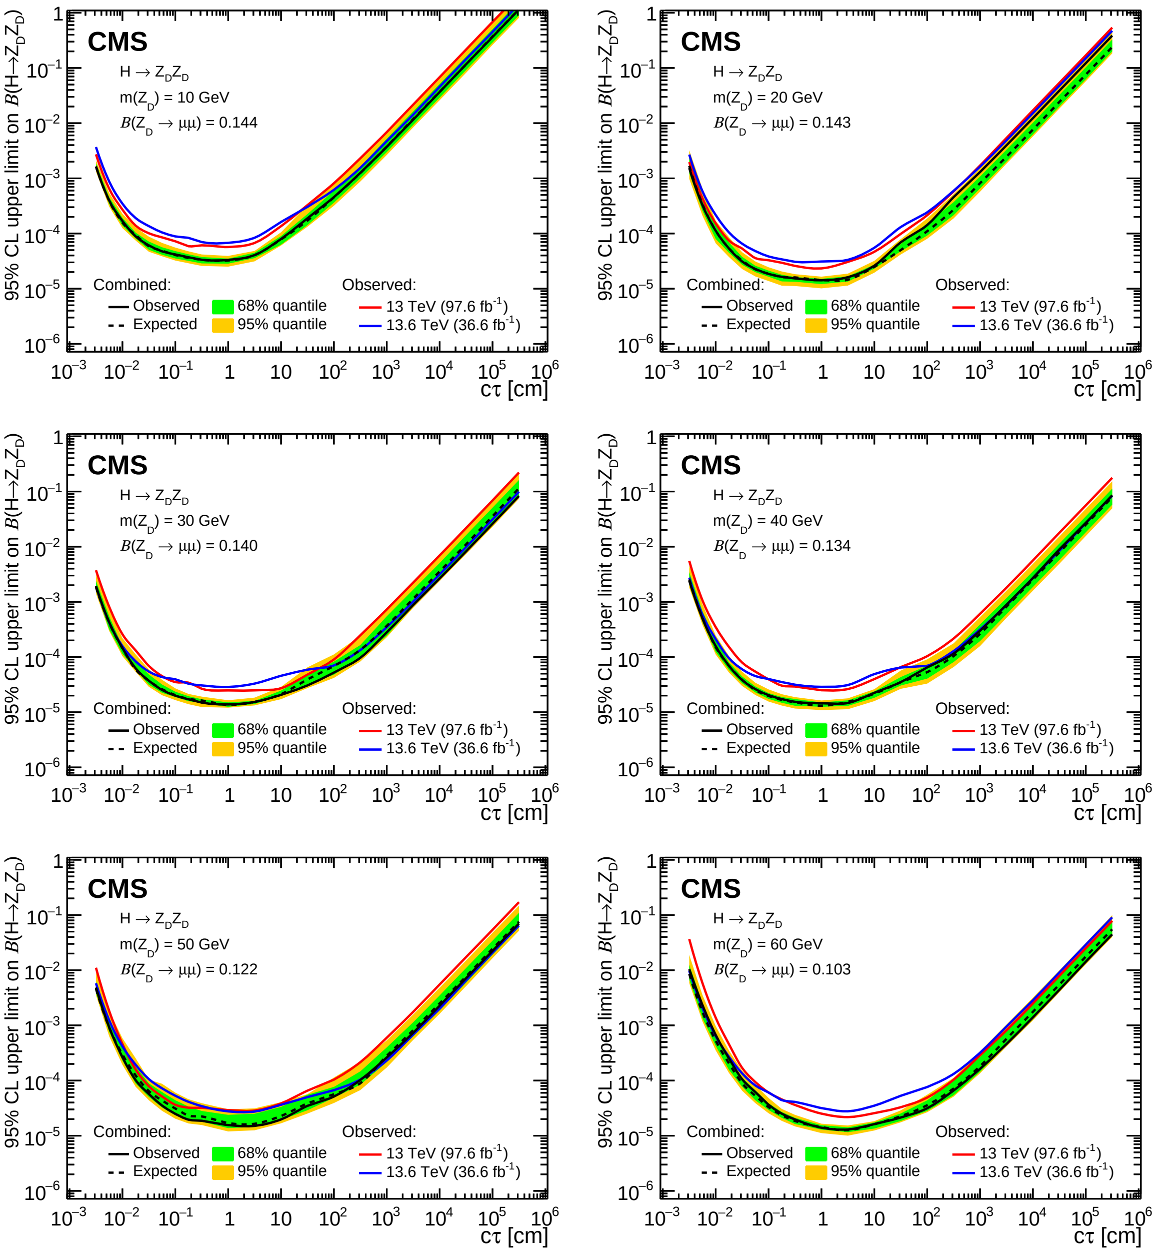

Figure 5.53: The 95% CL upper limits on $\mathcal{B}(H \to Z_D Z_D)$ as a function of $c\tau(Z_D)$ in the HAHM model for $m(Z_D)$ ranging from 10 GeV (top left) to 60 GeV (bottom right), obtained in this analysis, the Run 2 analysis [22], and their combination. The observed limits in this analysis and in the Run 2 analysis are shown as blue and red curves, respectively. The median combined expected limits are shown as dashed black curves, and the combined observed limits are shown as solid black curves. The green and yellow bands correspond, respectively, to the 68% and 98% quantiles for the combined expected limits. Reproduced from Ref. [23].

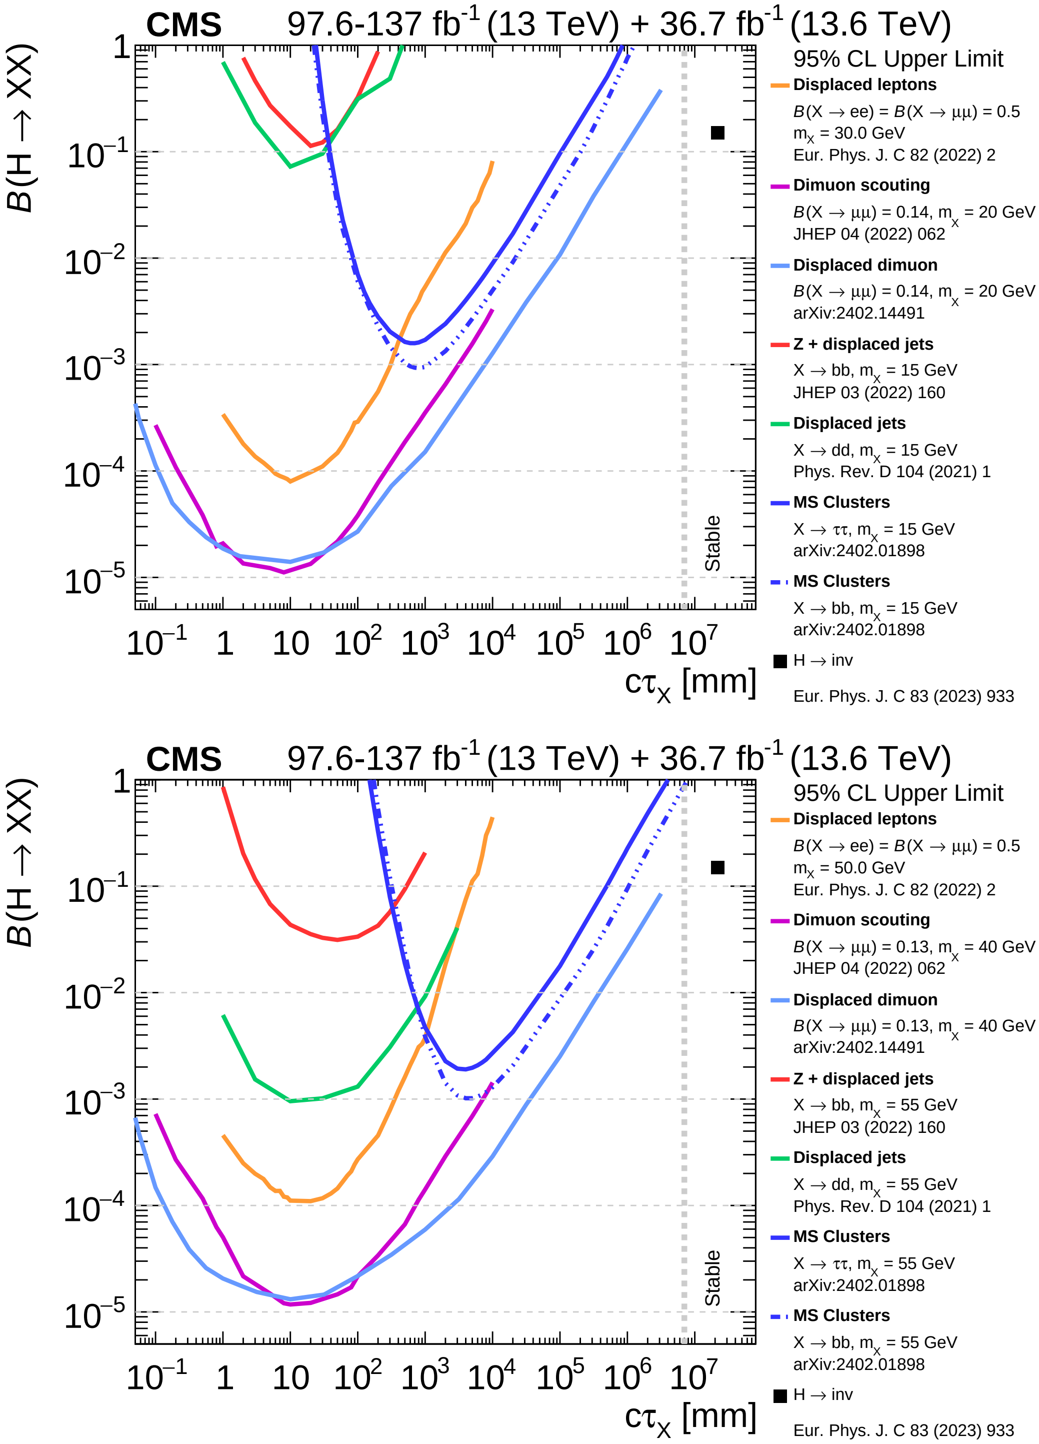

Figure 5.54: Observed 95% CL upper limits on $\mathcal{B}(H \to XX)$, for $15 < m(X) < 30$ GeV (top) and for $40 < m(X) < 55$ GeV (bottom). Reproduced from Ref. [24].

5.8.3 Upper limits in the RPV SUSY model

Figure 5.55 shows the 95% CL upper limits obtained in the ensemble of both dimuon categories in the framework of the RPV SUSY model. The limits are set on the product $\sigma(\mathrm{pp} \to \tilde{q}\bar{\tilde{q}}) \cdot \mathcal{B}(\tilde{q} \to q\tilde{\chi}_1^0)$, assuming $\mathcal{B}(\tilde{\chi}_1^0 \to \mu\mu\nu) = 0.5$. They are shown as functions of the mean proper decay length of $\tilde{\chi}_1^0$ for various combinations of $m(\tilde{q})$ and $m(\tilde{\chi}_1^0)$, as indicated in the legends, for $m(\tilde{q})$ values ranging from 125 GeV to 1.6 TeV, and compared to theoretical predictions. At $m(\tilde{q}) = 700$ GeV, the data excludes mean proper neutralino decay lengths between 30 μm and 350 m (at $m(\tilde{\chi}_1^0) = 50$ GeV), and between 30 μm and 1.5 km (at $m(\tilde{\chi}_1^0) = 500$ GeV). At $m(\tilde{q}) = 1.6$ TeV, the excluded $c\tau(\tilde{\chi}_1^0)$ range is between 0.07 and 4 cm for $m(\tilde{\chi}_1^0) = 50$ GeV, and between 70 μm and 2 m for $m(\tilde{\chi}_1^0) = 500$ GeV. The limits obtained in this analysis are tighter than those derived by CMS in the Run 1 search [57, 58] in all of the considered $(m(\tilde{q}), m(\tilde{\chi}_1^0), c\tau(\tilde{\chi}_1^0))$ parameter space, and are more stringent at $c\tau(\tilde{\chi}_1^0) \lesssim 1$ cm and $c\tau \gtrsim 1$ m than the limits on the λ_{122} RPV coupling set by the ATLAS Collaboration at $m(\tilde{q})$ of 700 GeV and 1.6 TeV [59], and are shown in Figure 5.56.

5.9 Reinterpretation

While the search is interpreted specifically in the context of the benchmark models described above, its results may be applied to other models predicting long-lived particles that decay to final states with a pair of oppositely charged muons. Thus, we provide a set of generator-level efficiency maps that approximate the reconstruction-level efficiencies of this analysis, and facilitate the reinterpretation of the results in the framework of other models. Figure 5.57 shows the signal efficiencies in the STA-STA dimuon category in the HAHM (right) and RPV SUSY (left) signal models, as a function of the minimum p_T and d_0 of the two generated muons with L_{xy}^{true} between 70-500 cm. Similar efficiency maps are provided for L_{xy}^{true} smaller than 20 cm, and between 20-70 cm, and are also provided for the TMS-TMS category. All figures and tables are also available as a HEPData [120] record, and may be found, along with instructions for the reinterpretation, in the supplementary material for the Run 3 search [23].

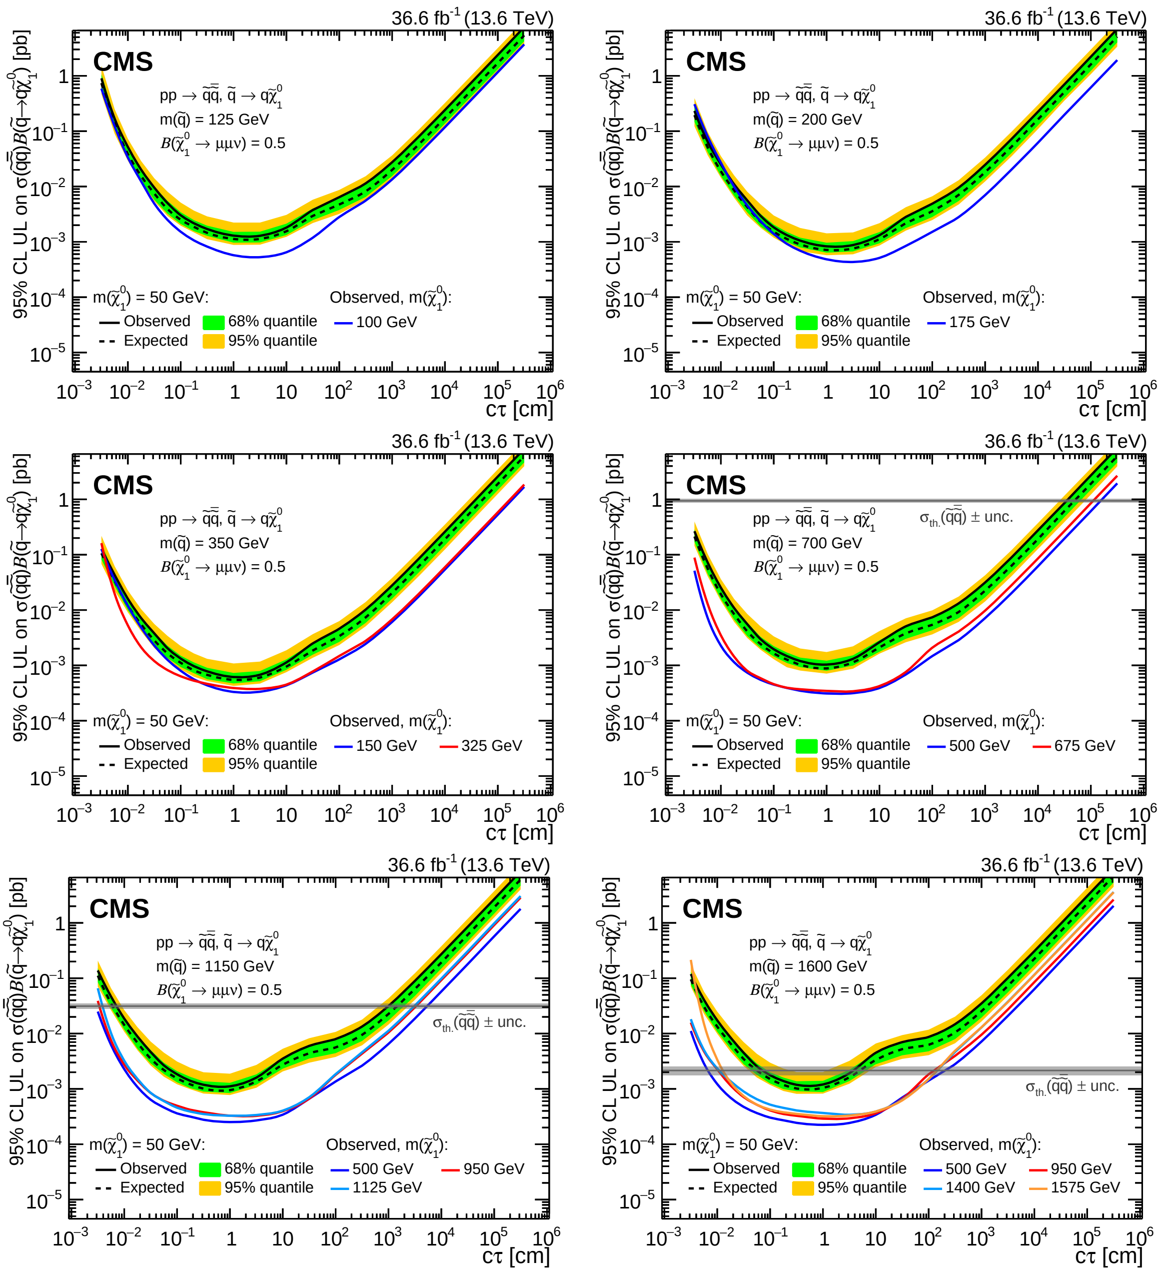

Figure 5.55: The 95% CL upper limits on $\sigma(\mathrm{pp} \to \tilde{q}\bar{\tilde{q}})\mathcal{B}(\tilde{q} \to q\tilde{\chi}_1^0)$ as a function of $c\tau(\tilde{\chi}_1^0)$ in the RPV SUSY model, for $\mathcal{B}(\tilde{\chi}_1^0 \to \mu\mu\nu) = 0.5$ and $m(\tilde{q})$ ranging from 125 GeV (top left) to 1.6 TeV (bottom right). The observed limits for various combinations of $m(\tilde{q})$ and $m(\tilde{\chi}_1^0)$ indicated in the legends are shown as solid curves, while the median expected limits are shown as dashed black curves. Their 68% and 95% quantiles are shown, respectively, as green and yellow bands for the case of $m(\tilde{\chi}_1^0) = 50$ GeV, and omitted for other neutralino masses for clarity. The gray horizontal lines indicate the theoretical values of the squark-antisquark production cross sections with the uncertainties shown as gray shaded bands. The predicted cross sections for $m(\tilde{q}) = 125$ GeV (7200 pb), 200 GeV (840 pb), and 350 GeV (50 pb) fall outside the y-axis range. Reproduced from Ref. [23].

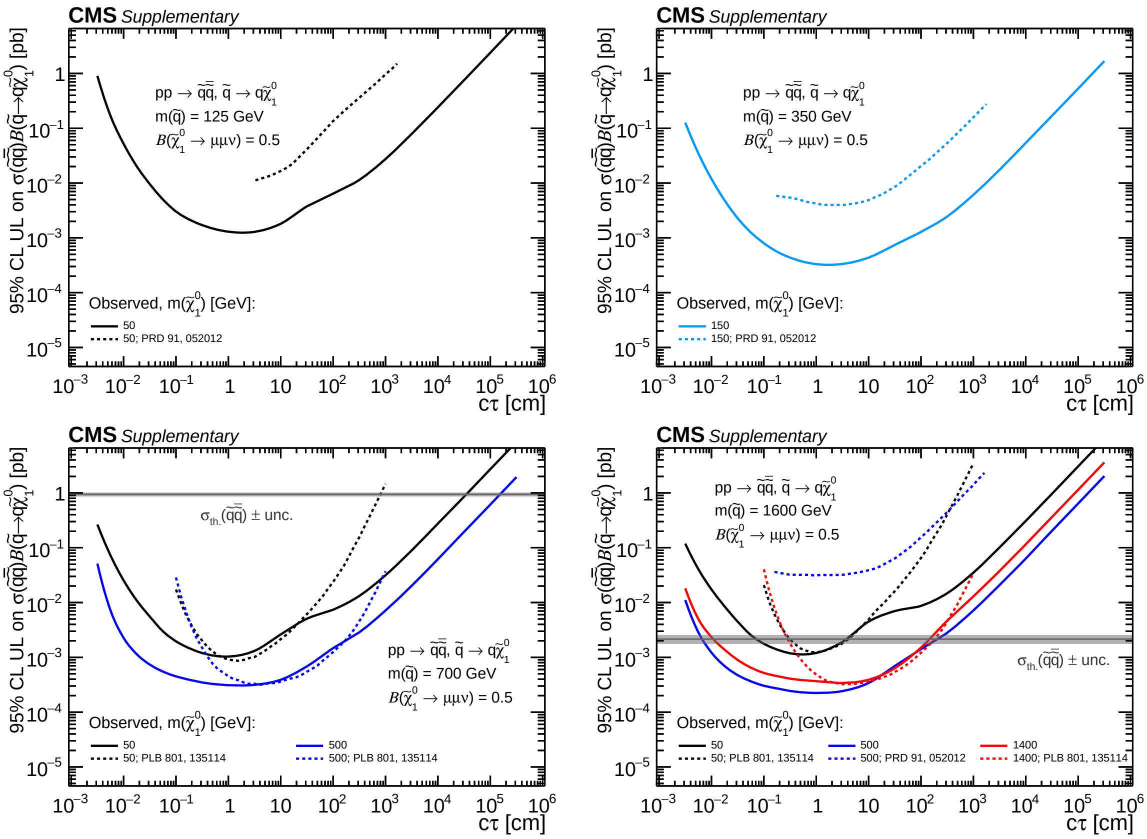

Figure 5.56: Observed 95% CL upper limits on $\sigma(\mathrm{pp} \to \tilde{q}\bar{\tilde{q}})\mathcal{B}(\tilde{q} \to q\tilde{\chi}_1^0)$ as a function of $c\tau(\tilde{\chi}_1^0)$ in the RPV SUSY model, for $\mathcal{B}(\tilde{\chi}_1^0 \to \mu\mu\nu) = 0.5$ and various combinations of $m(\tilde{q})$ and $m(\tilde{\chi}_1^0)$ indicated in the legends. The limits obtained in this analysis (solid) are compared with the corresponding limits (dashed) from the analysis of CMS Run 1 data[57, 58] rescaled by the ratio of $\sigma(\mathrm{pp} \to \tilde{q}\bar{\tilde{q}})$ at 13.6 TeV and 8 TeV, and from the ATLAS search for displaced vertices of oppositely charged leptons in the 2016 dataset [59] rescaled by the ratio of $\sigma(\mathrm{pp} \to \tilde{q}\bar{\tilde{q}})$ at 13.6 TeV and 13 TeV. The grey horizontal lines indicate the theoretical values of the squark-antisquark production cross-sections, with the uncertainties shown as grey shaded bands. The predicted cross sections for $m(\tilde{q}) =$ 125 and 350 GeV are, respectively, 7200 and 50 pb, and fall outside the y-axis range. Reproduced from the supplementary material provided for Ref. [23].

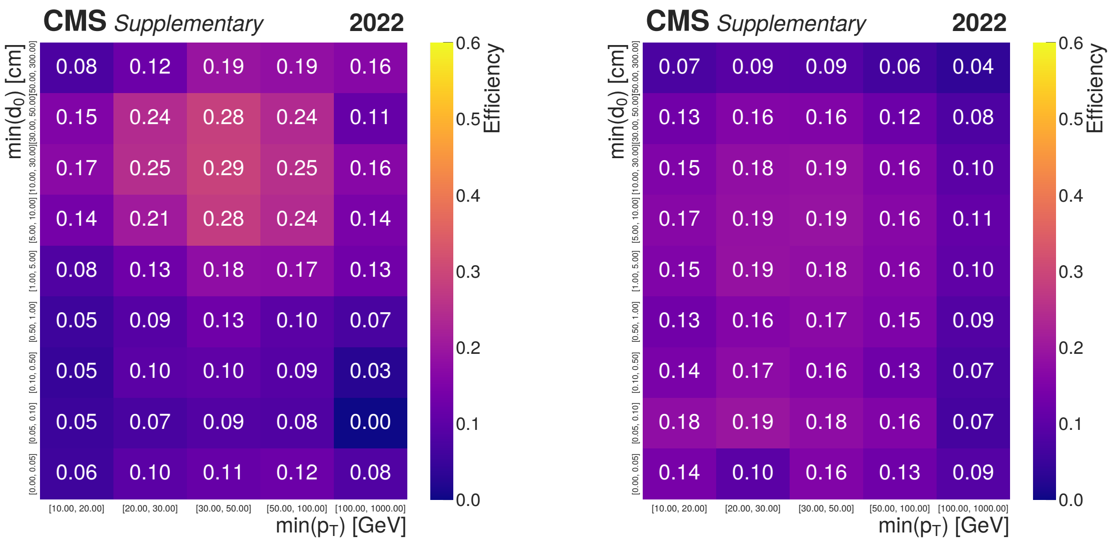

Figure 5.57: Signal efficiencies in the STA-STA dimuon category as a function of the smaller of the two values of generated muon p_T and d_0 in dimuons with L_{xy}^{true} between 70-500 cm in the HAHM (left) and the RPV SUSY (right) signal models. The efficiency in each bin is computed as the ratio of the number of simulated signal dimuons in that bin that pass the trigger requirements and selection criteria to the total number of simulated signal dimuons in that bin and within the geometric acceptance.

Chapter 6

Conclusions

In this thesis, I presented a search for long-lived particles that decay into a pair of muons, conducted using the CMS experiment at the LHC, at the record center of mass energy of $\sqrt{s} = 13.6$ TeV with data collected in 2022, during the first data taking year of Run 3. While we did not find evidence of such particles, or indeed, of any new physics beyond the predictions of the Standard Model, we can certainly claim to have made several remarkable achievements.

Motivated by the successful Run 2 search [22], we identified areas with potential for improvement. The beam spot constraint in the muon track finding at the Level 1 trigger would bias the L1 p_T measurement, causing displaced muon candidates to be reconstructed as low p_T prompt muons that would fail the trigger thresholds. This loss in efficiency was addressed by developing new L1 triggers that used muon track finders without this constraint in the CMS barrel. These triggers were complemented by also including existing L1 triggers with low p_T thresholds. This strategy improved the efficiency in signal by up to 80% at the Level 1 trigger.

To retain this increased efficiency at the High Level Trigger, new HLT paths were developed that lower the p_T thresholds from 23 GeV used in 2018, down to 10 GeV in Run 3. Elements of the offline selection used to suppress the backgrounds, such as displacement requirements and different types of muon reconstruction, were implemented in the trigger design itself. This resulted in a suite of complementary triggers that benefit from different muon reconstruction methods, using L2 muons to capture muons with large displacements while also benefiting from the more precise L3 muons to reject prompt backgrounds, to vastly increase the trigger efficiency for displaced muons by up to 4 times. These triggers allow us to cover a wide range of LLP decay lengths spanning the full detector at only a modest increase in the trigger rate. The triggers are largely model-independent, thus retaining sensitivity to a variety of LLP models.

However, despite their remarkable ability to suppress prompt muons, the lowered p_T thresholds in the triggers brought in more backgrounds with misreconstructed low-mass dimuons. To address these backgrounds a new track-based isolation for STA muons was developed and deployed for the first time in CMS. It is an effective variable for discriminating signal from the background, in the case where the signal muon are expected to be isolated, while also providing us with a background-rich control region to further study and understand events featuring non-isolated displaced dimuons. More generally, a purely data-driven background estimation method was developed in the Run 2 search, which was highly effective even in the Run 3 search. This is particularly important as the Monte Carlo simulation of background is not able to accurately reproduce the effects of misidentification and misreconstruction.

These efforts resulted in the first published BSM search using 13.6 TeV data among the four major LHC collaborations [23]. The results were used to obtain limits in terms of two widely used benchmark models featuring displaced dimuons, the Hidden Abelian Higgs Model and an R-parity violating SUSY model (which features non-resonant dimuon vertices), with only minimal modifications to the event selection, demonstrating its robustness. These limits could match the results set by the Run 2 search, constraining a large volume of model parameter space across several orders of magnitude in LLP lifetime, with only 40% as much data, and even exceed those limits in some regions. Combining the results of Run 2 and Run 3 improves the existing limits by a factor of 2, setting what are now the most stringent limits on the Higgs to LLP branching ratio, $\mathcal{B}(H \to XX)$, over almost all of the probed range of lifetimes. Finally, generator level efficiency maps that approximate the reconstruction-level efficiencies of the analysis are provided, as well as instructions to facilitate the reinterpretation of the results in the framework of other models that predict long-lived particles decaying to final states with a pair of oppositely charged muons.

While my contributions to the search were significant, as described in this thesis, particularly on the development of the triggers and the analysis of the STA-STA dimuon category, this search was ultimately a complete and cohesive team effort. This work has leveraged some of the best capabilities of the LHC and the CMS experiment to provide world-leading sensitivities to long-lived particles and BSM physics. We continue to push the boundaries of both physics theory and technology to provide the best possible understanding of the universe.

References

[1] Planck Collaboration. “Planck 2013 results. I. Overview of products and scientific results”. *Astronomy and Astrophysics* 571.A1 (2014). DOI: 10.1051/0004-6361/201321529.

[2] N. Jarosik et al. “Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Sky maps, systematic errors, and basic results”. *ApJS* 192.2 (2011). DOI: 10.1088/0067-0049/192/2/14.

[3] Particle Data Group. “Review of Particle Physics”. *Progress of Theoretical and Experimental Physics* Volume 2022, Issue 8 (Aug. 2022). DOI: 10.1093/ptep/ptac097.

[4] Oliver Sim Brüning et al. *LHC Design Report Vol 1 : The LHC main ring.* CERN Yellow Reports: Monographs. Geneva: CERN, 2004. DOI: 10.5170/CERN-2004-003-V-1.

[5] Lyndon Evans and Philip Bryant. “LHC Machine”. *Journal of Instrumentation* 3.08 (Aug. 2008), S08001. DOI: 10.1088/1748-0221/3/08/S08001.

[6] ATLAS Collaboration. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. *Physics Letters B* 716.1 (Sept. 2012), pp. 1–29. DOI: 10.1016/j.physletb.2012.08.020.

[7] CMS Collaboration. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. *Physics Letters B* 716.1 (2012), pp. 1–29. DOI: https://doi.org/10.1016/j.physletb.2012.08.020.

[8] CMS Collaboration. “Measurements of Higgs boson properties in the diphoton decay channel in proton-proton collisions at $\sqrt{s} = 13$ TeV”. *Journal of High Energy Physics* 2018.11 (Nov. 2018). DOI: 10.1007/jhep11(2018)185.

[9] CMS Collaboration. “Measurements of properties of the Higgs boson decaying into the four-lepton final state in pp collisions at $\sqrt{s} = 13$ TeV”. *Journal of High Energy Physics* 2017.11 (Nov. 2017). DOI: 10.1007/jhep11(2017)047.

[10] CMS Collaboration. “Combined measurements of Higgs boson couplings in proton–proton collisions at $\sqrt{s} = 13$ TeV”. *The European Physical Journal C* 79.5 (May 2019). DOI: 10.1140/epjc/s10052-019-6909-y.

[11] CMS Collaboration. “Observation of Higgs Boson Decay to Bottom Quarks”. *Phys. Rev. Lett.* 121 (12 Sept. 2018), p. 121801. DOI: 10.1103/PhysRevLett.121.121801.

[12] CMS Collaboration. “Observation of $t\bar{t}H$ Production”. *Phys. Rev. Lett.* 120 (23 June 2018), p. 231801. DOI: 10.1103/PhysRevLett.120.231801.

[13] CMS Collaboration. “Observation of the Higgs boson decay to a pair of τ leptons with the CMS detector”. *Physics Letters B* 779 (Apr. 2018), pp. 283–316. DOI: 10.1016/j.physletb.2018.02.004.

[14] The Muon g-2 Collaboration. “Measurement of the Positive Muon Anomalous Magnetic Moment to 0.20 ppm”. *Physical Review Letters* 131.16 (Oct. 2023). DOI: 10.1103/physrevlett.131.161802.

[15] SNO Collarboration. “Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory”. *Physical Review Letters* 89.1 (June 2002). DOI: 10.1103/physrevlett.89.011301.

[16] Stephen P. Martin. “A Supersymmetry Primer”. *Advanced Series on Directions in High Energy Physics*. World Scientific, July 1998, pp. 1–98. DOI: 10.1142/9789812839657_0001.

[17] Anthony Zee. *Quantum field theory in a nutshell*. Vol. 7. Princeton University Press, 2010.

[18] Bob Holdom. “Two U(1)’s and ϵ charge shifts”. *Physics Letters B* 166.2 (1986), pp. 196–198. DOI: https://doi.org/10.1016/0370-2693(86)91377-8.

[19] Maxim Pospelov, Adam Ritz, and Mikhail Voloshin. “Secluded WIMP dark matter”. *Physics Letters B* 662.1 (Apr. 2008), pp. 53–61. DOI: 10.1016/j.physletb.2008.02.052.

[20] LHCb Collaboration. “Search for $A' \to \mu^+\mu^-$ Decays”. *Phys. Rev. Lett.* 124 (4 Jan. 2020), p. 041801. DOI: 10.1103/PhysRevLett.124.041801.

[21] CMS Collaboration. “Search for a Narrow Resonance Lighter than 200 GeV Decaying to a Pair of Muons in Proton-Proton Collisions at $\sqrt{s} = 13$ TeV”. *Phys. Rev. Lett.* 124 (13 Apr. 2020), p. 131802. DOI: 10.1103/PhysRevLett.124.131802.

[22] CMS Collaboration. “Search for long-lived particles decaying to a pair of muons in proton-proton collisions at $\sqrt{s} = 13$ TeV”. *Journal of High Energy Physics* 2023.5 (May 2023). DOI: 10.1007/jhep05(2023)228.

[23] The CMS Collaboration. “Search for long-lived particles decaying to final states with a pair of muons in proton-proton collisions at $\sqrt{s} = 13.6$ TeV”. *Journal of High Energy Physics* 2024.5 (May 2024). Supplementary material : https://cms-results.web.cern.ch/cms-results/public-results/publications/EXO-23-014/. HEPData Record : https://www.hepdata.net/record/ins2760892. DOI: 10.1007/jhep05(2024)047.

[24] CMS Collaboration. “Dark sector searches with the CMS experiment” (2024). DOI: https://arxiv.org/abs/2405.13778. arXiv: 2405.13778 (hep-ex).

[25] J. Dumarchez E. Augé and J. Trân Thanh Vân, eds. *Search for long-lived particles decaying to a pair of muons in proton-proton collisions at $\sqrt{s} = 13$ TeV*. Rencontres de Moriond 2022: Proceedings of the ElectroWeak Session, La Thuile. ARISF, 2022. DOI: https://doi.org/10.58027/9qzt-x624.

[26] Mangesh Sonawane on behalf of the CMS Collaboration. *Exploring long-lived particles decaying into Displaced Dimuons at $\sqrt{s} = 13.6$ TeV*. 2024. URL: https://indico.cern.ch/event/1336414/contributions/5783068/.

[27] Mangesh Sonawane on behalf of the CMS Collaboration. *Exploring long-lived particles decaying into Displaced Dimuons at $\sqrt{s} = 13.6$ TeV*. 2024. URL: https://agenda.infn.it/event/38205/timetable/?view=standard#73-exploring-long-lived-partic.

[28] Mangesh Sonawane on behalf of the CMS Collaboration. *Exploring long-lived particles decaying into Displaced Dimuons at $\sqrt{s} = 13.6$ TeV*. 2024. URL: https://indico.cern.ch/event/1381368/contributions/5963316/.

[29] Sebastian Templ. “Search for displaced dimuons in proton-proton collisions at $\sqrt{s} = 13$ TeV with the CMS muon system”. PhD thesis. T.U. Wien, 2021. URL: https://cds.cern.ch/record/2801732.

[30] Alberto Escalante del Valle. “Measurement of associated Z+charm production and Search for W’ bosons in the CMS experiment at the LHC”. PhD thesis. Universidad Complutense de Madrid, 2017. URL: https://cds.cern.ch/record/2316144.

[31] Jakob Schwichtenberg. *Physics from Symmetry*. Undergraduate Lecture Notes in Physics. Springer, 2018. URL: http://www.stat.ucla.edu/~ywu/ps.pdf.

[32] Antonio Pich. *The Standard Model of Electroweak Interactions*. 2012. arXiv: 1201.0537 (hep-ph). URL: https://arxiv.org/abs/1201.0537.

[33] Matt Strassler. *Of Particular Significance*. URL: https://profmattstrassler.com/articles-and-posts/particle-physics-basics/.

[34] *Standard Model*. URL: https://en.wikipedia.org/wiki/Standard_Model.

[35] Planck Collaboration. "Planck 2015 results. XIII. Cosmological parameters." *Astronomy and Astrophysics* 594.A13 (2016). DOI: 10.1051/0004-6361/201525830.

[36] Douglas Clowe, Anthony Gonzalez, and Maxim Markevitch. "Weak-Lensing Mass Reconstruction of the Interacting Cluster 1E 0657–558: Direct Evidence for the Existence of Dark Matter." *The Astrophysical Journal* 604.2 (2004), p. 596. DOI: 10.1051/0004-6361/201525830.

[37] M. Markevitch et al. "Direct Constraints on the Dark Matter Self-Interaction Cross Section from the Merging Galaxy Cluster 1E 0657–56". *Astronomy and Astrophysics* 606.2 (2004), p. 819. DOI: 10.1086/383178.

[38] Stephen P. Martin. "A Supersymmetry Primer". *Perspectives on Supersymmetry*. World Scientific, July 1998, pp. 1–98. DOI: 10.1142/9789812839657_0001.

[39] Lawrence Lee et al. "Collider searches for long-lived particles beyond the Standard Model". *Progress in Particle and Nuclear Physics* 106 (2019), pp. 210–255. DOI: https://doi.org/10.1016/j.ppnp.2019.02.006.

[40] Juliette Alimena and et. al. Beacham James. "Searching for long-lived particles beyond the Standard Model at the Large Hadron Collider". *Journal of Physics G: Nuclear and Particle Physics* 47.9 (Sept. 2020), p. 090501. DOI: 10.1088/1361-6471/ab4574.

[41] M.J. Ramsey-Musolf. "Long Lived Particles at the LHC". Oct. 2017. URL: https://indico.cern.ch/event/649760/contributions/2689321/attachments/1542431/2419468/MJRM_LLP_Trieste_17.pdf.

[42] David Curtin et al. "Illuminating dark photons with high-energy colliders". *Journal of High Energy Physics* 2015.2 (Feb. 2015). DOI: 10.1007/jhep02(2015)157.

[43] James D. Wells. *How to Find a Hidden World at the Large Hadron Collider*. 2008. arXiv: 0803.1243 (hep-ph). URL: https://arxiv.org/abs/0803.1243.

[44] Z. Chacko, Hock-Seng Goh, and Roni Harnik. "The Twin Higgs : Natural Electroweak Breaking from a Mirror Symmetry". *Physical Review Letters* 96.23 (June 2006). DOI: 10.1103/physrevlett.96.231802.

[45] Gustavo Burdman et al. "Folded supersymmetry and the LEP paradox". *Journal of High Energy Physics* 2007.02 (Feb. 2007), p. 009. DOI: 10.1088/1126-6708/2007/02/009.

[46] Nathaniel Craig and Kiel Howe. "Doubling down on naturalness with a supersymmetric twin Higgs". *Journal of High Energy Physics* 2014.3 (Mar. 2014). DOI: 10.1007/jhep03(2014)140.

[47] Nathaniel Craig, Simon Knapen, and Pietro Longhi. "Neutral Naturalness from Orbifold Higgs Models". *Physical Review Letters* 114.6 (Feb. 2015). DOI: 10.1103/physrevlett.114.061803.

[48] Gustavo Burdman et al. "Colorless top partners, a 125 GeV Higgs boson, and the limits on naturalness". *Physical Review D* 91.5 (Mar. 2015). DOI: 10.1103/physrevd.91.055007.

[49] Vanda Silveira and A. Zee. “Scalar Phantoms”. *Physics Letters B* 161.1 (1985), pp. 136–140. DOI: https://doi.org/10.1016/0370-2693(85)90624-0.

[50] Daniel Feldman, Boris Kors, and Pran Nath. “Extra-weakly interacting dark matter”. *Physical Review D* 75.2 (Jan. 2007). DOI: 10.1103/physrevd.75.023503.

[51] David E Morrissey and Michael J Ramsey-Musolf. “Electroweak baryogenesis”. *New Journal of Physics* 14.12 (Dec. 2012), p. 125003. DOI: 10.1088/1367-2630/14/12/125003.

[52] Matthew J. Strassler and Kathryn M. Zurek. “Echoes of a hidden valley at hadron colliders”. *Physics Letters B* 651.5–6 (Aug. 2007), pp. 374–379. DOI: 10.1016/j.physletb.2007.06.055.

[53] David Curtin et al. “Exotic decays of the 125 GeV Higgs boson”. *Physical Review D* 90.7 (Oct. 2014). DOI: 10.1103/physrevd.90.075004.

[54] CMS Collaboration. “Search for long-lived particles decaying into muon pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV collected with a dedicated high-rate data stream”. *Journal of High Energy Physics* 2022.4 (Apr. 2022). DOI: 10.1007/jhep04(2022)062.

[55] Daniel Dercks et al. “R-parity violation at the LHC”. *The European Physical Journal C* 77.12 (Dec. 2017). DOI: 10.1140/epjc/s10052-017-5414-4.

[56] Herbi Dreiner. “An introduction to explicitR-parity violation”. *Pramana* 51.1–2 (July 1998), pp. 123–133. DOI: 10.1007/bf02827485.

[57] CMS Collaboration. “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”. *Physical Review D* 91.5 (Mar. 2015). DOI: 10.1103/physrevd.91.052012.

[58] CMS Collaboration. “Search for long-lived particles that decay into final states containing two muons, reconstructed using only the CMS muon chambers”. *CMS Physics Analysis Summary CMS-PAS-EXO-14-012* (2015).

[59] CMS Collaboration. “Search for displaced vertices of oppositely charged leptons from decays of long-lived particles in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. *Physics Letters B* 801 (2020), p. 135114. DOI: https://doi.org/10.1016/j.physletb.2019.135114.

[60] RPV SUSY cross sections. URL: https://twiki.cern.ch/twiki/bin/view/LHCPhysics/SUSYCrossSections#Cross_sections_for_various_subpr.

[61] Matthew J. Strassler and Kathryn M. Zurek. “Discovering the Higgs through highly-displaced vertices”. *Physics Letters B* 661.4 (Mar. 2008), pp. 263–267. DOI: 10.1016/j.physletb.2008.02.008.

[62] The CMS Collaboration. “The CMS experiment at the CERN LHC”. *Journal of Instrumentation* 3.08 (Aug. 2008), S08004. DOI: 10.1088/1748-0221/3/08/S08004.

[63] CMS Collaboration. *CMS Physics: Technical Design Report Volume 1: Detector Performance and Software*. Technical design report. CMS. Geneva: CERN, 2006. URL: https://cds.cern.ch/record/922757.

[64] Oliver Sim Brüning et al. *LHC Design Report Vol 2 : The LHC infrastructure and general services*. CERN Yellow Reports: Monographs. Geneva: CERN, 2004. DOI: 10.5170/CERN-2004-003-V-2.

[65] Michael Benedikt et al. *LHC Design Report Vol 3 : The LHC injector chain*. CERN Yellow Reports: Monographs. Geneva: CERN, 2004. DOI: 10.5170/CERN-2004-003-V-3.

[66] The ATLAS Collaboration. "The ATLAS Experiment at the CERN Large Hadron Collider". *Journal of Instrumentation* 3.08 (Aug. 2008), S08003. DOI: 10.1088/1748-0221/3/08/S08003.

[67] The LHCb Collaboration. "The LHCb Detector at the LHC". *Journal of Instrumentation* 3.08 (Aug. 2008), S08005. DOI: 10.1088/1748-0221/3/08/S08005.

[68] The ALICE Collaboration. "The ALICE experiment at the CERN LHC". *Journal of Instrumentation* 3.08 (Aug. 2008), S08002. DOI: 10.1088/1748-0221/3/08/S08002.

[69] Hannes Bartosik and Giovanni Rumolo. *Performance of the LHC injector chain after the upgrade and potential development*. 2022. arXiv: 2203.09202 (physics.acc-ph). URL: https://doi.org/10.48550/arXiv.2203.09202.

[70] Stephane Fartoukh et al. *LHC Configuration and Operational Scenario for Run 3*. Tech. rep. Geneva: CERN, 2021. URL: https://cds.cern.ch/record/2790409.

[71] Jacques Gareyte. "LHC main parameters". *Part. Accel.* 50 (1995), pp. 61–68. URL: https://cds.cern.ch/record/304825.

[72] CMS Collaboration. *Public CMS luminosity information*. 2023. URL: https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults#Summary_proton_proton_collisions.

[73] The computational results presented were obtained using the CLIP cluster. URL: https://clip.science.

[74] CMS Collaboration. "Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV". *Journal of Instrumentation* 15.10 (Oct. 2020), P10017. DOI: 10.1088/1748-0221/15/10/P10017.

[75] CMS Collaboration. "The CMS trigger system". *Journal of Instrumentation* 12.01 (Jan. 2017), P01020. DOI: 10.1088/1748-0221/12/01/P01020.

[76] CMS Collaboration. "Performance of electron reconstruction and selection with the CMS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV". *Journal of Instrumentation* 10.06 (June 2015), P06005. DOI: 10.1088/1748-0221/10/06/P06005.

[77] CMS Collaboration. "Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV". *Journal of Instrumentation* 13.06 (June 2018), P06015. DOI: 10.1088/1748-0221/13/06/P06015.

[78] CMS Collaboration. "Performance of photon reconstruction and identification with the CMS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV". *Journal of Instrumentation* 10.08 (Aug. 2015), P08010. DOI: 10.1088/1748-0221/10/08/P08010.

[79] The CMS Collaboration. "Description and performance of track and primary-vertex reconstruction with the CMS tracker." *Journal of Instrumentation* 9.10 (Oct. 2014), P10009. DOI: 10.1088/1748-0221/9/10/P10009.

[80] CMS. *The CMS magnet project: Technical Design Report*. Technical design report. CMS. Geneva: CERN, 1997. DOI: 10.17181/CERN.6ZU0.V4T9.

[81] The CMS Collaboration. "The CMS detector magnet". *IEEE Transactions on Applied Superconductivity* 10.1 (2000), pp. 389–394. DOI: 10.1109/77.828255.

[82] V Karimäki et al. *The CMS tracker system project: Technical Design Report*. Technical design report. CMS. Geneva: CERN, 1997. URL: https://cds.cern.ch/record/368412.

[83] CMS. *The CMS tracker: addendum to the Technical Design Report*. Technical design report. CMS. Geneva: CERN, 2000. URL: https://cds.cern.ch/record/490194.

[84] The CMS Collaboration. "The CMS Phase-1 Pixel Detector Upgrade". *JINST* 16.02 (2021), P02027. DOI: 10.1088/1748-0221/16/02/P02027. arXiv: 2012.14304.

[85] The CMS collaboration. “Development of the CMS detector for the CERN LHC Run 3”. *Journal of Instrumentation* 19.05 (May 2024), P05064. DOI: 10.1088/1748-0221/19/05/P05064.

[86] The CMS Collaboration. *The CMS electromagnetic calorimeter project: Technical Design Report*. Technical design report. CMS. Geneva: CERN, 1997. URL: https://cds.cern.ch/record/349375.

[87] The CMS Collaboration. *The CMS hadron calorimeter project: Technical Design Report*. Technical design report. CMS. Geneva: CERN, 1997. URL: https://cds.cern.ch/record/357153.

[88] The CMS collaboration. “Performance of the CMS muon trigger system in proton-proton collisions at $\sqrt{s} = 13$ TeV”. *Journal of Instrumentation* 16.07 (July 2021), P07001. DOI: 10.1088/1748-0221/16/07/P07001.

[89] Supratik Mukhopadhyay and et. al. “Studies on the upgrade of the muon system in the forward region of the CMS experiment at LHC with GEMs”. *Journal of Instrumentation* 9 (Jan. 2014). DOI: 10.1088/1748-0221/9/01/C01053.

[90] CMS Collaboration. “Particle-flow reconstruction and global event description with the CMS detector”. *Journal of Instrumentation* 12.10 (Oct. 2017), P10003. DOI: 10.1088/1748-0221/12/10/P10003.

[91] The CMS Collaboration. “Particle-flow reconstruction and global event description with the CMS detector”. *Journal of Instrumentation* 12.10 (Oct. 2017), P10003. DOI: 10.1088/1748-0221/12/10/P10003.

[92] R. Frühwirth. “Application of Kalman filtering to track and vertex fitting”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 262.2 (1987), pp. 444–450. DOI: https://doi.org/10.1016/0168-9002(87)90887-4.

[93] A Tapper and Darin Acosta. *CMS Technical Design Report for the Level-1 Trigger Upgrade*. Tech. rep. CERN, 2013. URL: https://cds.cern.ch/record/1556311.

[94] The CMS Collaboration. *Performance of Muon Reconstruction in the CMS High Level Trigger using pp Collision data at $\sqrt{s} = 13.6$ TeV in 2023*. Tech. rep. CMS, 2024. URL: https://cds.cern.ch/record/2888302.

[95] Costas Foudas et al. “Upgrade of the CMS Barrel Muon Track Finder for HL-LHC featuring a Kalman Filter algorithm and an ATCA Host Processor with Ultrascale+ FPGAs”. *PoS* TWEPP2018 (2019), p. 139. DOI: 10.22323/1.343.0139.

[96] Abhigyan Dasgupta. “Search for Long-Lived Particles Decaying to Displaced Dimuons at 13 TeV and Study of Neutron-Induced Background Hits in the Muon System of the Compact Muon Solenoid.” PhD thesis. UCLA, 2019. URL: https://escholarship.org/uc/item/32s9r84f.

[97] The CMS Collaboration. *Luminosity measurement in proton-proton collisions at $\sqrt{s} = 13.6$ TeV in 2022 at CMS*. Tech. rep. Geneva: CERN, 2024. URL: https://cds.cern.ch/record/2890833.

[98] J. Alwall et al. “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations”. *Journal of High Energy Physics* 2014.7 (July 2014). DOI: 10.1007/jhep07(2014)079.

[99] Jon Butterworth et al. “PDF4LHC recommendations for LHC Run II”. *Journal of Physics G: Nuclear and Particle Physics* 43.2 (Jan. 2016), p. 023001. DOI: 10.1088/0954-3899/43/2/023001.

[100] CMS Collaboration. “Extraction and validation of a new set of CMS pythia8 tunes from underlying-event measurements”. *The European Physical Journal C* 80.1 (Jan. 2020). DOI: 10.1140/epjc/s10052-019-7499-4.

[101] Alexander Karlberg et al. *Ad interim recommendations for the Higgs boson production cross sections at $\sqrt{s} = 13.6$ TeV*. 2024. arXiv: 2402.09955 (hep-ph).

[102] Christian Bierlich et al. *A comprehensive guide to the physics and usage of PYTHIA 8.3*. 2022. arXiv: 2203.11601 (hep-ph).

[103] Wim Beenakker et al. "NNLL-fast: predictions for coloured supersymmetric particle production at the LHC with threshold and Coulomb resummation". *Journal of High Energy Physics* 2016.12 (Dec. 2016). DOI: 10.1007/jhep12(2016)133.

[104] Richard D Ball et al. "The PDF4LHC21 combination of global PDF fits for the LHC Run III". *Journal of Physics G: Nuclear and Particle Physics* 49.8 (July 2022), p. 080501. DOI: 10.1088/1361-6471/ac7216.

[105] The CMS Collaboration. "Description and performance of track and primary-vertex reconstruction with the CMS tracker". *Journal of Instrumentation* 9.10 (Oct. 2014), P10009. DOI: 10.1088/1748-0221/9/10/P10009.

[106] Particle Data Group. "Review of Particle Physics". *Progress of Theoretical and Experimental Physics* 2022.8 (Aug. 2022), p. 083C01. DOI: 10.1093/ptep/ptac097. eprint: https://academic.oup.com/ptep/article-pdf/2022/8/083C01/49175539/ptac097.pdf.

[107] William Arthur Nash. "Search for Long-Lived Particles that Decay to Muon Pairs at 13 TeV and an Algorithm to Improve Resolution in the Muon System of the Compact Muon Solenoid". PhD thesis. UCLA, 2022. URL: https://escholarship.org/uc/item/3gw9h84p.

[108] The LHCb Collaboration. "Identification of beauty and charm quark jets at LHCb". *Journal of Instrumentation* 10.06 (June 2015), P06013–P06013. DOI: 10.1088/1748-0221/10/06/p06013.

[109] The CMS Collaboration. "Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV". *Journal of Instrumentation* 13.05 (May 2018), P05011–P05011. DOI: 10.1088/1748-0221/13/05/p05011.

[110] The CMS Collaboration. "First measurement of the top quark pair production cross section in proton-proton collisions at $\sqrt{s} = 13.6$ TeV". *Journal of High Energy Physics* 2023.8 (Aug. 2023). DOI: 10.1007/jhep08(2023)204.

[111] The CMS Collaboration. "Performance of CMS muon reconstruction in pp collision events at $\sqrt{s} = 7$ TeV". *Journal of Instrumentation* 7.10 (Oct. 2012), P10002–P10002. DOI: 10.1088/1748-0221/7/10/p10002.

[112] M. Cepeda and et al. *Higgs Physics at the HL-LHC and HE-LHC*. 2019. arXiv: 1902.00134 (hep-ph). URL: https://arxiv.org/abs/1902.00134.

[113] A L Read. "Presentation of search results: the CLs technique". *Journal of Physics G: Nuclear and Particle Physics* 28.10 (Sept. 2002), p. 2693. DOI: 10.1088/0954-3899/28/10/313.

[114] Thomas Junk. "Confidence level computation for combining searches with small statistics". *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 434.2–3 (Sept. 1999), pp. 435–443. DOI: 10.1016/s0168-9002(99)00498-2.

[115] CMS Collaboration. *The CMS statistical analysis and combination tool: COMBINE*. 2024. arXiv: 2404.06614 (physics.data-an). URL: https://arxiv.org/abs/2404.06614.

[116] The ATLAS Collaboration, The CMS Collaboration, and The LHC Higgs Combination Group. *Procedure for the LHC Higgs boson search combination in Summer 2011*. Tech. rep. Geneva: CERN, 2011. URL: https://cds.cern.ch/record/1379837.

[117] The CMS Collaboration. "A search for decays of the Higgs boson to invisible particles in events with a top-antitop quark pair or a vector boson in proton-proton collisions at $\sqrt{s} = 13$ TeV". *The European Physical Journal C* 83.10 (Oct. 2023). DOI: 10.1140/epjc/s10052-023-11952-7.

[118] The CMS Collaboration. "Search for invisible decays of the Higgs boson produced via vector boson fusion in proton-proton collisions at $\sqrt{s} = 13$ TeV". *Physical Review D* 105.9 (May 2022). DOI: 10.1103/physrevd.105.092007.

[119] The CMS Collaboration. "Combined measurements of Higgs boson couplings in proton–proton collisions at $\sqrt{s} = 13$ TeV". *The European Physical Journal C* 79.5 (May 2019). DOI: 10.1140/epjc/s10052-019-6909-y.

[120] Eamonn Maguire, Lukas Heinrich, and Graeme Watt. "HEPData: a repository for high energy physics data". *Journal of Physics: Conference Series* 898 (Oct. 2017), p. 102006. DOI: 10.1088/1742-6596/898/10/102006.

[121] E. Noether. "Invariante Variationsprobleme". ger. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* 1918 (1918), pp. 235–257. URL: http://eudml.org/doc/59024.

[122] The CMS Collaboration. "A portrait of the Higgs boson by the CMS experiment ten years after the discovery". *Nature* 607.7917 (July 2022), pp. 60–68. DOI: 10.1038/s41586-022-04892-x.

Appendix A

Quantum Field Theory of the Standard Model

Every physical theory or mathematical framework relies on a minimal body of assumptions - called axioms - that cannot be derived from more fundamental principles. The SM, too, begins with certain assumptions that serve as the starting ingredients, required to be taken as true, to derive a theory that accurately describes experimental results. However, considering the vast breadth and depth of the SM and the accuracy of its predictions, it requires a surprisingly small number of assumptions and 'free parameters' to be input.

The mathematical structure of quantum field theories follows the *Lagrangian formalism*. The notion is that a quantity called the *Lagrangian* encodes the physics in a given system. The Lagrangian, $\mathcal{L}(t)$, is a function of generalized coordinates $q(t)$ and their time derivatives $\dot{q}(t)$, and is used to define the *action*, S, a functional.

$$S\left[q(t)\right] = \int \mathcal{L} dt \tag{A.1}$$

where the Lagrangian $\mathcal{L} = \mathcal{L}\left(q(t), \partial_t q(t)\right)$. In the SM, the Lagrangian describes fields rather than the particles themselves. Thus the Lagrangian density, $\mathscr{L}$, is expressed in terms of fields ϕ by

$$\mathscr{L} = \mathscr{L}\left(\phi(\vec{x}, t), \partial_\mu \phi(\vec{x}, t)\right) \tag{A.2}$$

and is related to the Lagrangian by

$$\mathcal{L} = \int \mathscr{L} d^3x \tag{A.3}$$

The action, therefore, can be written as

$$S = \int d^4x \mathscr{L}(\phi^i, \partial_\mu \phi^i) \tag{A.4}$$

A physical system is defined by equations of motion, which arise when the action S is minimized. These equations are known as 'Euler-Lagrange equations' and describe the dynamics of the theory. Therefore, we seek a Lagrangian that minimizes the action while describing the configuration and interactions of the field components, giving rise to all the known particles and forces (except gravity).

The most fundamental assumptions for the SM Lagrangian are as follows:

- The speed of light, c, is the same in all reference frames - called inertial frames.
- The physical laws are identical in all inertial frames.

The set of all spacetime transformations that respect these symmetry requirements is called the Poincaré group. This symmetry group consists of the Lorentz group, which governs rotations and boosts, as well as translations. In the context of Lagrangian formalism, imposing the invariance of physical laws across inertial reference frames means requiring that the Lagrangian is invariant under transformations of the Poincaré group.

One then needs to identify and exploit the other symmetries of the Lagrangian that describe spin-½ fields (associated with the Dirac equation), spin-1 fields (associated with the Proca equation), and spin-0 fields (associated with the Klein-Gordon equation). The SM can be described by the symmetry group $\mathrm{SU}(3)_c \otimes \mathrm{SU}(2)_L \otimes \mathrm{U}(1)_Y$. Each symmetry in the SM Lagrangian is also associated with a conserved quantity - called a conserved *charge* - in accordance with Noether's theorem [121]. The charges associated with each force in the SM are as follows:

- Electromagnetic force : the *electric* charge.
- Weak force : the *weak hypercharge*.
- Strong force : *colour* charge.

Elementary particles can only participate in an interaction via a given force if they carry the corresponding charge. Conversely, if a particle has zero charge of a given type, it is not affected by the corresponding interaction.

The SU(3) group corresponds to the strong force and is described by Quantum Chromodynamics (QCD), the SU(2) group corresponds to the weak force, and the U(1) group corresponds to the electromagnetic force described by Quantum Electrodynamics (QED). However, it turns out that the electromagnetic and weak forces are aspects of a single, unified force called the *electroweak* force, described by the $\mathrm{SU}(2) \otimes \mathrm{U}(1)$ group structure. This symmetry is spontaneously broken (i.e., SSB) below a certain energy, called the *electroweak scale*, by the Higgs mechanism, giving rise to the mass of the W and Z bosons. The subscripts refer to the charges that mediate the respective interactions: C refers to the color charge, L refers to the fact that only left-handed particles (and right-handed anti-particles) participate in the weak interaction, and Y refers to the weak hypercharge.

A.1 Quantum Electrodynamics and the U(1) Gauge Group

Quantum Electrodynamics (QED) is an Abelian theory described by the U(1) gauge group. It is the quantum and relativistic version of classical electrodynamics, describing how light (photons) and matter (electrically charged fermions) interact. It was the first theory in which the frameworks of quantum mechanics and special relativity were unified, and is capable of predicting phenomena such as the Lamb shift of the energy levels of hydrogen with precision better than a few parts in a billion.

The QED Lagrangian is constructed from Lagrangians describing free spin-½ fields ψ,

$$\mathcal{L}_{Dirac} = \bar{\psi}(i\gamma^{\mu}\partial_{\mu} - m)\psi \tag{A.5}$$

free spin-1 fields A^{μ},

$$\mathcal{L}_{Maxwell} = \partial^{\mu}A^{\nu}\partial_{\mu}A_{\nu} - \partial^{\nu}A^{\mu}\partial_{\nu}A_{\mu} \tag{A.6}$$

and imposing Lorentz invariance as well as invariance under global and local U(1) transformations. Both A.5 and A.6 are readily invariant under global transformations of the form

$$\psi \rightarrow \psi' = e^{i\alpha}\psi \tag{A.7}$$

$$A^{\mu} \rightarrow A'^{\mu} = A^{\mu} + \alpha^{\mu} \tag{A.8}$$

where α is an arbitrary constant.

However, under local transformations, where $\alpha(x)$ is a function of spacetime coordinates x, the Lagrangians are, in fact, not invariant. As it turns out, adding a term coupling the Dirac and Maxwell fields, $A_{\mu}\bar{\psi}\gamma^{\mu}\psi$ solves this issue and makes the Lagrangian invariant under local transformations as well.

$$\mathcal{L}_{QED} = \underbrace{\bar{\psi}(i\gamma^{\mu}\partial_{\mu} - m)\psi}_{\text{free spin-1/2 fields}} + \underbrace{\partial^{\mu}A^{\nu}\partial_{\mu}A_{\nu} - \partial^{\nu}A^{\mu}\partial_{\nu}A_{\mu}}_{\text{free spin-1 fields}} + \underbrace{A_{\mu}\bar{\psi}\gamma^{\mu}\psi}_{\text{interaction term}} \tag{A.9}$$

Introducing the *covariant derivative* $D_{\mu} \equiv \partial_{\mu} + igA_{\mu}$, where g is a parameter describing the strength of the electromagnetic interaction, and simplifying, one obtains

$$\mathcal{L}_{QED} = \bar{\psi}(i\gamma^{\mu}\not{D}_{\mu} - m)\psi - \frac{1}{4}F^{\mu\nu}F_{\mu\nu} \tag{A.10}$$

where $F^{\mu\nu} = \partial^{\mu}A^{\nu} - \partial^{\nu}A^{\mu}$, and Feyman's slash notation ($\not{D} \equiv D^{\mu}\gamma_{\mu}$) is adopted to make the equation compact.

Thus, we arrive at the Lagrangian for the quantum field theory of electrodynamics (QED). It accurately describes fermions (spin-½ particles) and their interactions via the exchange of the photon (spin-1 mediator).

A.2 The Strong Interaction and the SU(3) Gauge Group

One can repeat the same exercise, beginning with the symmetries described by the SU(3) gauge group. Consider a Lagrangian constructed from three fermion fields q_i, that is globally invariant under SU(3) transformations of the form

$$Q \to Q' = e^{iT^a \theta_a(x)} Q \tag{A.11}$$

where Q is a triplet of spin-½ fields $Q \equiv \begin{pmatrix} q_1 \\ q_2 \\ q_3 \end{pmatrix}$, and the eight generators of the SU(3) group are $T^a = \frac{\lambda^a}{2}$, and λ^a are conventionally chosen to be the 3×3 Gell-Mann matrices.

The resulting Lagrangian is globally invariant but not locally SU(3) invariant

$$\mathcal{L}_{\mathrm{SU(3)}} = \bar{Q}(i\partial_\mu \gamma^\mu - m)Q \tag{A.12}$$

To obtain a locally invariant SU(3) Lagrangian, we can introduce new spin-1 gauge boson fields, G_a^μ, corresponding to the gluons, which transform as follows :

$$G_a^\mu \to G_a'^\mu = G_a^\mu - \frac{1}{g_s}\partial^\mu \theta_a - f^{abc}\theta_b G_c^\mu \tag{A.13}$$

where g_s determines the strength of the strong interactions, and f^{abc} are the SU(3) structure constants, described by :

$$[T^a, T^b] = i f^{abc} T^c \tag{A.14}$$

These spin-1 fields, along with coupling terms between these and the spin-$\frac{1}{2}$ fields, to obtain a Lagrangian, $\mathcal{L}_{\mathrm{QCD}}$, that accurately describes strongly interacting particles in the framework of Quantum Chromodynamics :

$$\mathcal{L}_{\mathrm{QCD}} = \underbrace{-\frac{1}{4} G^a_{\mu\nu} G_a^{\mu\nu}}_{\text{free spin-1 fields}} + \underbrace{\bar{Q}(i\partial_\mu \gamma^\mu - m)Q}_{\text{free spin-½ fields}} - \underbrace{g_s \bar{Q}\gamma_\mu G_a^\mu T^a Q}_{\text{interaction terms}} \tag{A.15}$$

where $G_a^{\mu\nu}$ is the field strength, with the following gauge invariant form :

$$G_a^{\mu\nu} = \partial^\mu G_a^\nu - \partial^\nu G_a^\mu - g_s f^{abc} G_b^\mu G_c^\nu \tag{A.16}$$

The conserved charges arising from the global SU(3) symmetry, according to Noether's theorem, are the three color charges, and are represented as 'red' (r), 'green' (g), and 'blue' (b). The fermionic triplets Q are thus the quarks. Particles that do not interact via the strong force are SU(3) singlets and have zero color charge.

The non-Abelian nature of the generators of the SU(3) matrices give rise to cubic and quartic self-interactions of the gluon field and result in two important properties in the phenomenology of strong interactions :

- **Color confinement**: Quarks cannot be observed individually but always in color-neutral bound states. *Baryons* are bound states of 3 quarks, such as the proton (uud) or the neutron (udd), while *mesons* are bound states of a quark and an anti-quark ($q\bar{q}$), such that the overall state is color-neutral. Recent discoveries have demonstrated the existence of rarer bound states of four or five quarks, called *tetraquarks* ($q\bar{q}q\bar{q}$) or *pentaquarks* ($qqqq\bar{q}$) respectively. This phenomenon is called *color confinement*.

  Individual quarks produced in high energy particle collisions rapidly produce quarks and anti-quarks from the vacuum to form a color-neutral state – a process called *hadronization*.

- **Asymptotic Freedom**: Due to color confinement, free quarks or gluons cannot be observed over macroscopically detectable distances. However, due to the uncertainty principle, quarks may be treated as free particles over extremely short distances (or correspondingly, at very high energies), as the strong coupling constant decreases with an increases momentum transfer scale. This phenomenon is known as the *asymptotic freedom* and allows us to perform perturbative calculations at high energies. In contrast, perturbative calculations are possible for the other interactions only in the low-energy domain.

A.3 The Weak Interaction and the SU(2) Gauge Group

In a manner very much similar to what is described in Section A.1, one can obtain a Lagrangian by adding two copies of the $\mathcal{L}_{Dirac}$,

$$\mathcal{L}_{\psi_1+\psi_2} = i\bar{\Psi}\gamma^\mu\partial_\mu\Psi \tag{A.17}$$

where Ψ is a doublet, $\Psi \equiv \binom{\psi_1}{\psi_2}$. This Lagrangian is invariant under global SU(2) transformations,

$$\Psi \rightarrow \Psi' = e^{ia_i\frac{\sigma_i}{2}}\Psi \tag{A.18}$$

where the a_i are arbitrary constants, and $\sigma_i (i = 1, 2, 3)$, are the Pauli matrices. $\sigma_i/2$ are then the generators of the SU(2) group.

This Lagrangian is not locally SU(2) symmetric, however, and needs terms to restore the symmetry, giving rise to three new fields. The extra term needs to be $i\bar{\Psi}\gamma_\mu W^i_\mu \sigma_i \Psi$, and describes the interaction between the two spin ½ fields $\begin{pmatrix} \psi_i \\ \psi_2 \end{pmatrix}$, and the three new spin-1 fields W^a, defined by

$$W^a_{\mu\nu} = \partial_\mu W^a_\nu - \partial_\nu W^a_\mu + g\varepsilon^{ijk} W^b_\mu W^c_\nu \tag{A.19}$$

where the index a runs over the generators of the SU(2) group, g is the gauge coupling constant for the SU(2) interaction

Putting together the components and simplifying, we obtain the Lagrangian for the weak interaction as

$$\mathcal{L}_{Weak} = \underbrace{i\bar{\Psi}\gamma^\mu \partial_\mu \Psi}_{\text{free spin-1/2 field}} \underbrace{- \frac{1}{4}(W^{\mu\nu})_i (W_{\mu\nu})_i}_{\text{free spin-1 fields}} + \underbrace{\bar{\Psi}\gamma^\mu \frac{\sigma_i}{2}(W_\mu)_i \Psi}_{\text{interaction term}} \tag{A.20}$$

where k are the field strength tensors.

Naively, we would associate the spin-½ fields as the fermions and the spin-1 fields as the $W^\pm$ and Z bosons. However, experimental observations demonstrate the following unique properties of weak interactions :

- Parity is violated in weak interactions: only left-handed fermions and right-handed anti-fermions participate in weak interactions. The left and right-handedness of fermions is defined as the projections of the Dirac fermion spinors :

  $$\psi_{L,R} = \frac{1 \mp \gamma^5}{2}\psi \tag{A.21}$$

  Introducing a mass term for the fermionic fields Ψ in Equation A.20 would explicitly mix the left and right-handed components, which would contradict observations.

- The bosons that mediate the weak interaction are observed to be massive. The introduction of a mass term for the mass terms would violate the SU(2) symmetry.

It is possible to formulate a Lagrangian that does satisfy the observations, and results in what is known as the *electroweak unification*, while the mass of the bosons is generated by spontaneous symmetry breaking (SSB), described by the **Higgs mechanism**.

A.4 Electroweak Unification

In the electroweak formulation, the electromagnetic and weak interactions are combined in a unified description. In this formulation, the left-handed fermions, L transform as

$SU(2)_L \otimes U(1)_Y$ doublets, and right-handed fermions, R, transform as $U(1)_Y$ singlets, due to the fact that only left-handed fermions (and right-handed anti-fermions) participate in the weak interaction. We can represent the left- and right-handed fermions, and their corresponding transformations, as follows:

$$L = \begin{pmatrix} e^- \\ \nu_e \end{pmatrix}_L, \begin{pmatrix} u \\ d' \end{pmatrix}_L, \ldots \qquad L \to L' = e^{i(Y\beta(x)+\frac{\sigma_i}{2}\alpha_i(x))}L \tag{A.22}$$

$$R = e_R^-, u_R, d_R, \ldots \qquad R \to R' = e^{iY\beta(x)}R \tag{A.23}$$

where Y is the weak hypercharge. The experimental observations of CP violation and quark flavor mixing can be included in the theory if the weak eigenstates d', s', b' are not the same as the mass eigenstates d, s, b, but are related by the Cabibbo-Kobayashi-Maskawa (CKM) unitary matrix, whose elements must be experimentally determined.

The resulting Lagrangian $\mathcal{L}_{EW}$ is invariant under local $SU(2)_L \otimes U(1)_Y$ transformation, and requires the introduction of 4 spin-1 fields.

$$\mathcal{L}_{EW} = -\frac{1}{4}W_a^{\mu\nu}W_{\mu\nu}^a - \frac{1}{4}B^{\mu\nu}B_{\mu\nu} + \sum_{\psi_i=L,R} i\bar{L}_i \not{D} L_i + i\bar{R}_i \not{D} R_i \tag{A.24}$$

where $\not{D} = \gamma^\mu D_\mu$, with D_μ as the partial derivative

$$D_\mu L = [\partial_\mu + ig\tilde{W}_\mu + ig'YB_\mu]L \tag{A.25}$$

$$D_\mu R = [\partial_\mu + ig'YB_\mu]R \tag{A.26}$$

Here $\tilde{W}_\mu$ can be expressed in terms of the three new $SU(2)_L$ invariant massless spin-1 fields $W_\mu^i (i = 1, 2, 3)$ as :

$$\tilde{W}_\mu = \sigma_i W_\mu^i = \frac{1}{2}\begin{bmatrix} W_\mu^3 & W_\mu^1 - iW_\mu^2 \\ W_\mu^1 + iW_\mu^2 & -W_\mu^3 \end{bmatrix} \tag{A.27}$$

and B_μ is another massless spin-1 field that is invariant under local U(1) transformations. Their corresponding field strengths are given by

$$W_{\mu\nu}^i = \partial_\mu W_\mu^a - \partial_\nu W_\nu^a - g\epsilon^{ijk}W_\mu^j W_\nu^k \tag{A.28}$$

$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu \tag{A.29}$$

.

As the Pauli matrices do not commute, $[\sigma^i, \sigma^j] = i\epsilon^{ijk}\sigma^k$, the kinetic term for the SU(2) spin-1 free fields, $-\frac{1}{4}W_a^{\mu\nu}W_{\mu\nu}^a$, gives rise to cubic and quartic self-interactions of the W_μ^i fields.

These new gauge fields are not the known physical $W^{\pm}$, Z bosons, and the photon. However, we can obtain the neutral currents by defining a linear transformation between the neutral fields W_μ^3 and B_μ, and the physical Z_μ and A_μ fields.

$$\begin{bmatrix} W_\mu^3 \\ B_\mu \end{bmatrix} = \begin{bmatrix} \cos\theta_W & \sin\theta_W \\ -\sin\theta_W & \cos\theta_W \end{bmatrix} \begin{bmatrix} Z_\mu \\ A_\mu \end{bmatrix} \tag{A.30}$$

where θ_W is a free parameter in the theory called the weak mixing angle, or the Weinberg angle, that can be measured with high experimental precision from weak decay rates.

Similarly, We can obtain the charged currents and associate the physical $W^{\pm}$ bosons with the W^1 and W^2 fields via

$$W^{\pm} = \frac{W^1 \mp iW^2}{\sqrt{2}} \tag{A.31}$$

The resulting fields are still massless due to gauge invariance, however, and we still cannot introduce a fermionic mass term into Equation A.26 due to the left and right-handed terms having different properties. This problem is addressed in the following section.

A.5 Spontaneous Symmetry Breaking and the Higgs Mechanism

While $\mathcal{L}_{EW}$ is invariant under local gauge invariance, it does not account for the measured masses of the W bosons, $M_W = 80.3692 \pm 0,0133$ GeV, and the Z boson, $M_Z = 91.1880 \pm 0.0020$ GeV, or the measured masses of the fermions [106]; only the photon is truly massless.

This problem can be solved by introducing a new complex scalar doublet field Φ

$$\Phi = \begin{bmatrix} \phi^+ \\ \phi^0 \end{bmatrix} \tag{A.32}$$

We can add a new $\mathrm{SU}(2)_L \otimes \mathrm{U}(1)_Y$ gauge invariant term, $\mathcal{L}_{\text{Higgs}}$, to the electroweak Lagrangian,

$$\mathcal{L}_{\text{Higgs}} = \underbrace{(\partial^\mu\Phi)^\dagger(\partial_\mu\Phi)}_{\text{Kinetic term}} - \underbrace{\mu^2(\Phi^\dagger\Phi) - \lambda(\Phi^\dagger\Phi)^2}_{\text{Higgs Potential term, } V(\Phi)} \tag{A.33}$$

The combined electroweak Lagrangian is then

$$\mathcal{L}_{EW} = [(\partial_\mu - ig'\sigma_i W_\mu^i - igB_\mu)\Phi^\dagger] \cdot [(\partial^\mu - ig'\sigma_i W^{\mu i} + igB^\mu)\Phi] - V(\Phi) \tag{A.34}$$

In the Higgs potential term, $V(\Phi)$, the parameters μ and λ are assumed to be temperature-dependent and to have evolved with the universe. If the self-interaction term is negative $\mu^2 < 0$, and the quartic coupling term is positive, $\lambda > 0$, this results in a potential with a degenerate set of states with a non-zero minimum energy, called the *vacuum expectation value* (VEV) of the Higgs potential, famously knows as the 'Mexican hat' potential, as shown in Figure A.1.

$$\langle 0| \Phi |0\rangle = \sqrt{\frac{-\mu}{2}} e^{i\varphi} = \frac{v}{2} e^{i\varphi} \tag{A.35}$$

where the VEV, $v \approx 246$ GeV, and can be determined from measurements of the muon lifetime. This value sets the energy scale at which the Higgs mechanism played in breaking the electroweak symmetry in the early universe, and is therefore known as the electroweak energy scale. To illustrate this, in the early era of the universe in which the temperature at a time t, $T(t) >> v$, the SM was invariant under $\mathrm{SU}(2)_L \otimes \mathrm{U}(1)_Y$ transformations, and $V(\Phi)$ had a degenerate set of energy states with minimum energy zero, $< \Phi >= 0$. However, as the universe cooled into an era where $T(t) << v$, these symmetries would spontaneously break, and the Higgs potential $V(\Phi)$ would acquire a non-zero VEV.

We have the freedom to choose a Higgs potential such that $H = \Phi - v/\sqrt{2}$, giving us the Higgs field

$$H = \frac{1}{\sqrt{2}} e^{i\sigma_i \Theta_i} \begin{bmatrix} 0 \\ v \end{bmatrix} \tag{A.36}$$

This freedom is caused by the presence of three non-physical massless Goldstone bosons Θ_i. Choosing a particular unitary gauge where $\Theta_i = 0$ breaks the $\mathrm{SU}(2)_L \otimes \mathrm{U}(1)_Y$ symmetry of the Lagrangian, removing the Goldstone bosons, and leaves the SM Lagrangian unbroken under the $\mathrm{U}(1)_{EM}$ symmetry.

The W and Z bosons acquire mass and interactions with the Higgs field through the kinetic term in $\mathcal{L}_{EW}$,

$$\mathcal{L}_{EW} \supset (v + H)^2 [\frac{g^2}{4} W_\mu^\dagger W^\mu + \frac{g^2}{8\cos^2\theta_W} Z_\mu Z^\mu] \tag{A.37}$$

The masses of the bosons are them given by $\frac{1}{2} M_W W_\mu^\dagger W^\mu$ and $\frac{1}{2} M_Z Z_\mu Z^\mu$ respectively. This leads to the prediction that

$$\cos\theta_W = M_W / M_Z \tag{A.38}$$

This relation fixes the mass of the W boson with respect to the mass of the Z boson, and has been confirmed experimentally to be true with high precision in collider experiments. The weak mixing angle θ_W also relates the coupling constants g and g' by

$$\sin\theta_W = \frac{g'}{\sqrt{g^2 + g'^2}} \approx 0.215 \tag{A.39}$$

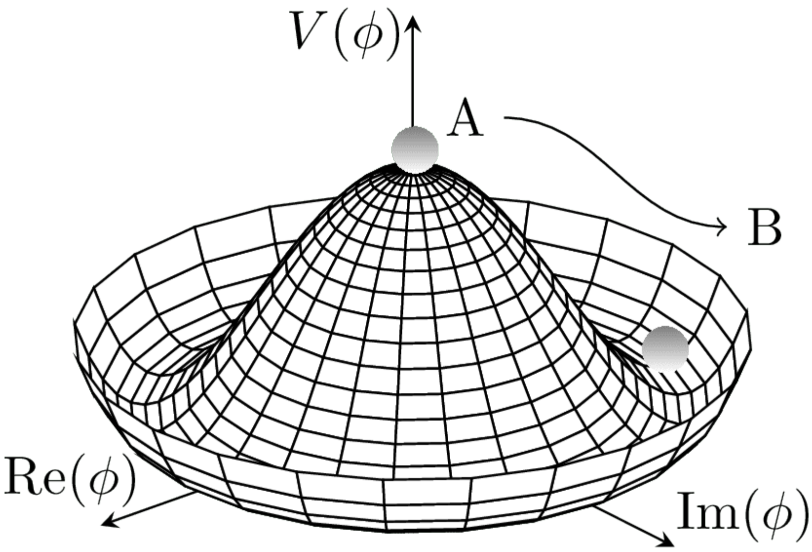

Figure A.1: Illustration of the shape of the Higgs potential in 3 dimensions, also known as the 'Mexican hat' potential due to its iconic shape. Point A represents the field configuration when $\phi = 0$, while B represents the breaking of the symmetry, leading to the field attaining a preferred minimum energy state.

The Higgs potential contains higher order terms, representing the self-interactions of the Higgs boson, which give rise to its mass, M_H. However, this quantity can only be measured experimentally as long as λ is a free parameter, because $M_H = \sqrt{2\lambda}\nu$. The Higgs boson was discovered in 2012 by the CMS and ATLAS experiments at the LHC after decades of searches, and its mass was measured to be $M_H = 125.2 \pm 0.11$ GeV.

Additionally, it is also possible to introduce a Yukawa interaction term, $\mathcal{L}_{\text{Yukawa}}$ between the new scalar field and the fermions

$$\mathcal{L}_{\text{Yukawa}} = Y_{ij}[\bar{L}_i \phi R_j + h.c.] \rightarrow^{SSB} -\sum_{i,j} Y_{ij} \left(\frac{v + H}{\sqrt{2}} \bar{f}_i f_i \right) \tag{A.40}$$

If this Yukawa matrix is chosen to be diagonal, the masses of the fermions can be obtained as a function of the Yukawa couplings y_i of the Higgs field with the fermions

$$m_i = \frac{y_i v}{\sqrt{2}} \tag{A.41}$$

where the index i runs over the leptons and quarks. The coupling strength is given by

$$c_i = \frac{y_i H}{2} = \frac{m_i H}{v} \tag{A.42}$$

This result reveals that the mass of a charged fermion is proportional to its coupling strength to the Higgs field: the stronger the coupling, the heavier the fermion. This is illustrated in Figure A.2.

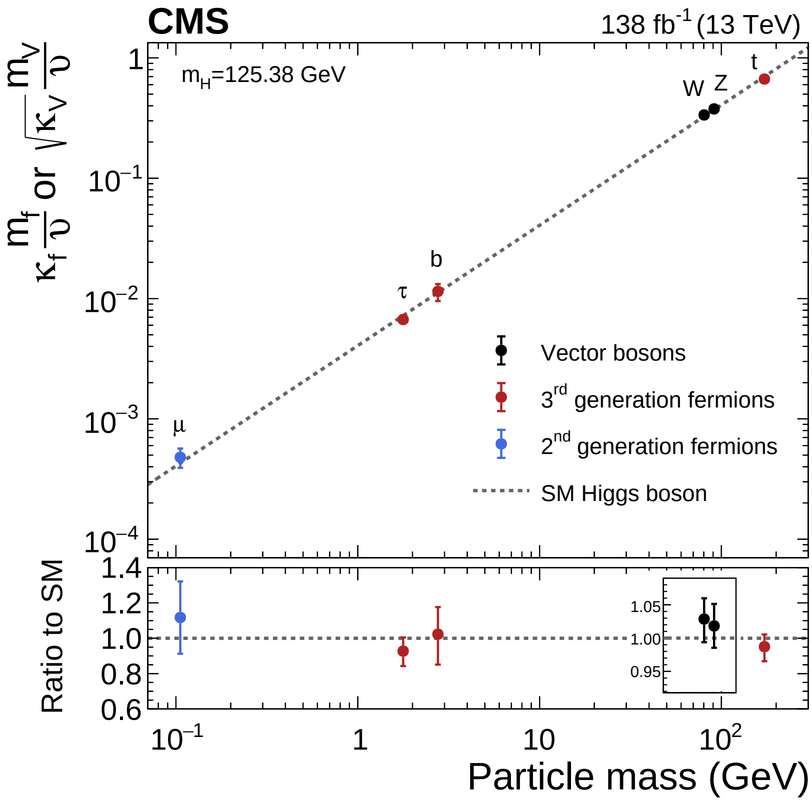

Figure A.2: CMS result [119] showing the measured coupling strengths of the SM particles as a function of their mass (red), compared with SM predictions. Sourced from [122].

A.6 Summary

The gauge principle allows us to formulate the theoretical framework of the different fundamental interactions (excluding gravity), namely the strong, electromagnetic, and weak interactions, using only symmetry arguments. However, this results in a theory with only interactions mediated by massless spin-1 fields, which contradicts experimental observations of the massive W and Z bosons, as well as the fermions.

We are able to introduce a scalar field to the theory that is able to spontaneously break the original symmetry and generate masses for the different particles. This scalar field predicts the existence of the Higgs boson, which was experimentally discovered in 2012, concluding the search for SM particles – as all of them have now been found.

To build a theory as complex and comprehensive as the SM, we require *only* 18 different free parameters whose values cannot be predicted by theory and need to be experimentally determined. These are three masses for the leptons (leaving the neutrinos massless in the SM), six masses for the quarks, three mixing angles and 1 CP violating phase in the CKM matrix, three gauge couplings, a vacuum expectation value for the Higgs field, and the mass of the Higgs boson itself.

Appendix B

TMS muon and dimuon selection

B.1 TMS muon identification

We begin with an initial list of TMS muon candidates consisting of all global and arbitrated tracker muons with $p_T > 10$ GeV that are matched to STA muons by the STA to TMS muon association procedure discussed in Section 5.5.2. The position resolution of TMS muons is significantly superior to that of STA muons due to the excellent resolution of the tracker, resulting in a much better L_{xy} resolution (as shown in Figure 5.20), as well as p_T and d_0 resolution. The p_T resolution of TMS muons is in the range of 1-3%, more than an order of magnitude better than that of STA muons, which is typically between 20 and 30%. This translates into a large improvement in the invariant mass of resolution of TMS-TMS dimuons.

The standard identification criteria for global and tracker muons, designed for prompt muons, result in poor efficiency for muons not emerging from the PV. Therefore, the following identification requirements, developed and optimized in the Run 2 analysis [22] using simulated signal and background samples, as well as samples of loosely isolated TMS muons forming fake displaced dimuons in signal-free regions, are applied to TMS muons in our analysis :

- relative p_T uncertainty $\sigma(p_T)/p_T < 1$, where $\sigma(p_T)$ is the internal uncertainty of the muon track fit;
- is a tracker muon;
- more than one matched CSC or DT segment.

For the 2022 analysis, the cut values remain the same.

B.2 TMS muon isolation

Both TMS muons are required to be isolated from other activity in the event, to suppress large backgrounds from dijet and multijet QCD events. The relative tracker isolation $I^{\mathrm{rel}}_{\mathrm{trk}}$

gives the best performance for the search. This algorithm calculates the scalar sum of the p_T of tracker tracks reconstructed in a cone of radius $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$ centered on the muon track direction. The tracker tracks in the cone must satisfy four requirements, optimized for this search, to be included in the sum.

- The distance along the beamline, $|\Delta z|$, between the reference point of the track and that of the muon defining the cone, must be smaller than 5 mm;
- The impact parameter d_0 of the track with respect to the beam spot must be smaller than 2 mm.
- ΔR between the track and muon defining the cone must be larger than 0.01 and smaller than 0.3.
- If the isolation cone around a given muon, μ_1, also contains the tracker track for the second muon forming the dimuon, μ_2, it is not included in the $\sum p_T$ for μ_1 (and vice versa). This requirement dramatically improves the efficiency for signal events with boosted LLPs yielding a collimated pair of muons, while having a negligible impact on background rejection.

The p_T of the muon track itself is not included in the sum. The muon is considered isolated if the ratio of the p_T sum to the muon track p_T is below 0.075. The track direction and p_T values are computed at the point of closest approach to the beam spot. Thus, optimal requirement on the cone size and isolation threshold

$$I^{\mathrm{rel}}_{\mathrm{trk}}(\Delta R < 0.3) < 0.075 \tag{B.1}$$

is applied to each of the two TMS muons in a dimuon

B.3 TMS-TMS event selection

- An important discriminating variable in the TMS-TMS category is the d_0 **significance**, $d_0/\sigma(d_0)$. This quantity is usually large for muons produced in the decays of long-lived particles and small for promptly-produced muons pointing towards the PV. The measurement of $\sigma(d_0)$ includes both track and PV uncertainties.

  Dimuons in the TMS-TMS category are required to be displaced with respect to the PV by requiring that both the L_{xy} significance, $L_{xy}/\sigma(L_{xy})$, as well as the d_0 significance, $d_0/\sigma(d_0)$ be larger than 6. The signal region is further subdivided into three $\min(d_0/\sigma(d_0))$ bins : [6, 10], [10, 20], and [20, inf).
- To suppress backgrounds consisting of dimuons formed from unrelated muons, a threshold is imposed on the χ^2 of the vertex fit, $\chi^2_{\mathrm{vertex}} > 10$.

- Backgrounds from unrelated or mismeasured TMS muons are further suppressed by requiring both muons forming the TMS-TMS dimuon to have a similar number of pixel hits. The difference between the number of pixel hits on the two TMS muons in a dimuon must be smaller than 3.

$$\Delta N(\text{pixel hits}) = \left| N(\text{pixel hits})_{\mu_1} - N(\text{pixel hits}_{\mu_2}) \right| < 3 \tag{B.2}$$

  For TMS-TMS dimuons arising from LLP decays outside the pixel detector, this quantity is 0 and is 100% efficient, while for other signals, it reduces the efficiency by only a few percent.

- To suppress cosmic ray muons, in addition to cuts on timing and direction, in the TMS-TMS category, back-to-back muons are suppressed by imposing a threshold the 3D opening angle, $\alpha < 3.0$ rad ($\cos\alpha > -0.99$). The cut value reflects the much better angular resolution of TMS muons.

- To suppress prompt Drell Yan type backgrounds in the TMS-TMS category, thresholds $|\Delta\Phi|_\text{C}$ are applied on the transverse collinearity angle $|\Delta\Phi|$ such that $|\Delta\Phi| < |\Delta\Phi|_\text{C}$. This threshold value is defined to be $|\Delta\Phi|_\text{C} = \pi/30$ for the HAHM model, reflecting the superior dimuon momentum and L_{xy} resolution in this category, and $|\Delta\Phi|_\text{C} = \pi/4$ for the RPV SUSY model. Figure B.1 shows the $|\Delta\Phi|$ distributions in both HAHM and RPV signal samples, as well as a prompt DY background enriched sample in data.

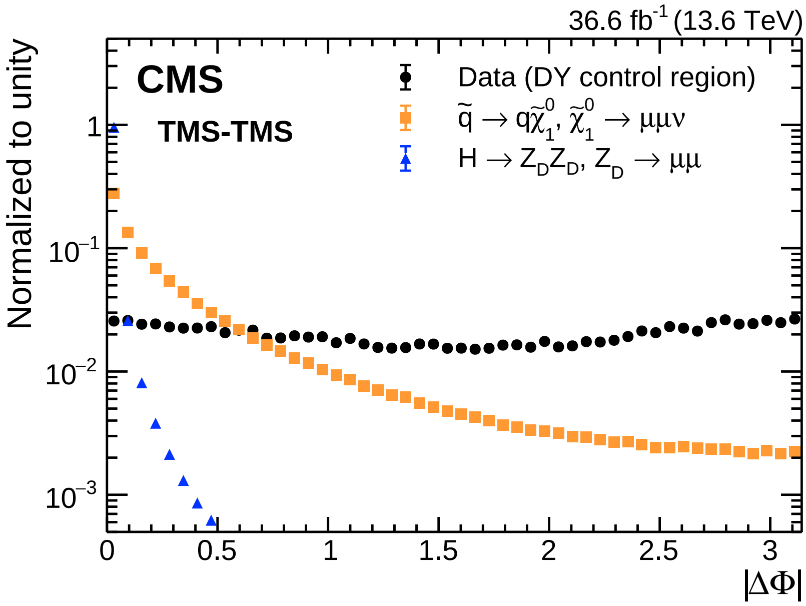

Figure B.1: Distribution of $|\Delta\Phi|$ for TMS-TMS dimuons in a data sample enriched in Drell-Yan events (black circles), and for events passing all selection criteria, except for the requirement on $|\Delta\Phi|$, in all HAHM (blue triangles) and RPV SUSY (orange squares) generated signal samples combined. All distributions are normalized to unity.

- To suppress QCD dijet and multijet events yielding background dimuons that are formed from genuine or fake muons belonging to different jets, we require p_T thresholds on the muons forming the dimuon in the TMS-TMS category. The leading muon is required to have a $p_T > 25$ GeV, while the subleading muon is required to have a $p_T > 10$ GeV. These thresholds were optimized in the Run 2 search and were found to be effective even in Run 3, even after adding lower p_T HLT paths, as shown in Figure B.2.

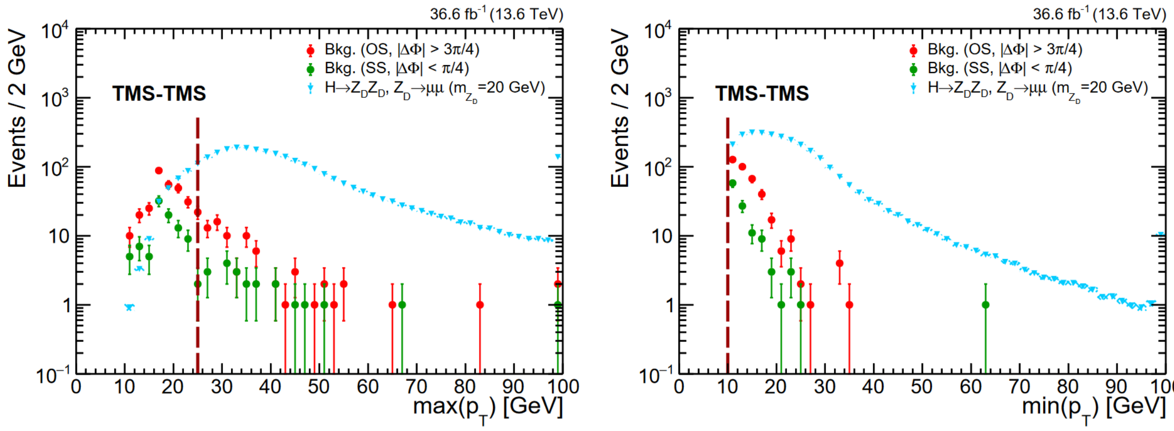

Figure B.2: Distribution of the p_T of the leading (left) and subleading (right) muon forming the TMS-TMS dimuons for events in the HAHM sample with $m(Z_D) = 20$ GeV (light blue), and in the OS $|\Delta\Phi| > 3\pi/4$ (red) and SS $|\Delta\Phi| < \pi/4$ (green) control regions in 2022 data. The full selection is applied except for the requirement on $\max(p_T)$. The last bin contains the overflow.

- To suppress low mass backgrounds such as those from QCD processes, we require the reconstructed invariant mass of the dimuon in the TMS-TMS category, $m_{\mu\mu} > 10$ GeV. Taking advantage of the superior invariant mass resolution in the TMS-TMS category, which is between 1-3%, we perform the search for LLPs in this category using narrow mass intervals to suppress remaining backgrounds, which are typically equal to $\pm 6 - 8$ times the mass resolution. These are listed in Table B.1.

A summary of the selection is provided in Table 5.3. The offline event selection efficiencies for TMS-TMS dimuons are shown in Figure B.3 as a function of their simulated L_{xy} and $|z|$ in the HAHM signal samples, with the approximate geometry of the CMS tracker modules overlaid on top, showing displaced vertices reconstructed up to 60 cm in L_{xy} and 150 cm in $|z|$.

B.4 TMS-TMS signal region

The background estimation method in the TMS-TMS dimuon category is performed in generally the same manner as in the STA-STA category and is described in [23], and

Mass interval [GeV]	$(m(H), m(X))$ [GeV]	$m(Z_D)$ [GeV]
$10 < m_{\mu\mu} < 12$	-	10
$18 < m_{\mu\mu} < 22$	(125, 20), (200, 20), (400, 20), (1000, 20)	20
$27 < m_{\mu\mu} < 33$	-	30
$37 < m_{\mu\mu} < 43$	-	40
$46 < m_{\mu\mu} < 54$	(125, 50), (200, 50), (400, 50), (1000, 50)	50
$55 < m_{\mu\mu} < 65$	-	60
$140 < m_{\mu\mu} < 160$	(400, 150), (1000, 150)	-
$326 < m_{\mu\mu} < 374$	(400, 150), (1000, 150)	-

Corrected mass interval [GeV]	$m(\tilde{q})$ [GeV]	$m(\tilde{\chi}_1^0)$ [GeV]
$30 < m_{\mu\mu}^{\text{corr}} < 55$	125, 200, 350, 700, 1150, 1600	50
$60 < m_{\mu\mu}^{\text{corr}} < 110$	125	100
$90 < m_{\mu\mu}^{\text{corr}} < 165$	350	150
$105 < m_{\mu\mu}^{\text{corr}} < 192$	200	175
$195 < m_{\mu\mu}^{\text{corr}} < 357$	350	325
$300 < m_{\mu\mu}^{\text{corr}} < 550$	700, 1150, 1600	500
$405 < m_{\mu\mu}^{\text{corr}} < 742$	700	675
$570 < m_{\mu\mu}^{\text{corr}} < 1045$	1150, 1600	950
$675 < m_{\mu\mu}^{\text{corr}} < 1237$	1150, 1600	1125
$840 < m_{\mu\mu}^{\text{corr}} < 1540$	1600	1400
$945 < m_{\mu\mu}^{\text{corr}} < 1732$	1600	1575

Table B.1: Correspondence between the invariant mass intervals in the TMS-TMS category and the parameters of the simulated signal samples.

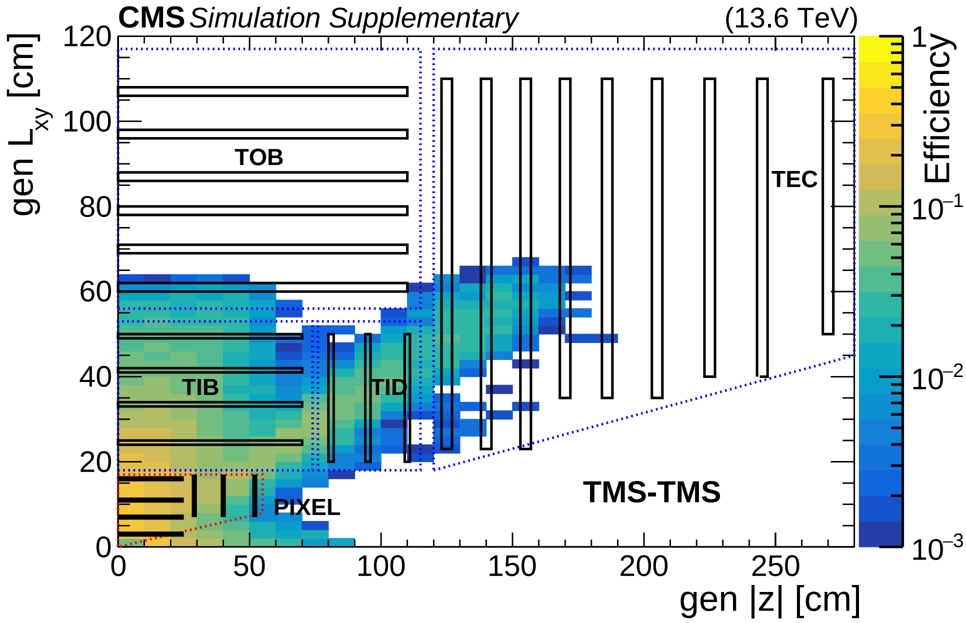

Figure B.3: Offline selection efficiency of TMS-TMS dimuons as a function of simulated L_{xy} and $|z|$ decay positions of the dark photon Z_D, in all generated HAHM signal samples combined. Approximate locations of the silicon pixel and strip tracker modules, namely the tracker inner barrel (TIB) and disks (TID), tracker outer barrel (TOB), and tracker endcap (TEC), are drawn superimposed as black boxes.

therefore omitted here in the interest of brevity. However, the final background predictions and yields in the TMS-TMS signal region are included here for posterity.

Figure B.4 shows the unblinded $\min(d_0/\sigma(d_0))$ distributions in the signal region, in the TMS-TMS category for the HAHM analysis with $|\Delta\Phi| < \pi/30$ (left) and the RPV SUSY model analysis with $|\Delta\Phi| < pi/4$ (right), showing the predicted background yields, as well as the observed number of events. Also shown are the expected signal yields, based on the median expected 95% CL exclusion limits obtained from the ensemble of both the TMS-TMS and STA-STA categories, scaled up to improve visibility.

Figures B.5 and B.6 similarly show the unblinded $m_{\mu\mu}$ and $m_{\mu\mu}^{\text{corr}}$ distributions, in each of the $\min(d_0/\sigma(d_0))$ bins, for the HAHM and RPV SUSY model analyses respectively. The predicted background yields, as well as the observed number of events, are shown, along with the expected signal yields, based on the median expected 95% CL exclusion limits obtained from the ensemble of both the TMS-TMS and STA-STA categories, scaled up to improve visibility.

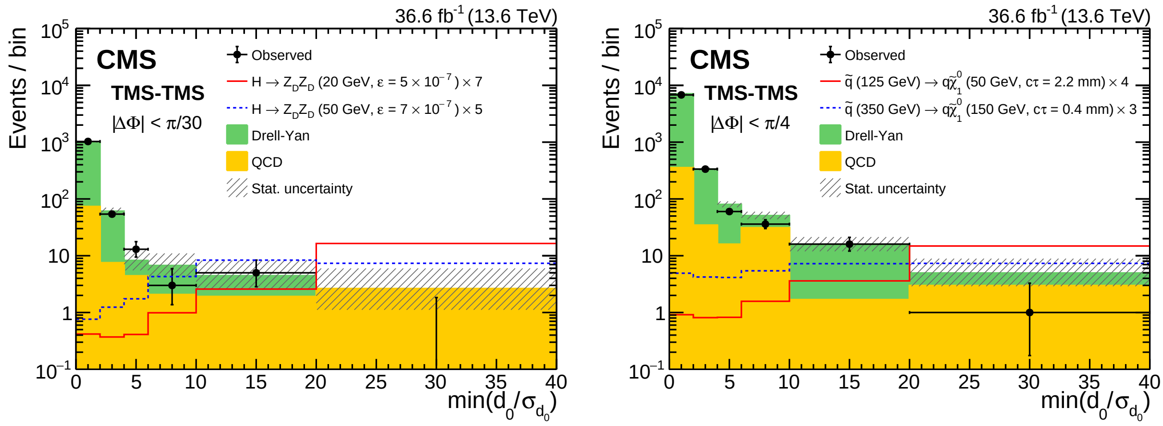

Figure B.4: Distributions of $\min(d_0/\sigma(d_0))$ for TMS-TMS dimuons with (left) $|\Delta\Phi| < \pi/30$ and (right) $|\Delta\Phi| < \pi/4$, for events in all mass intervals combined. The number of observed events (black circles) is overlaid with the stacked histograms showing the expected numbers of QCD (yellow) and DY (green) background events. Statistical uncertainties in the total expected background are shown as hatched histograms. Signal contributions expected from simulated signals indicated in the legends are shown in red and blue. Their yields are set to the corresponding median expected 95% CL exclusion limits obtained from the ensemble of both the TMS-TMS and STA-STA dimuon categories, scaled up as indicated in the legend to improve visibility. Events are required to satisfy all nominal selection criteria with the exception of the $d_0/\sigma(d_0)$ requirement. The last bin includes events in the histogram overflow.

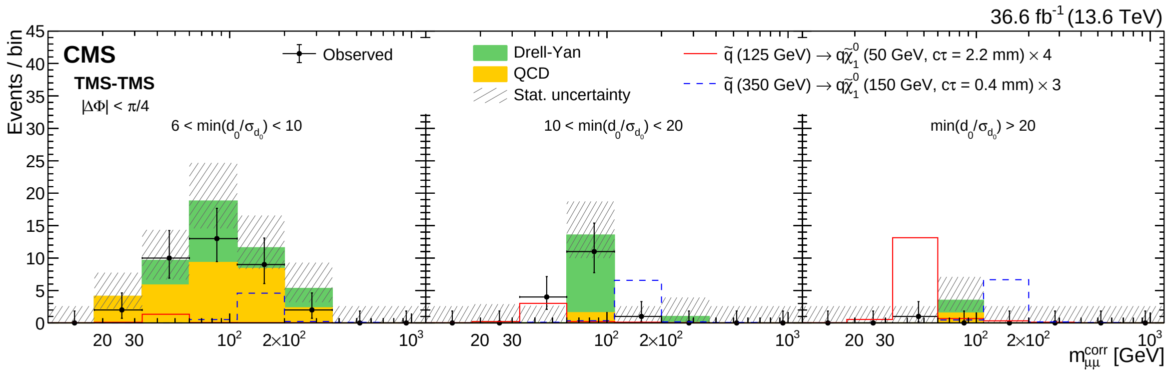

Figure B.5: Comparison of observed and expected numbers of events in the TMS-TMS dimuon category, for the HAHM analysis requiring $|\Delta\Phi| < \pi/30$, in bins of $m_{\mu\mu}$. The number of observed events (black circles) is overlaid with the stacked histograms showing the expected numbers of QCD (yellow) and DY (green) background events in bins of $m_{\mu\mu}$ in three $\min(d_0/\sigma(d_0))$ bins: (left) 6-10, (center) 10-20, and (right) > 20. Hatched histograms show statistical uncertainties in the total expected background. Signal contributions expected from simulated $h \to Z_D Z_D$ events (with the parameters indicated in the legends) are shown as red and blue histograms. Their yields are set to the corresponding median expected 95% CL exclusion limits obtained from the ensemble of both the TMS-TMS and STA-STA dimuon categories, scaled up as indicated in the legend to improve visibility. The last bin includes events in the histogram overflow.

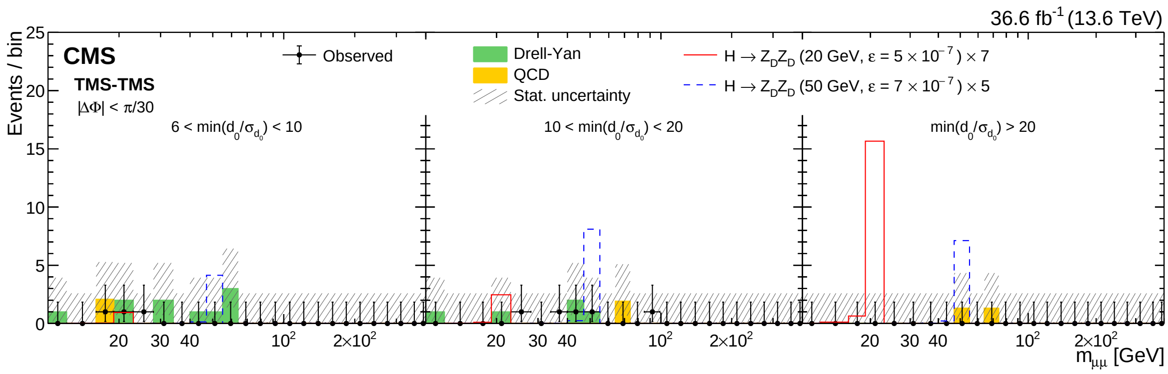

Figure B.6: Comparison of observed and expected numbers of events in the TMS-TMS dimuon category, for the RPV SUSY model analysis requiring $|\Delta\Phi| < \pi/4$, in bins of $m_{\mu\mu}^{\text{corr}}$. The number of observed events (black circles) is overlaid with the stacked histograms showing the expected numbers of QCD (yellow) and DY (green) background events in bins of $m_{\mu\mu}^{\text{corr}}$ in three $\min(d_0/\sigma(d_0))$ bins: (left) 6-10, (center) 10-20, and (right) > 20. Hatched histograms show statistical uncertainties in the total expected background. Signal contributions expected from simulated events in the RPV SUSY model (with the parameters indicated in the legends) are shown as red and blue histograms. Their yields are set to the corresponding median expected 95% CL exclusion limits obtained from the ensemble of both the TMS-TMS and STA-STA dimuon categories, scaled up as indicated in the legend to improve visibility. The last bin includes events in the histogram overflow.

Appendix C

Lifetime Reweighting

The signal samples used in this analysis were produced over a discrete lifetime grid. To obtain the signal efficiencies and corresponding limits at lifetime values in between the generated lifetime points, we use a lifetime reweighting procedure described in this appendix.

The decay time of the LLP X can be expressed in terms of reconstructed quantities

$$t = \frac{m_X \cdot L_{xy}}{p_T^X} \tag{C.1}$$

A particle with mean proper lifetime τ will have its decay time distributed according to an exponential probability distribution given by Equation 2.4.

The signal samples were produced for three or four simulated lifetime points for each combination of masses in the HAHM, RPV SUSY, and BSM heavy scalar models. To interpolate between these lifetime points, we can define a meta-sample probability distribution

$$\mathcal{M}(t) = \frac{\sum\limits_{i=1}^{n} \mathcal{P}(t|\tau_i) N_{\tau_i}}{\sum\limits_{i=1}^{n} N_{\tau_i}} \tag{C.2}$$

where $\mathcal{P}(t|\tau_i)$ and N_{τ_i} represent the decay probability distribution and number of generated dimuons for a generated sample i with mean proper lifetime τ_i respectively, among n total samples. The meta-sample can then be multiplied by a weight $w(t,\tau)$ such that it results in a probability distribution with an arbitrary lifetime $\mathcal{P}(t|\tau)$

$$w(t|\tau) = \frac{\mathcal{P}(t|\tau)}{\mathcal{M}(t)} \tag{C.3}$$

Our selection efficiency, $\tilde{\epsilon}(t)$ is then dependent only on the decay time t of the LLP, not the generated lifetimes of the samples, allowing us to write the (non-normalized) lifetime distribution of selected LLP $\mathcal{S}$ as

$$\mathcal{S}(t) = \tilde{\epsilon}(t)\mathcal{M}(t) \tag{C.4}$$

To obtain the selection efficiency $\epsilon(\tau)$ of the full meta-sample for an arbitrary lifetime τ, we can use

$$\int \tilde{\epsilon}(t)w(t,\tau)\mathcal{M}(t)dt = \int w(t,\tau)\mathcal{S}(t)dt = \epsilon(\tau) \tag{C.5}$$

This allows us to derive the selection efficiency for lifetimes other than those generated for the analysis.