DISPLACED VERTICES
A SEARCH FOR LONG-LIVED PARTICLES AT CMS

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Dan Quach

December 2020

DISPLACED VERTICES

A SEARCH FOR LONG-LIVED PARTICLES AT CMS

Dan Quach, Ph.D.

Cornell University 2020

We report the results of a search for long-lived particles produced in pairs in proton-proton collisions at the LHC operating at a center-of-mass energy of $\sqrt{s} = 13\,\text{TeV}$. The data were collected by the CMS detector during the full Run 2 data taking period from 2015 through 2018, corresponding to a total integrated luminosity of $140\,\text{fb}^{-1}$. This search targets pair-produced long-lived particles with a mean proper decay length between 0.1 and 100 mm that each decay into at least two quarks. The signature is a pair of displaced vertices each formed from many tracks. This search extends a previous CMS search using the 2015 and 2016 dataset, with improvements in background rejection, background estimation techniques, as well as uncertainty estimation. Results are compared with *R*-parity violating supersymmetry models that predict pair-produced long-lived particles, each decaying into multijet or dijet final states. No events are observed with two reconstructed high-track-multiplicity vertices. For models of long-lived pair-produced neutralinos, gluinos, and top squarks, pair-production cross sections larger than 0.08 fb at 95% confidence level are excluded for masses between 800 and 3000 GeV and mean proper decay lengths between 1 and 25 mm. In the gluino model, masses up to 2500 GeV are excluded for mean proper decay lengths between 0.6 and 90 mm. In the neutralino model, masses up to 1100 GeV are excluded for mean proper decay lengths between 0.6 and 70 mm. For the top squark model, masses up to 1600 GeV are excluded for mean proper decay lengths between 0.4 and 80 mm.

BIOGRAPHICAL SKETCH

Dan Quach was born into a Vietnamese immigrant family, the youngest of three, in the southwestern U.S. city of Phoenix, AZ. He lived in Phoenix for the first 23 years of his life, attending Arizona State University for his undergraduate studies. At ASU, he undecidedly bounced around different majors from computer systems engineering to aerospace engineering, and finally wound up graduating with a degree in physics and math at the end of 2012.

In 2013, he made his way to the opposite end of the continental U.S. to Ithaca, NY where he continued his physics education at Cornell University. Initially, he dabbled in astrophysics research, studying counter-rotating stellar accretion disks, before finally switching over to experimental high-energy physics to work on the CMS experiment under Professor Ritchie Patterson. During this time, while doing his particle physics research, he was also hopping around geographically. He spent a couple of years in France and Switzerland to work on the detector hardware at CERN. Following this, he spent a year in the Chicagoland area working on his analysis out of Fermilab before finally returning to Ithaca to finish up his studies. When not focusing on his analysis, Dan can sometimes be found trying to get swole in the weight room or getting choked out on the mats in Brazilian jiu-jitsu, barring any injuries he may be facing that prevent him from doing either. In addition, he has acquired random hobbies over this time, learning how to ski in the Alps, 3-ball juggling, and solving various twisty cube puzzles (not all simultaneously). He has tried to add unicycling to his repertoire, but so far this skill has eluded his best attempts. He will be defending this document in August 2020, in the middle of a global pandemic. What lies ahead in the next chapter of Dan's life? ¯_(ツ)_/¯

To the beans.

ACKNOWLEDGEMENTS

At the beginning of 2015, I asked my soon-to-be-advisor, Ritchie Patterson, if I could join her research group—a request she graciously (perhaps foolishly 😜) accepted. By this point, I was a lost and confused graduate student in the middle of my second year of study. Now, as I write this passage, I am nearing the conclusion of this now seven-year long journey, emerging on the other side slightly less lost and confused but a far better scientist due in large part to her guidance and support. I want to thank Ritchie for her patience in our countless discussions over the years. She brings a high level of clarity in communicating her thoughts and ideas and encourages her students to do the same; she was the quantitative foil to my qualitative conjecturing. Furthermore, despite her incredibly packed schedule, Ritchie would still always find time for her students and was completely engaged in the matter at hand. I want to thank her for the independence and flexibility she extended to me while I was stationed at CERN to pursue the things I felt were important while also gently reining me back in when I briefly lost sight of what I needed to accomplish. I am extremely fortunate to have had her as my advisor.

I want to thank my post-doc, Jordan Tucker, a second advisor, a mentor, and friend. I deeply appreciate his willingness to endure all of my questions as a novice experimental particle physicist from the beginning; helping me navigate through the labyrinthine code base authored by hundreds of other physicists; and sharing his broad and deep wealth of knowledge and technical expertise, which has helped me become far more technically competent and proficient.

I want to thank all of the other people involved in my analysis over the years: Jennifer Chu, Joey Reichert, Shaun Hogan, and Werner Sun. Each one, along with Ritchie and Jordan, has dedicated enormous amounts of their time and effort into helping craft this analysis. They have helped made it, in my *completely unbiased* opinion, one of the coolest analyses coming out of CMS.

I'd like to also thank my other committee members, Julia Thom-Levy and Maxim Perelstein for reading and providing their comments on the various manuscripts I've written for their consumption (and my education) over the years, including this one. I am additionally grateful to Julia for introducing me to the world of pixels. I may have over-enthusiastically jumped into it in lieu of other things, but it has been an extremely meaningful experience during my graduate career, and I am grateful for having had the opportunity.

On the topic of pixels, there have been a large number of individuals involved in the phase-1 pixel detector upgrade project during the time that I was deeply embroiled in its activities at CERN, including but not limited to: Will Johns, Satoshi Hasegawa, Weinan Si, Karl Ecklund, Danek Kotlinski, Bora Akgün, Stefanos Leontsinis, Benedikt Vormwald, Jory Sonneveld, Atanu Modak, Klaas Padeken, Hannsjörg Weber, and Matt Kilpatrick. It was a privilege to work with all of these people on developing, commissioning, operating, debugging, fixing, and maintaining our occasionally dysfunctional detector. Despite all its warts, it is still unambiguously my favorite subdetector of CMS.

I'd like to thank the Cornell CMS community as a whole who have made the last several years truly fulfilling, both at CERN and Cornell: Ritchie Patterson, Julia Thom-Levy, Peter Wittich, Jim Alexander, Anders Ryd, Jordan Tucker, Joey Reichert, Louise Skinnari, Aurelijus Rinkevicius, Livia Soffi, Yangyang Cheng, Rui Zou, Jose Monroy, Nathan Mirman, Jorge Chaves, Susan Dittmer, Shao Min Tan, Jennifer Chu, Kevin McDermott, Margaret Zientek, Zhengcheng Tao, Andre Frankenthal, Abhisek Datta, Derek Cranshaw, Shaun Hogan, Divya Gadkari, Tres Reid, and Sam Bright-Thonney. It's been an absolute pleasure working with everybody, attending boring meetings together, having lunch, having dinner, cruising around in the Cornell car at CERN, skiing at Chamonix, discussing dumb stuff, discussing less dumb stuff, eating cheese, putting on juggling competitions in the office, and the list goes on.

Thank you to my best friends from Phoenix, who are now scattered around the world: Ishan Joshipura, Alexander Tiplea, and Lance McDiffett. You mean a lot more to me than you probably know.

Finally, I would like to thank my family, Mom, Dad, Cheese, Pom, and all the beans. Thank you for all your love and support over these seven years. And thank you to Julia Win for her continued encouragement and patience with me. I love you all.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

The study of high energy particle physics is a story of the human endeavor spanning millennia to understand the nature of the universe at ever shrinking length scales. It is the embodiment of reductionism, the philosophical idea that the behavior of a complex system may be fundamentally understood with a deep understanding of its constituent parts. Over the past century, rapid developments in the field propelled by advancements in technology and fueled by the consistent interplay of theoretical modeling and experimental discovery, have culminated in what is now known as the standard model (SM) of particle physics. At present, the SM encapsulates the best knowledge and understanding of the elementary building blocks of matter and the interactions between them. This model has been tested and cross-validated on both theoretical and experimental fronts to staggering levels of precision in many different contexts [26, 52] and has led to predictions of several particles that were later discovered experimentally, the most recent of which was the 2012 discovery of the Higgs boson at the Large Hadron Collider (LHC).

But despite its numerous successes, there are still missing pieces to the puzzle, apparent gaps in the explanatory power of the model. These range from the lack of treatment of gravitational interactions between the particles, no explanation for the observed matter-antimatter asymmetry, no potential candidates for dark matter and energy, and no explanation of neutrino masses. In addition to these, the model requires many ad hoc parameters to be specified, the values for which must be finely tuned to match experimental observations. For these reasons among several others, the SM is considered incomplete and has stimulated many efforts to find a more descriptive and foundational theory of nature that addresses these problems. These are referred to as searches for physics beyond the standard model (BSM).

Models of BSM physics can predict the existence of particles exhibiting a wide gamut of phenomenological features. These features lead to a number of distinct signatures when measured in the detector that analyses can exploit by developing techniques to explicitly target the handful of events with such features among tens of billions in data. A signature of particular interest to this thesis is one in which a new proposed particle has a long lifetime during which it may travel a macroscopically observable distance before decaying. Many earlier searches for BSM physics have focused on either the production of new states that decay promptly into visible final states or the production of new states that decay into final states with unseen particles that manifest as missing energy in the detector. More recently, increased interest has been devoted toward exploring the more generic and intermediate possibility of long-lived particles that may decay within the detector but away from the prompt interaction region. One needs to look no further than the SM to find evidence of such behavior to be well-motivated as a characteristic, as can be seen in Fig. 1.1, which plots the mass and measured lifetime of various particles that exist in the SM.

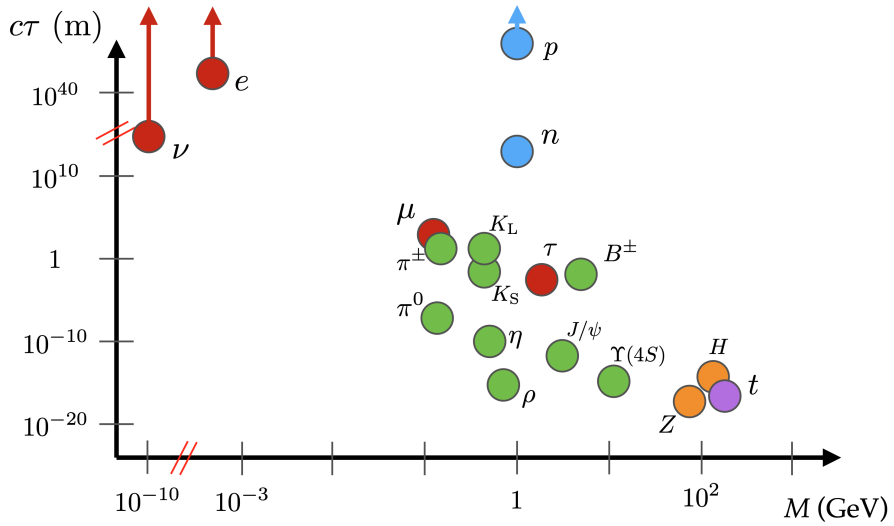

Figure 1.1: A two-dimensional scatter plot of particle lifetime vs. particle mass for various particles in the standard model [11].

A diverse array of BSM physics models predict the production of long-lived particles

at the CERN LHC including, but not limited to: R-parity violating (RPV) supersymmetry (SUSY) [21, 37, 38, 50], split SUSY [15, 14, 16, 17, 47, 48, 53], hidden valley models [51, 68, 69], stealth SUSY [43, 44], as well as models giving rise to dark matter candidates [29, 36, 39, 41, 40, 49, 55, 56]. Searches for long-lived particles therefore probe a broad BSM parameter space for new physics as there are many different mechanisms that may suppress the decay and lead to increased lifetimes. Among these are small couplings, limited phase space in the final state, massive mediators, or approximate symmetries.

The expansive parameter space calls for an inclusive and model-independent search. Accordingly, this analysis targets long-lived particles that are produced in pairs and decay into final states with multiple charged particles, primarily in the form of jets. More concretely, this analysis parses through the the billions of events to find those that contain the unique, but generic, experimental signature of two vertices, formed from the intersection of multiple charged-particle trajectories, displaced from the interaction region, but within the radius of the LHC beam pipe.

As benchmarks, this analysis explores two SUSY signal models with distinct final states. Modern colliders such as the LHC have finally reached the energy capabilities to test weak-scale SUSY theories experimentally from their phenomenological predictions. Unfortunately, no conclusive evidence of SUSY in the simplest supersymmetric extension of the SM have been observed, with increasingly stringent lower bounds being placed on the production cross section for larger regions of parameter space. This has led to re-evaluations of various assumptions of SUSY, an example of which are the conservation of R-parity, and its necessity within the theory. Such directions have led to alternative theoretical formulations that relax these assumptions, which therefore sidestep the experimental constraints while also proposing new phenomenological consequences

to be explored. The first benchmark model for this analysis represents a minimal flavor violating model of RPV SUSY [37] in which the lightest SUSY particle (LSP) is a long-lived neutralino or gluino that is produced in pairs. This long-lived particle decays into three or more jets, resulting in a "multijet" final state signal topology. The second benchmark model is another RPV model in which the pair-produced top squark is the long-lived LSP [38]. Each squark decays into a pair of down antiquarks, resulting in a "dijet" final state signature.

The displaced vertices are reconstructed from charged particle tracks using a custom vertex reconstruction algorithm. To discriminate the signal from the SM background, we use the separation between the vertex pairs in the plane transverse to the beam direction. This variable distinguishes the signal events, which tend to have well-separated vertex pairs, from background events whose vertices originate from track misreconstruction and have marginal displacements from the beam axis.

We target signals with lifetimes corresponding to a mean proper decay length $c\tau$ range of 0.1 to 100 mm. Longer lifetimes are generally less accessible since we require the vertices to be reconstructed within the beam pipe radius to prevent background vertices from material interactions. In addition, this search is primarily sensitive to models in which the mass of each pair-produced particle exceeds 600 GeV because of a trigger requiring large total energy projected into the plane perpendicular to the beam direction.

This analysis is a continuation of the previous CMS displaced vertex search based on data collected during 2015 and 2016 [61]. The data used for this analysis were collected with the CMS detector in 2017 and 2018, corresponding to an integrated luminosity of 101 fb^{-1}. The CMS collaboration upgraded its inner pixel tracking detector during the winter technical stop between the 2016 and 2017 run periods, providing tracking improvements that benefit this analysis since the vertex reconstruction relies heavily

on track quality. While the overall analysis strategy remains largely the same as the previous analysis, improved techniques between these two iterations have further reduced background and also reduced systematic uncertainties. The results from the 2017 and 2018 data are combined with those of the previous 2015 and 2016 CMS displaced vertex search for a total integrated luminosity of $140\,\text{fb}^{-1}$. Similar searches have been performed by the ATLAS and CMS experiment [8, 64] that probe longer lifetimes exceeding several mm. This analysis complements these searches with greater sensitivity to shorter lifetimes.

This dissertation is organized by first providing theoretical context of the search in Chapter 2, after which a description of the machinery of the experiment used will be presented in Chapter 3, and the event reconstruction from the data collected from the detectors is described in Chapter 4. The detailed description of the search for displaced vertices is found in Chapter 5. A brief interlude comparing the analysis with other similar analyses is provided in Chapter 6 with a short discussion of future directions for the field. The summary of results and conclusions are provided in Chapter 7.

CHAPTER 2

THEORETICAL BACKGROUND

While the phenomenological signature of a long-lived particle decaying into jets may arise in multiple theories, the search described in this dissertation applies a pair of SUSY models as benchmarks to evaluate the search strategy employed. This chapter provides the theoretical motivation for studying SUSY with a brief historical overview along with details of early simplified models and their subsequent shortcomings that led to alternative formulations to address the experimental discrepancies. This culminates in a discussion of SUSY models with R-parity violation, which provides one mechanism for displaced signatures in experiments that this analysis is investigating.

2.1 Why SUSY?

SUSY is one of the most fully developed theories of BSM physics. The central feature of these theories is the fundamental relation postulated between the two basic classes of elementary particles, the bosons and fermions, theorizing the existence of a superpartner particle, or sparticle, of opposite spin statistics for each known particle in the SM. In a perfectly supersymmetric universe, a particle and its sparticle would have identical properties with the exception of its spin. Such a scenario was contrary to contemporary experimental data where decades of research had shown no evidence for the existence of any sparticle of equivalent mass to its SUSY partner. This suggested that SUSY, if it were a theory consistent with nature and reality, would have to be a spontaneously broken symmetry to produce the mass asymmetry where supersymmetric particles have larger masses than their respective SM counterparts. Superficially, it would seem the formulation of such theories introduces unnecessary complications including a dramatic increase in the number of tunable free parameters of the theory, from 26 in the SM

to over a hundred in the simplest extension. However, many remarkable features of SUSY provide satisfying resolutions to some of the fundamental problems of the SM, which kept SUSY in position as one of the leading candidates for BSM physics in recent history.

Following a period in the mid-20th century, during which a slew of new particles and resonances were appearing left and right at accelerator experiments, the physicists of the 1960's experienced a revolution in their attempts to wrangle this new metaphorical zoo of particles into a coherent framework. They had found tremendous success in applying symmetry principles for the underlying structure of the theory when organizing all the new particle data. In particular, application of the special unitary group, SU(3) in what is now referred to as SU(3) "flavor" symmetry, allowed classes of similar mass particles to be characterized by multiplets based on their properties such as isospin and "strangeness". However, subsequent attempts to further generalize the structure by incorporating SU(2) spin with this new SU(3) flavor symmetry into larger multiplets of a single hybrid group were largely unsuccessful. This ultimately culminated in a 'no-go' theorem published by Sidney Coleman and Jeffrey Mandula, proving the impossibility of combining space-time and internal symmetries except in trivial cases.

Other physicists around this time, however, were exploring radically new symmetries in the context of string theory searching for a fundamental relation between the two spin statistics of particles: bosons and fermions. This symmetry, now called a supersymmetry, was later shown by Haag, Lopuszanski, and Sohnius to circumvent the Coleman-Mandula theorem. The following years saw a burst of interest within the physics community of supersymmetric theories when many attractive properties of this symmetry began to emerge, providing further motivation for continued investigation [18]:

1. Since supersymmetry implies the existence of a superpartner of opposite spin

statistics for each SM particle, it was shown to provide a candidate solution to the hierarchy problem as a result of the relative minus sign between the fermionic and bosonic loop contributions, which would provide mutual cancellation to the Higgs mass. Furthermore, to avoid fine-tuning issues, the scale of the new particles are expected to be of $O(\text{ TeV})$ [57].

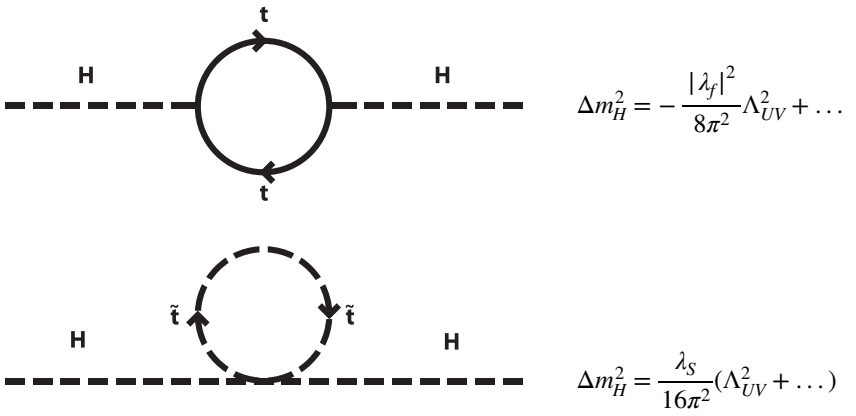

Figure 2.1: Feynman diagrams of fermionic (top) and scalar (bottom) loop corrections to the Higgs mass. Note the relative minus sign of the quadratic Λ_{UV} term in the Higgs correction.

2. The inclusion of the superpartner masses also modifies the evolution of the running gauge couplings of the Standard Model, resulting in a surprising unification of the gauge couplings when extrapolated to higher energies as shown in Figure 2.2, providing a viable grand unified theory model [57].

3. Baryon and lepton number violating processes have yet to be observed, prompting the imposition of an additional symmetry called R-parity to conform to existing observational bounds on the proton decay lifetime. This symmetry, if conserved, would imply the stability of a lightest supersymmetric particle that cannot decay into SM particles. This new stable particle, if it exists, is a potential candidate for cold dark matter in specific models of SUSY breaking [18].

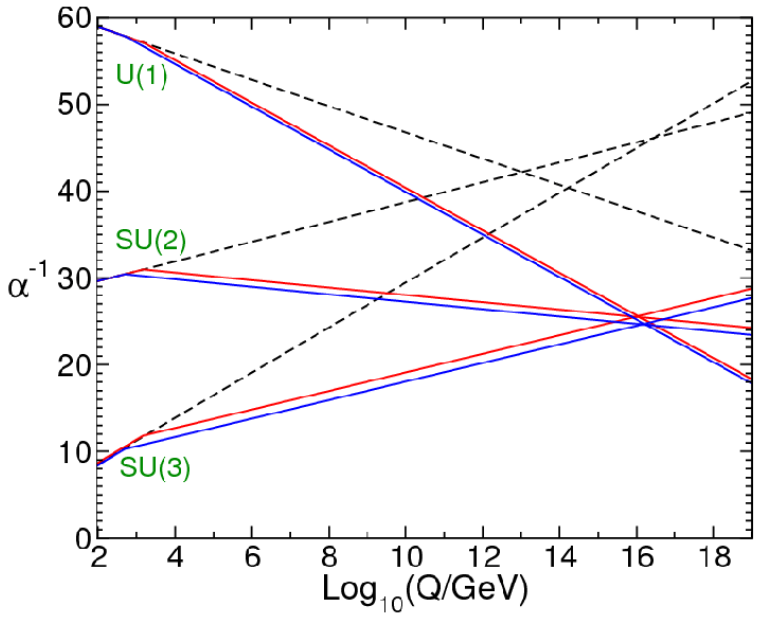

Figure 2.2: Evolution of gauge couplings for the Standard Model (dashed lines) and the MSSM (solid lines). The colors of the solid lines correspond to variations of the sparticle masses

Individually, any one of these features would justify further study and understanding, but all three packaged within a mathematically elegant structure that also illuminates a deep relation between fermions and bosons provided enormously compelling reasons to probe deeper into the theory.

2.2 Minimal Supersymmetric Standard Model

The simplest supersymmetric extension of the SM is known as the $\mathcal{N} = 1$ Minimal Supersymmetric Standard Model (MSSM), where $\mathcal{N}$ refers to the number of supersymmetry generators of the theory, Q and $Q^\dagger$. The theory is minimal in its addition of new particles and interactions to make the model phenomenologically viable.

2.2.1 Field Content

Constructing the MSSM begins with choosing the gauge symmetry group, i.e. the Standard Model gauge group $SU(3)_C \times SU(2)_L \times U(1)_Y$ and promoting the gauge fields to gauge superfields:

Name	Spin 1	Spin 1/2	$SU(3)_C \times SU(2)_L \times U(1)_Y$
gluon/gluino	g	$\tilde{g}$	$(8, 1, 0)$
W, Wino	$W^{\pm}, W^0$	$\tilde{W}^{\pm}, \tilde{W}^0$	$(1, 3, 0)$
B, Bino	B^0	$\tilde{B}^0$	$(1, 1, 0)$

Table 2.1: Table of vector supermultiplets of the MSSM.

The particle content of the gauge supermultiplets as tabulated in Table 2.1 consists of the SM gauge bosons along with their fermionic superpartner, a Weyl fermion called a gaugino. Specifying the matter content is the next step in constructing the MSSM, with the fermion fields of the SM being promoted to chiral scalar superfields:

Name	Symbol	Spin 1/2	Spin 0	$SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$
quarks/squarks	Q	$(u_L \;\; d_L)$	$(\tilde{u}_L \;\; \tilde{d}_L)$	$(3, 2, \frac{1}{6})$
	$\overline{u}$	$u_R^{\dagger}$	$\tilde{u}_R^*$	$(\overline{3}, 1, -\frac{2}{3})$
	$\overline{d}$	$d_R^{\dagger}$	$\tilde{d}_R^*$	$(\overline{3}, 1, \frac{1}{3})$
leptons/sleptons	L	$(\nu \;\; e_L)$	$(\tilde{n} \;\; \tilde{e}_L)$	$(1, 2, -\frac{1}{2})$
	$\overline{e}$	$e_R^{\dagger}$	$\tilde{e}_R^*$	$(1, 1, 1)$
Higgsinos/Higgs	H_u	$(\tilde{H}_u^+ \;\; \tilde{H}_u^0)$	$(H_u^+ \;\; H_u^0)$	$(1, 2, \frac{1}{2})$
	H_d	$(\tilde{H}_d^0 \;\; \tilde{H}_d^-)$	$(H_d^0 \;\; H_d^-)$	$(1, 2, -\frac{1}{2})$

Table 2.2: Table of chiral supermultiplets in the MSSM.

The bosonic superpartners of the fermions must be represented by scalar spin-0 particles, as opposed to 1 or higher, since the left and right handed components of the SM

fermions transform differently, which would not be allowed with the larger spins. Thus, the chiral superfields are comprised of the Weyl fermion and the complex scalar bosons known as sfermions.

The Higgs also resides in the chiral supermultiplet, where the scalar Higgs doublet is promoted to a chiral supermultiplet. However, as it turns out, a single Higgs doublet is insufficient for the theory as it introduces a gauge anomaly from the introduction of new fermions. This issue is resolved by adding a second doublet with the opposite $U(1)$ hypercharge, which restores the cancellation. Coincidentally, the inclusion of a second doublet resolves another problem where although the SM Higgs bestowed mass upon both the up and down type quarks, the $Y = +1/2$ supermultiplet can give masses only to the up-type quarks in the MSSM. The second doublet with $Y = -1/2$ can then grant mass to the down type quarks.

2.2.2 Interactions

With the field content established, the interactions of the MSSM can be specified via the superpotential W:

$$W = (\overline{u} y_u Q H_u - \overline{d} y_d Q H_d - \overline{e} y_e L H_d + \mu H_u H_d) + W_{\not{R}\not{L}} \tag{2.1}$$

The first three terms in parentheses of W correspond to the chiral superfields specified in Table 2.2 with the dimensionless Yukawa couplings y_e, y_u, y_d, which are 3×3 matrices corresponding to the 3 generations. These terms provide the SM fermions their masses. The μ term provides the supersymmetric Higgs mass for the two Higgs fields.

$W_{\not{R}\not{L}}$ consists of other valid gauge invariant and holomorphic terms, but violate either

baryon (B) or lepton (L) number:

$$W_{\not{R}\not{L}} = \frac{1}{2}\lambda^{ijk} L_i L_j \overline{e}_k + \lambda'^{ijk} L_i Q_j \overline{d}_k + \mu'^{i} L_i H_u + \frac{1}{2}\lambda''^{ijk} \overline{u}_i \overline{d}_j \overline{d}_k \tag{2.2}$$

These terms are troublesome since the phenomenon of B or L violation is tightly constrained by experimental observation. Searches for proton decay have placed limits longer than $\sim 10^{34}$ years on the lifetime of the proton, resulting in very small coefficients for B-violation terms in the superpotential. Thus, in constructing a SUSY model, these terms must be somehow removed. Manually arranging the couplings to vanish would not provide a satisfying resolution without an underlying reason for their smallness. Moreover, directly postulating B and L conservation is also misguided as $U(1)_B$ and $U(1)_L$ global symmetries are anomalous in the SM and are known to be broken by non-perturbative electroweak effects. Instead, a discrete symmetry called R-parity is imposed that more naturally prevents the undesirable B and L violating terms. R-parity can be encoded into a quantum number in terms of B, L, and the spin s of the particle as:

$$R \equiv (-1)^{3(B-L)+2s} \tag{2.3}$$

An immediate observation of this quantity is that it assigns a value $R = +1$ for ordinary SM fields on one hand, and $R = -1$ for the superpartner fields. Thereby implicitly conserving B and L with an explicit requirement of conserving R.

2.2.3 Soft supersymmetry breaking

Having specified the matter content and interactions of the theory, the full supersymmetric Lagrangian of the theory can now be defined. If supersymmetry were exact,

mass degenerate boson-fermion pairs would populate the universe, contrary to observations. This suggests an analogy to the spontaneous electroweak symmetry breaking mechanism within the SM. There are many viable mechanisms of SUSY breaking, none of which are clearly preferable to any other. This has led to the use of an effective Lagrangian augmented with explicit supersymmetry breaking terms as an unbiased approach that can, in principle, cover the full spectrum of possibilities. One must proceed with caution as these breaking terms should not reintroduce quadratic divergences. Referring to the Higgs correction terms as given with Figure 2.1, broken SUSY terms would yield the following correction to the Higgs mass:

$$\Delta m_H^2 = \frac{1}{8\pi^2}(\lambda_S - |\lambda_f^2|)\Lambda_{UV}^2 \tag{2.4}$$

It is clear that these SUSY breaking terms must be "soft", i.e. have a mass dimension of 2 or 3 in order to prevent large modifications that produce additional divergent contributions, which would impair the theory's ability to reconcile the hierarchy problem of the SM that provided the original motivation to develop the theory. The soft breaking terms of the MSSM take the form of:

$$\begin{aligned}
\mathcal{L}_{\text{soft}} = &-\frac{1}{2}(M_1\tilde{B}\tilde{B} + M_2\tilde{W}\tilde{W} + M_3\tilde{g}\tilde{g} + \text{c.c.}) \\
&- (\tilde{\bar{u}}\mathbf{a}_u\tilde{Q}H_u - \tilde{\bar{d}}\mathbf{a}_d\tilde{Q}H_d - \tilde{\bar{e}}\mathbf{a}_e\tilde{L}H_d) \\
&- \tilde{Q}^\dagger\mathbf{m}_Q^2\tilde{Q} - \tilde{L}^\dagger\mathbf{m}_L^2\tilde{L} - \tilde{\bar{u}}\mathbf{m}_{\bar{u}}^2\tilde{\bar{u}}^\dagger - \tilde{\bar{d}}\mathbf{m}_{\bar{d}}^2\tilde{\bar{d}}^\dagger - \tilde{\bar{e}}\mathbf{m}_{\bar{e}}^2\tilde{\bar{e}}^\dagger \\
&- m_{H_u}^2H_u^*H_u - m_{H_d}^2H_d^*H_d - (bH_uH_d + \text{c.c.})
\end{aligned} \tag{2.5}$$

The equation is divided into four distinct types of terms that can break SUSY softly:

1. Gaugino masses, as shown in the first line of this equation

2. Cubic scalar interactions in the second line
3. Scalar mass terms of the squarks and sleptons in the third line
4. Contributions from Higgs potential in the final line

These terms introduce undesirable behaviors into the theory including flavor mixing or CP-violating processes that remain unobserved in experiments. The assumption of flavor universality in which the couplings are independent of flavor has been used to restrict such processes and comply with experimental data.

2.3 R-parity

2.3.1 R-parity conserving models

Many of the earliest models of SUSY explicitly required R-parity conservation from which several important physical consequences arise:

- The lightest $R = -1$ particle would necessarily be stable since it would be unable to decay into any $R = +1$ SM particle. Typically, this has been named the "lightest supersymmetric particle" or LSP and, in many models, the LSP is electrically neutral and colorless, inspiring the name "neutralino," and has many excellent properties corresponding to a potential WIMP dark matter candidate.
- Sparticles can only be produced in SM collisions in even numbers via pair production, and will subsequently decay into the LSP

Since it is assumed the LSP is stable and very weakly interacting, its production in either case, whether directly or at the end of a sparticle decay chain, would escape the

detector manifesting as an imbalance in the transverse momentum. The characteristic missing transverse momenta or sometimes "transverse energy", E_T^{miss}, provides a distinct signature for SUSY that motivated searches at different colliders including LEP, the Tevatron, and the LHC. These searches have, in short, turned up empty-handed, setting lower bounds exceeding 1 TeV on some sparticle masses. This led to increased interest in R-parity violating (RPV) SUSY models that attempt to reconcile the same phenomenological observations without the explicit requirement of conserved R-parity.

2.3.2 R-parity violating models

One alternative approach toward building a viable SUSY model begins with the same field content and superpotential of the MSSM as given in Section 2.2, but does not immediately assume R-parity conservation nor flavor universality. Instead, the model imposes the minimal flavor violation (MFV) hypothesis, which neutralizes flavor mixing and CP-violating processes automatically, and also exhibits an adherence to the experimental bounds placed on B and L violating processes, thus imposing an accidental R-parity symmetry in the low energy regime [37]. At the same time, the model permits R-parity violation allowing for the decay of the LSP, offering a natural mechanism for it to elude the missing transverse energy signal searches.

The gauge interactions of the Standard Model do not depend on flavor, which implies that if the Yukawa coupling matrices were zero, the SM would be invariant under the following global symmetry group [42, 58]:

$$G_F \equiv \mathrm{SU}(3)_Q \times \mathrm{SU}(3)_u \times \mathrm{SU}(3)_d \times \mathrm{SU}(3)_L \times \mathrm{SU}(3)_e \times \mathrm{U}(1)^5 \tag{2.6}$$

In the SM, the quark and lepton masses and CKM mixing all break this symmetry. Furthermore, extending the SM into the MSSM introduces many additional sources that also

break this symmetry from terms in the superpotential and soft breaking terms. Thus, the model of MFV SUSY begins with the same chiral fields given in Table 2.2 and the superpotential defined in Equation 2.1, but the assumptions of R-parity conservation and flavor universality are replaced by the single assumption of minimal flavor violation. This 'minimal' assumption restricts the number of Yukawa couplings allowed to break the $SU(3)^5$ flavor symmetry, where $SU(3)^5$ corresponds to $SU(3)_Q\times$ $SU(3)_u\times$ $SU(3)_d\times$ $SU(3)_L\times$ $SU(3)_e$. In particular, this implies all flavor violating terms are those proportional to the coupling matrices y_u, y_d, and y_e in Equation 2.1. As a result of this structure, the largest couplings, and therefore the RPV effects, are strongest for the third generation (s)quarks. These three Yukawa couplings are promoted to holomorphic spurion fields used to parametrize the MSSM flavor-breaking couplings [37].

In the case of massless neutrinos, a symmetry $\mathbb{Z}_3^L \in SU(3)_L\times SU(3)_e$ arises of the form:

$$L \to \omega L, \quad \overline{e} \to \omega^{-1}\overline{e}, \quad Y_e \to Y_e \tag{2.7}$$

where $\omega \equiv e^{2\pi i/3}$. This symmetry restricts any L violating processes to only occur in multiples of 3, which are not possible from any of the soft terms, implying a strong suppression of lepton violation. Proton decay, $p \to e^+ + \pi^0$, requires L violation to occur. A suppression of L violation thus, effectively, renders the proton stable. The neutrino masses can be restored to the theory through a see-saw mechanism framework introducing three new right-handed neutrino fields $\bar{N}$, resulting in the following modification to the superpotential [37]:

$$W_{\text{lept}} = Y_e L H_d \bar{e} + Y_N L H_u \bar{N} + \frac{1}{2} M_N \bar{N}\bar{N} \tag{2.8}$$

where Y_e, Y_N and M_N are the Yukawa couplings for the neutral leptons promoted to spurion fields like those in the quark and charged lepton sector. In this case, lepton violating terms are allowed, which therefore must place constraints on the right-handed

neutrinos due to the experimental bound on proton decay.

2.4 Experimental signatures

As mentioned, the characteristic experimental signature of the MSSM assuming R-parity conservation is that of missing E_T due to the escaping LSP. The LSP is almost always assumed to be the electrically neutral and colorless neutralino, as an otherwise exotic heavy charged particle would easily reveal its existence within the detector, which has not been the case. Thus, searches for MSSM require understanding the decay channels of heavier sparticles into the lightest neutralino. There are many potential decays that can occur, a few phenomenologically interesting possibilities are highlighted here:

- **Chargino or Neutralino Decays**

$$\tilde{C}_1^{\pm} \to \ell^{\pm}\nu\tilde{N}_1, \qquad \tilde{N}_2 \to \ell^+\ell^-\tilde{N}_1 \tag{2.9}$$

$$\tilde{C}_1 \to jj\tilde{N}_1, \qquad \tilde{N}_2 \to jj\tilde{N}_1 \tag{2.10}$$

- **Slepton Decays**

$$\tilde{\ell} \to \ell\tilde{N}_1, \qquad \tilde{\nu} \to \nu\tilde{N}_1 \tag{2.11}$$

- **Squark Decays**

$$\tilde{q} \to q\tilde{g}, \qquad \tilde{q} \to q\tilde{N}_i \tag{2.12}$$

- **Gluino Decays**

$$\tilde{g} \to q\tilde{q}, \qquad \tilde{g} \to qq\tilde{N}_i \tag{2.13}$$

To summarize, the LSP will always appear at the end of every decay chain for the heavier sparticles and accompanied with, typically, jets or leptons.

RPV MFV SUSY models have unambiguously different phenomenology from R-parity conserving ones. The most apparent distinction is the absence of a stable LSP, which therefore removes the effectiveness of reconstructing missing E_T as a method for SUSY detection. Where before, the stability of the LSP also assumed electrical and color neutrality to conform to experimental observations, these assumptions are unnecessary in this model since the LSP may decay. Thus, all cases, whether the LSP is a squark, slepton, neutralino, chargino, or gluino can be considered. The MFV hypothesis regulates the squark and slepton masses through the Yukawa couplings, which most naturally accommodates either the stop or sbottom squarks or the stau as the LSP rather than any of the other squarks or sleptons. A representative selection of possible LSPs with their dominant decay mode diagram and lifetimes are shown in the following table. The gluino is similar to the neutralino, while the chargino has a similar lifetime to the neutralino with a slightly modified final state [37].

Diagrams of the lifetime behavior for the different LSPs' dependence upon the parameters are shown in Figure 2.3.

Large areas of parameter space for the stop correspond to prompt decays near the primary vertex. In such cases, referring to the leading order diagram from the table, the experimental signal would most likely emerge in the form of a resonance of multi-jet final states with many b jets near the stop mass. An analysis in search for these signatures would necessarily encounter enormous background sources of QCD interactions in hadron colliders, making such signals troublesome to pursue.

However, the interesting region of parameter space to consider for each of the possible LSPs has smaller $\tan\beta$, corresponding to longer-lived LSPs. This would manifest as a macroscopically displaced vertex with several jets within the detector. This distinctive attribute provides a useful handle for discriminating a potential SUSY signal above

LSP	Leading Order Decay	Lifetime
Stop	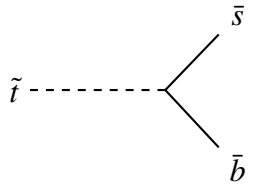	$\tau_{\tilde{t}} \sim (2\ \mu\text{m})\left(\frac{10}{\tan\beta}\right)^4\left(\frac{300\ \text{GeV}}{m_{\tilde{t}}}\right)\left(\frac{1}{2\sin^2\theta_{\tilde{t}}}\right)$
Sbottom	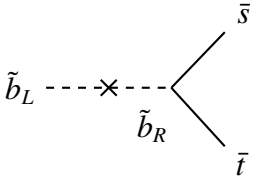	$\tau_{\tilde{b}_L} \sim (41\ \mu\text{m})\left(\frac{10}{\tan\beta}\right)^6\left(\frac{300\ \text{GeV}}{m_{\tilde{b}_L}}\right)$
Neutralino	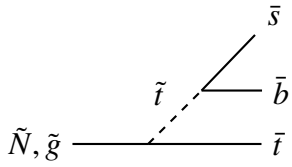	$\tau_{\tilde{N}} \sim (12\ \mu\text{m})\left(\frac{20}{\tan\beta}\right)^4\left(\frac{300\ \text{GeV}}{m_{\tilde{N}}}\right)$
Stau	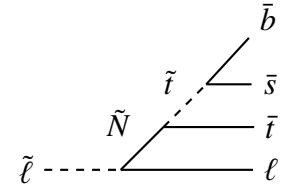	$\tau_{\tilde{\tau}} \sim (44\ \mu\text{m})\left(\frac{45}{\tan\beta}\right)^4\left(\frac{500\ \text{GeV}}{m_{\tilde{\tau}}}\right)$

Table 2.3: Table of select possible LSPs in the MFV SUSY model with the leading order diagrams for their decays and respective lifetime scaling.

the SM background in the relatively unexplored intermediate region between prompt searches and E_T^{miss} searches.

This particular analysis considers models where the LSP is the gluino/neutralino or the top squark, which decay into either "multijet" final states or "dijet" final states, respectively. The production modes for these superpartners are similar to those in the MSSM and assumed to occur via R-parity conserving couplings. The gluino and top squark pair production cross sections are computed at $\text{NNLO}_{\text{approx}}$+NNLL precision [25, 23, 24] while the neutralino pair production cross section is computed at NLO+NLL precision in a limit of mass-degenerate higgsino states $\tilde{\chi}_1^{\pm}, \tilde{\chi}_1^0$, and $\tilde{\chi}_2^0$ where all the other sparticles assumed to be heavy and decoupled [45, 46].

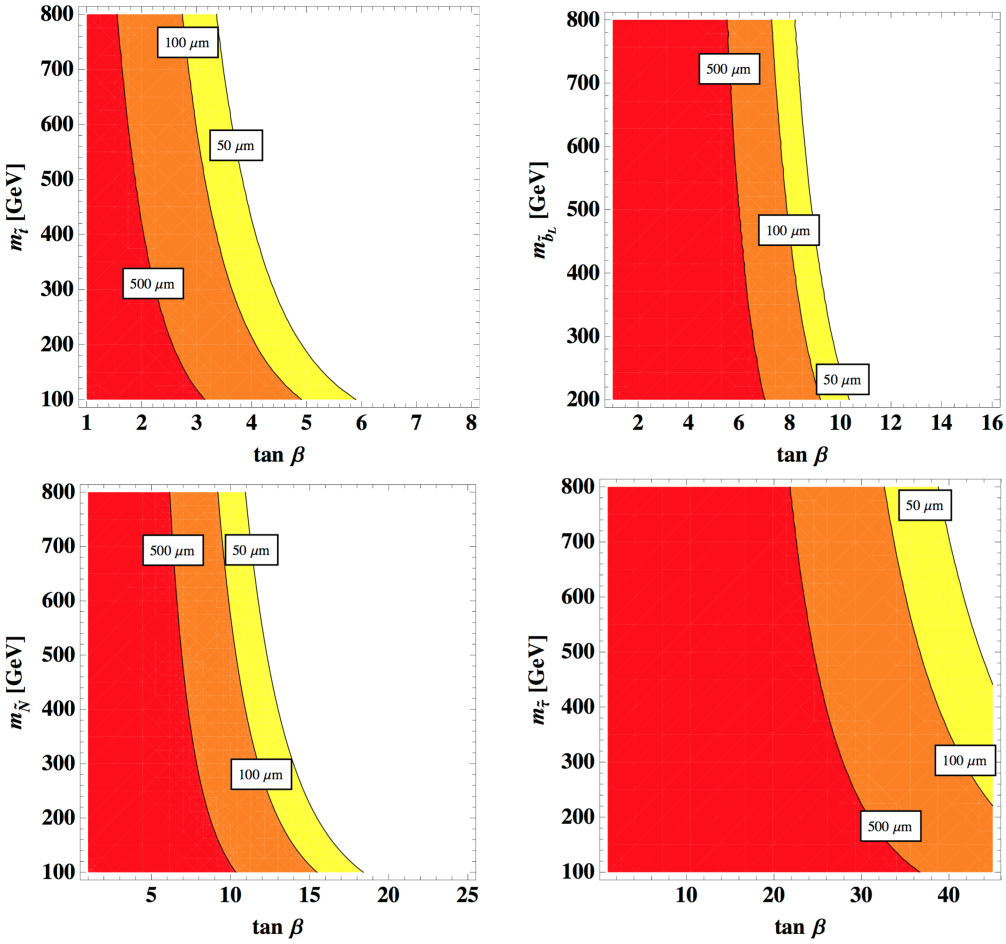

Figure 2.3: Lifetime plots as a function of sparticle mass and $\tan\beta$ for different possible LSPs as given in Table 3: stop (top left), left-handed sbottom (top right), neutralino (bottom left), and stau (bottom right).

CHAPTER 3

EXPERIMENTAL APPARATUS

In order to probe the interaction length scales and access the energies required to pair-produce the set particles predicted by SUSY, experiments build larger and larger accelerators to ramp up the energy of charged SM particles and collide them inside of giant detectors. These machines create the conditions in which rare processes may occur, doing so with an astonishing rate of collisions to increase the probability of observing them. This chapter provides a glimpse of both the collider used to produce these high-energy collisions as well as the detector surrounding the collision point that records and reconstructs the interesting physics processes occurring in the interaction.

3.1 The Large Hadron Collider

The LHC is the world's largest circular particle accelerator located near the city of Geneva, Switzerland on the Franco-Swiss border. It lies approximately 100 m beneath the surface in a tunnel 27 km long in circumference. With two separate rings circulating counterrotating beams of particles, the LHC accelerates and collides primarily beams of protons, but is also capable of colliding heavy nuclei such as lead or even xenon. When operating at full design capacity, the collider can deliver proton-proton collisions at a peak center-of-mass energy of $\sqrt{s} = 14$ TeV, although in the data analyzed in this thesis, the center-of-mass energy of the collisions was 13 TeV.

Figure 3.1 shows a schematic of the CERN accelerator complex, which includes the LHC and the various injectors used in the multistage process of accelerating the protons up to full energy. The process of accelerating the beams of protons to their full collision energy occurs throughout a cascade of these different accelerator machines beginning

first with a tank of hydrogen gas. The gas particles are injected into a duoplasmatron, which uses an electric field to strip the electrons and disassemble the diatomic hydrogen molecules into its constituent protons. These protons enter into a linear accelerator, Linac2, which accelerates them up to 50 MeV before injecting them into the Proton Synchrotron Booster. They continue to increase in energy to 1.4 GeV, after which they are injected into the Proton Synchrotron to accelerate them to 25 GeV, then to the Super Proton Synchrotron to 450 GeV. Finally, they are injected into the Large Hadron Collider into two separate beams circulating in opposite directions and the beam energy is ramped up to its full energy of 6.5 TeV and made to collide within the detectors. The entire process can last nearly two hours with each new fill of protons into the machine providing five to six hours, on average, of stable beam collision data taking.

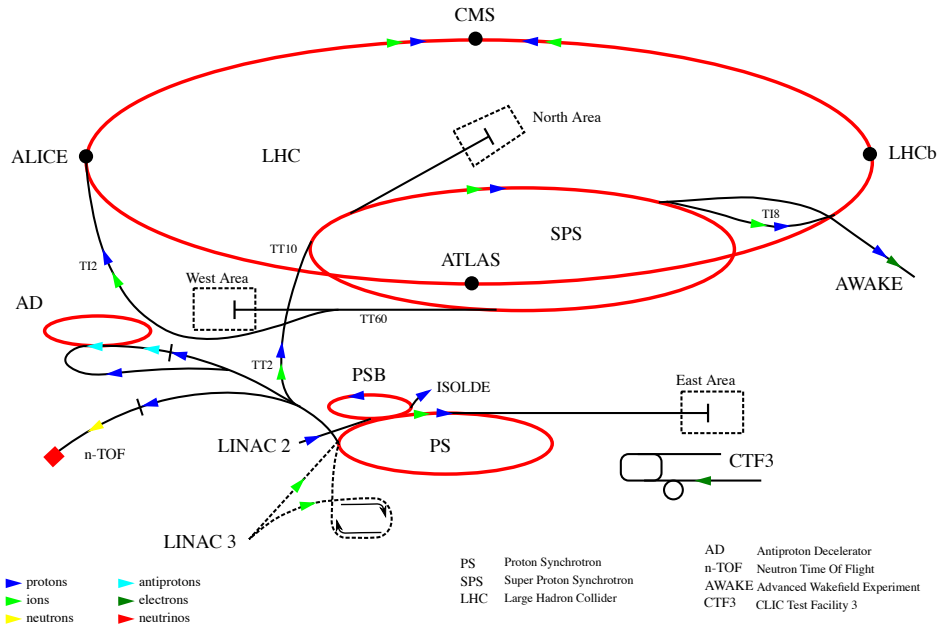

Figure 3.1: Schematic view of the CERN accelerator complex.

In the LHC, the beams are controlled by an intricate arrangement of radio-frequency

(RF) cavities and magnets throughout the entire ring. Eight superconducting RV cavities provide a 5 MV/m field to accelerate the particles around the ring. This field oscillates at 400 MHz resulting in the discretized *bunch* structure of the beam. Operating at full design capacity, each beam can be made up of up to 2808 colliding bunches of protons, where each bunch consists of approximately 1.1×10^{11} protons. A total of 9,593 powerful magnets distributed throughout the LHC steer and control the proton bunches around the ring. Dipole magnets constitute 1,232 of these, and perform the primary role of maintaining the circular trajectory around the ring with magnetic fields up to 8.3 T, a schematic diagram of the cross section of an LHC dipole is shown in Fig. 3.2. These magnets are constructed from superconducting NbTi cables operating at a temperature of 1.9 K by superfluid helium cooling, which is necessary to sustain the 11,850 A of current required to produce the large magnetic fields. Additional quadrupole and many higher order multipole magnets provide finer adjustments of the beam characteristics in order to focus and squeeze the beam. This tuning of the beam parameters works to adjust the instantaneous luminosity described by equation 3.1.

$$\mathcal{L} = \frac{N_b^2 n_b f_{\text{rev}} \gamma}{4\pi\epsilon_n \beta^* \sqrt{1 + \left(\frac{\theta_c \sigma_z}{2\sigma^*}\right)^2}} \tag{3.1}$$

The instantaneous luminosity is dependent only on these beam parameters as specified in Table 3.1 with an instantaneous luminosity on the order of $\mathcal{L} = 10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$ for standard LHC operation under the listed design conditions. The number of events for a process with a cross section σ is thus given by:

$$N = \sigma \int \mathcal{L} dt \tag{3.2}$$

where $\int \mathcal{L} dt$ corresponds to the time-integrated luminosity. The LHC collides its

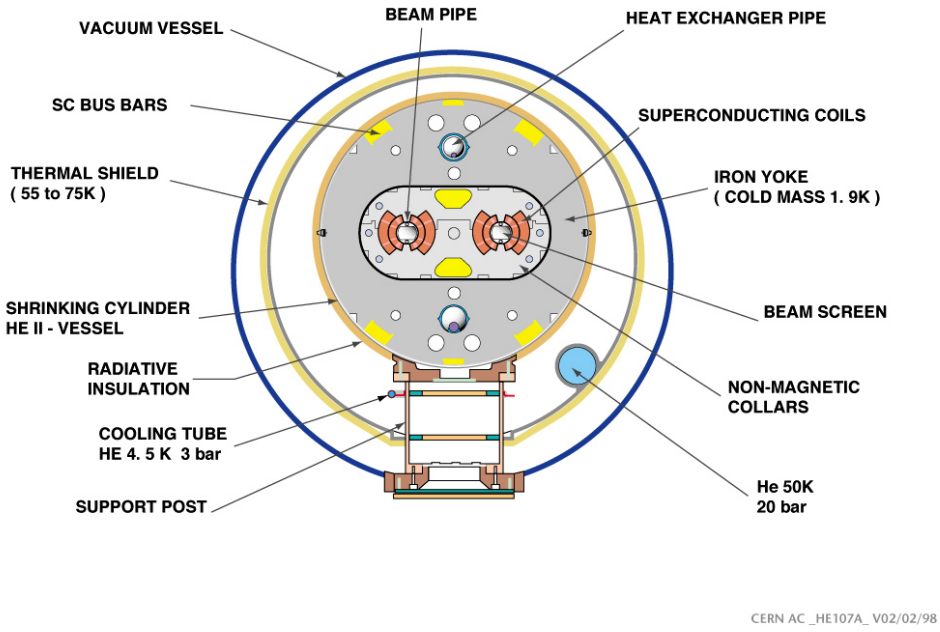

Figure 3.2: Schematic of the LHC dipole cross section.

Table 3.1: LHC beam parameters and the corresponding value under design conditions. [3]

Parameter	Description	Value
N_b	Number of protons per bunch	1.15×10^{11}
n_b	Number of bunches per beam	2808
f_{rev}	Bunch revolution frequency	11245 Hz
γ	Relativistic gamma factor	7461
ϵ_n	Normalized transverse beam emittance	3.75 μm rad
β^*	Optical β function evaluate at interaction point	0.55 m
θ_c	Beam crossing angle	300 μrad
σ_z	RMS of longitudinal bunch length	7.55 cm
σ^*	RMS of transverse beam size	16.7 μm

beams at four interaction points on its circumference once every ~25 ns, referred to as a *bunch crossing*, around which are positioned different detectors designed for different physics research goals. They are ALICE, ATLAS, CMS, and LHCb. Figure 3.3 shows the integrated luminosity of collisions delivered to CMS each year since the beginning of data taking with the LHC. This thesis details the analysis of data collected in 2017 and 2018, however the final result uses the full run 2 dataset using data from 2015 to 2018.

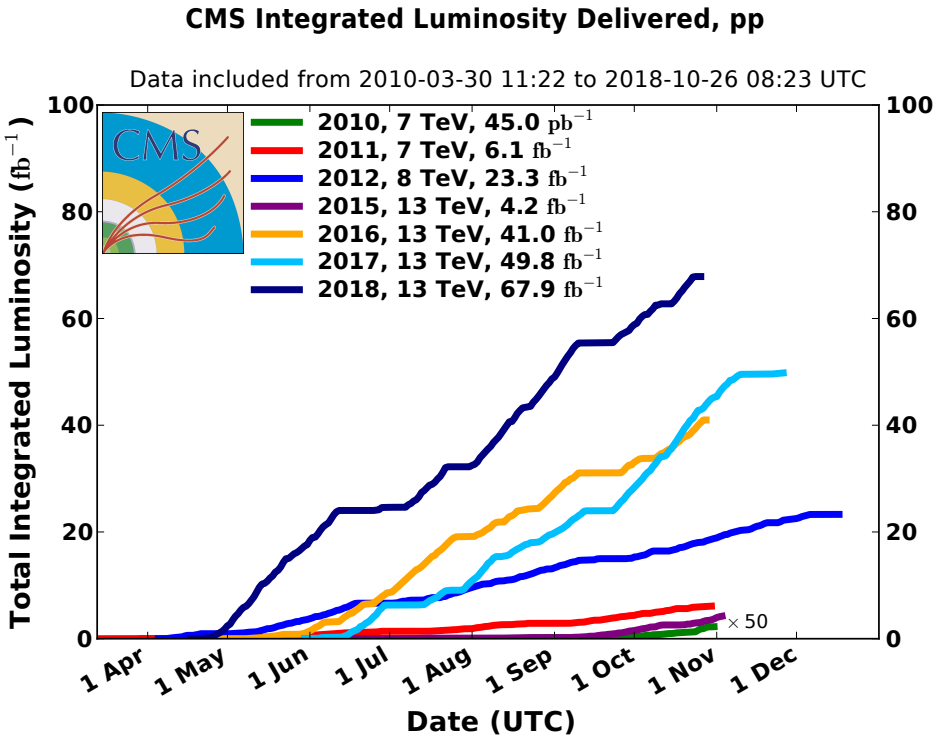

Figure 3.3: Integrated luminosity delivered to CMS throughout the year by the LHC for each year of datataking.

3.2 The Compact Muon Solenoid

The Compact Muon Solenoid, or CMS, broadly refers to multiple entities: the experiment itself and the particle physics research it aims to conduct; it also refers to the

collaboration of multiple thousands of scientists, engineers, and students across hundreds of nations and academic institutions working on the experiment; and of course, it refers to the physical detector employed by the collaboration in order to carry out its research program. This section will focus on the third meaning, namely, the apparatus that collects data.

The CMS detector is one of two general purpose detectors for particle physics research on the LHC. It is situated at Point 5 on the opposite end of the LHC ring from the main CERN campus, near the village of Cessy, France. CMS is a large hermetic detector, broadly cylindrical in shape (or dodecagonal prismatic), weighing 14,000 metric tonnes with a total length of 21 m and average diameter of 15.0 m–sizable, and yet compact, relative to its counterpart experiment, ATLAS, which occupies nearly six times the volume. The LHC beam pipe crosses directly through the central axis of the detector to collide the beams within an interaction region at the center of CMS. This nominal interaction point represents the geometric origin of the coordinate system within CMS with which to describe the subsequent behavior of particles produced in the interaction. In Cartesian form, the x-axis of this coordinate system points toward the center of the LHC ring, the y-axis points vertically upward, while the z-axis points along the beam axis under the constraint that the resulting coordinate system is right-handed (in this case, the axis points toward the Jura mountains). The x-y plane is often referred to as the transverse plane, which is frequently used in practice as it provides helpful descriptors of important particle properties including the transverse momentum p_{T}, the component of the particle momentum projected onto this plane, as well as the azimuthal angle ϕ, which is measured from the x-axis in the transverse plane. Pseudorapidity or η is the variable to describe angles in the forward direction toward the endcaps of the cylindrical volume and is defined as:

$$\eta = -\ln\left[\tan\left(\frac{\theta}{2}\right)\right] \tag{3.3}$$

where θ is the polar angle measured from the z-axis. The pseudorapidity is preferred over θ due to its approximate invariance under Lorentz transformations, which is useful in collider physics where boosts along the z direction varies per collision. The sign of η follows the sign convention for z. As the particle momentum increases to higher proportions compared to its mass, the pseudorapidity approaches the rapidity $y = \frac{1}{2}\ln\left(\frac{E+p_z}{E-p_z}\right)$. A Lorentz-invariant angular distance is also often used and defined as:

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} \tag{3.4}$$

The entire detector is comprised of multiple concentric layers of specialized subsystems built out of sensitive instrumentation and readout electronics surrounding the interaction region in order to measure different properties of particles materializing out of the colliding beams. Figure 3.4 shows a diagram of the whole CMS detector, with cutaways to observe the layers underneath the surface. While it aims to capture most of the particles produced, sometimes particles can successfully make their way past the physical structure and material comprising the detector without any interaction and may escape unseen entirely. This may be due to the weakly interacting particles such as the neutrino within the Standard Model, or perhaps new, exotic, weakly-interacting particles whose existence may be inferred in further analysis of the data. Nonetheless, the measurements from the sensors are collected from each detector component and each piece contributes to the reconstruction of the collision event and a picture of the emerging particle debris. The following sections will take an unconventional path in providing a description for each of the different subdetectors of CMS, starting from the outermost systems and peeling away the layers to move inward toward central interaction region.

See Refs. [31, 22] for a more detailed treatment on the CMS detector.

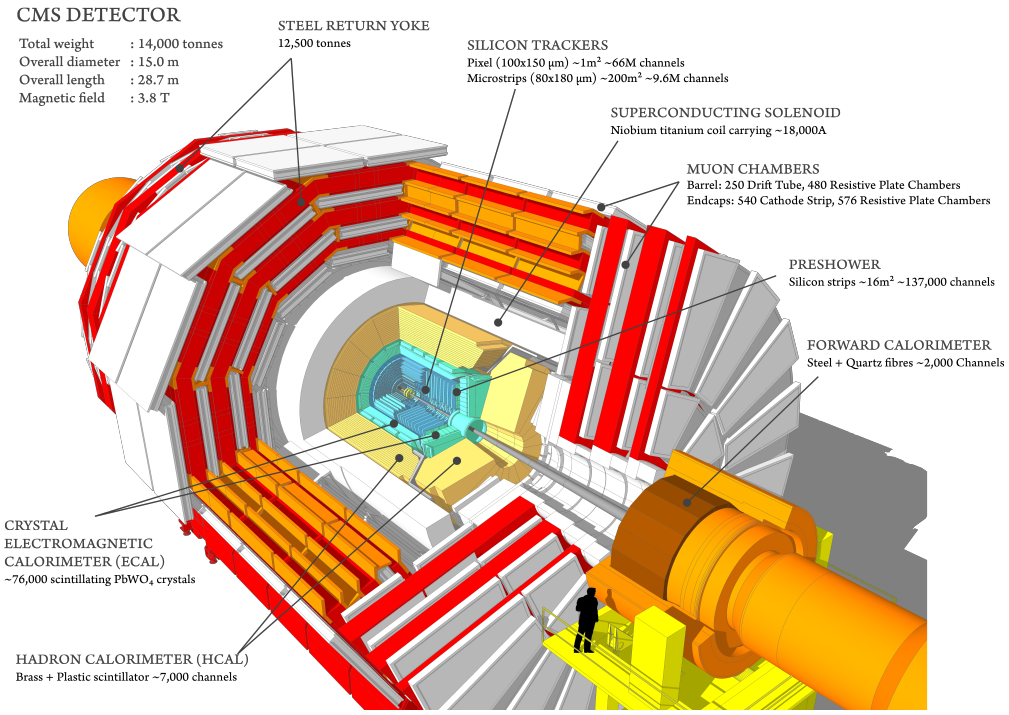

Figure 3.4: A diagram of the CMS detector with the layers exposed in a cutaway to reveal the different subdetectors [1].

3.2.1 Muon system

Particles detected after penetrating through many internal layers and surviving to the most external chambers are most often muons. The CMS detector has a very robust and dedicated muon system with 80% of the detector, by volume, in the outermost subdetectors were designed with the specific objectives of recording the passage of these particles, measure their momentum, and provide triggering to signal whether an event with a suspected muon should be kept for closer analysis [2, 34, 66]. The system consists of three separate subdetectors: the Drift Tubes (DT), the Cathode Strip Chambers (CSC), and the Resistive Plate Chambers (RPC), which are shown and labeled in Fig. 3.5. The muon system is separated into four concentric layers or *stations* to provide additional

redundancy. All three operate using gaseous tracking-based techniques but employ them with slightly different technologies to address the different magnetic field conditions or background levels intrinsic to their location on the detector. These technologies are described here.

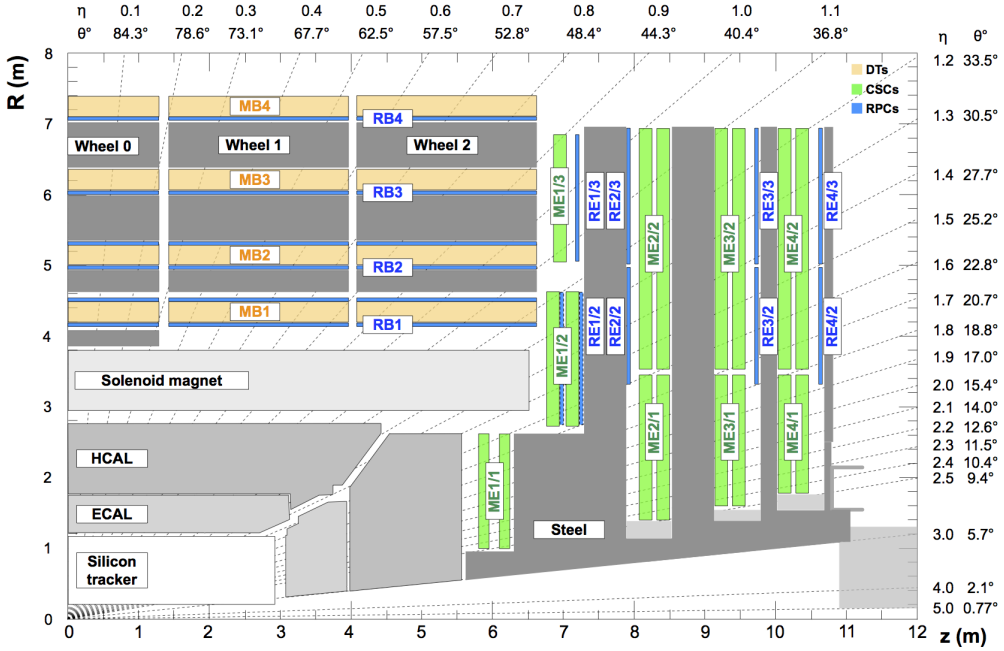

Figure 3.5: A schematic quarter view of the muon subdetectors in the r-z plane [66].

Drift Tubes

The DT chambers are found in the central barrel region with $|\eta| < 1.2$. An individual unit of the DT system is a rectangular aluminum tube, called a drift cell, with a transverse cross section of $42 \times 13\,\text{mm}^2$ containing a gas mixture of 85% argon (Ar) and 15% carbon dioxide (CO_2). These tubes are equipped with a $50\,\mu$m diameter gold-plated stainless steel wire stretched along its central axis, which is charged to a voltage of +3600 V. Additional electrode and cathode strips, charged to ±1800 V, are placed along the walls of the cell to provide finer control and shaping of the electric field within the volume of the chamber. As charged particles traverse through the gas within the cell

chamber, they ionize the gas, freeing electrons that will drift toward the wire. With precise timing measurements and calibrations, the ionizing particle's distance away from the anode wire, along the transverse axis, can be inferred from the measured duration of many drifting electrons within the well-tuned electric field. Rows of these chambers provide coverage in a single direction, either ϕ or z, but orienting multiple layers of these cell chambers perpendicularly will constrain the position in the two-dimensional plane orthogonal to the radial direction. Four layers of drift cells that are staggered by a half-cell form a *superlayer* (SL). Combining two or three SLs forms a single muon station, which provides several degrees of redundancy and offers much more precise measurements of the muon trajectory through the barrel.

Cathode Strip Chambers

The CSCs provide coverage in the forward endcaps within the range $0.9 < |\eta| < 2.4$. The chambers of the CSC detector are composed of alternating cathode strip and anode wire planes. Seven trapezoidal cathode panels, each containing 80 cathode strips along the radial direction, form six gaps in between the layers where a gas mixture of 50% CO_2, 40% Ar, and 10% carbon tetrafluoride (CF_4) is contained. The CSC operates under the same principle as the DT system in detecting the passage of a muon through its layers. The muon reveals its position from timing measurements of the freed electrons from the ionization as they drift to the anode. However, the system arrangement described here provides a much faster readout with greater segmentation, compared to the DTs, in order to handle the large background flux of radiation in the forward directions, and an increased resilience to non-uniformities in the magnetic field as well.

Resistive Plate Chambers

The RPCs are parallel-plate gaseous detectors located in both the barrel and endcaps with coverage spanning $|\eta| < 1.6$. The detector is composed of two 2 mm-thick plastic, highly resistive, Bakelite plates held at opposite voltages, ±9.6 kV, enclosing a gas mixture of 95.2% Freon ($C_2H_2F_4$), 4.5% isobutane (i-C_4H_10), and 0.3% hexafluoride SF_6. As a muon passes through and ionizes the gas mixture, the freed electrons accelerate in the direction of the electric field toward the anode, freeing additional electrons along the way, cascading into an avalanche. Since the electrodes have high resistivity, the detected signal does not come from the freed electrons themselves, but the induced movement of charge in the detection strips in the electrode. As there is no drift time associated with this detection mechanism, it has very fine timing resolutions of ~1 ns, much smaller than the 25 ns bunch crossing period. However, this comes at the expense of much greater imprecision in the spatial resolution of the incident particle. For this reason, the primary function of the RPC subdetector is to provide a complementary trigger system to the rest of the muon system. Its timing precision more reliably identifies the appropriate bunch crossing from which an associated charged particle emerged, and thus, which event to trigger.

3.2.2 Solenoid Magnet

Moving inward, one of the central features of CMS (and its final namesake) is its large superconducting solenoid magnet. The solenoid is 12.5 m long and 6 m in diameter made up of NbTi cables cooled using liquid helium to a temperature of 4 K. When operating, it carries over 18,000 A of current, equivalent to 2.3 GJ of stored energy, producing a stable axial field of 3.8 T within the volume of the solenoid, a diagram of which

can be seen in Fig. 3.6. The muon system described previously is embedded within large 10,000 ton steel structures that form the return yokes for the magnetic field, which control and guide the magnetic field outside the solenoid. The powerful field deflects the paths of charged particles by a magnitude proportional to their momenta, which therefore reframes the measurement of particle momenta to a geometric calculation of the curvature of its trajectory. The volume of the solenoid encloses the remaining subdetectors.

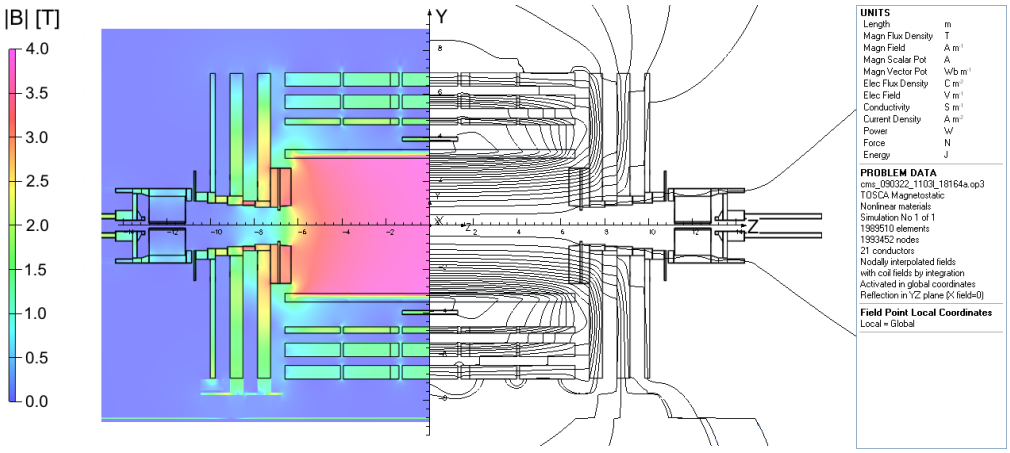

Figure 3.6: Diagram of the solenoid magnet within CMS, revealing the magnetic field strength and the magnetic field lines. [33]

3.2.3 Hadronic Calorimeter

Just inside the solenoid is the first of two calorimeters, the hadronic calorimeter (HCAL). Its purpose is to measure the energy of hadrons that interact via the strong interaction, which make it important in searches involving jets or missing transverse energy analyses. The HCAL is partitioned into four parts: the barrel (HB), the outer barrel (HO), the endcaps (HE), and the forward calorimeters (HF). Taken altogether, the full HCAL system provides very wide coverage in pseudorapidity, $|\eta| < 5.2$. The system is a sampling calorimeter consisting of alternating layers of brass absorber plates and plastic

scintillating layers. As an incident particle hits the absorber plates, this generates an interaction producing many secondary particles, which can further cascade or *shower* into more particles as additional interactions with the absorber material occur throughout the layers. Photons are produced as these shower particles traverse through the alternating scintillating layers, with the amount of light emitted proportional to the energy of the particles. These photons are collected by wavelength shifting fibers and funneled into a hybrid photodiode where the optical signals are converted to electrical signals to be delivered to the CMS data acquisition system. A quarter view of the HCAL is shown in Figure 3.7.

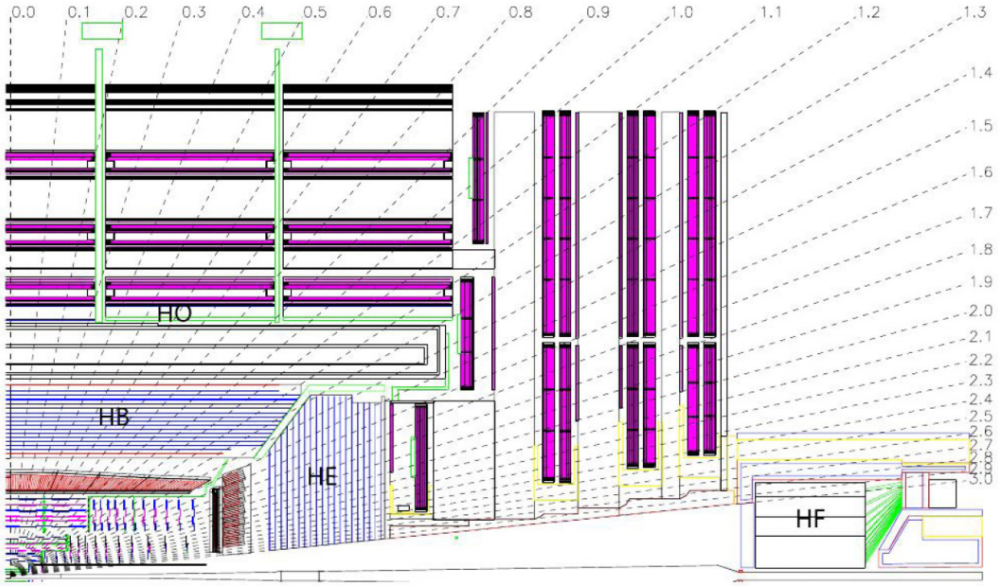

Figure 3.7: Quarter view of the CMS hadronic calorimeter [31].

3.2.4 Electromagnetic Calorimeter

Immediately within the HCAL volume sits the electromagnetic calorimeter (ECAL), designed to measure the energy of electrons and photons. Unlike the HCAL, the ECAL is a homogeneous calorimeter with lead tungstate ($PbWO_4$) crystals as its scintillating ma-

terial. A total of 61,200 crystals form the barrel of the ECAL (EB), covering $|\eta| < 1.479$, which are aligned with a quasi-projective geometry that reduces the space between adjacent crystals. The crystals are tapered at 230 mm in length with a $22 \times 22\,\text{mm}^2$ front face and $26 \times 26\,\text{mm}^2$ back face. The forward endcaps of the ECAL (EE) cover $1.479 < |\eta| < 3.0$ with 7,324 crystals on each end. In addition to the barrel and endcaps, the ECAL has an additional system called the preshower (ES) that sits just in front of the endcaps on either side, covering $1.653 < |\eta| < 2.6$. The ES enhances the spatial resolution and helps to improve the identification of photons and neutral pions in the endcaps. A schematic of the ECAL is shown in Figure 3.8.

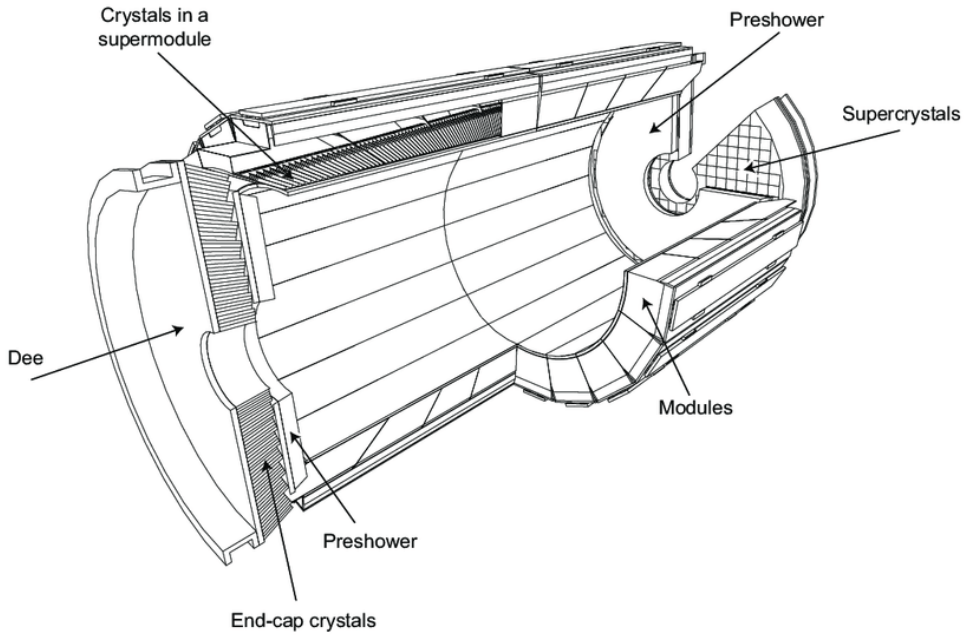

Figure 3.8: Schematic diagram of the CMS electromagnetic calorimeter [32]

.

The use of $PbWO_4$ provides many desirable qualities that make it ideal for electromagnetic calorimetry within CMS. While they are very dense and metallic, they are still highly transparent. They have a short radiation length of 0.89 cm and a small Molière radius of 2.2 cm, which can restrain the electromagnetic shower size and improve the position measurement within a smaller calorimeter. The crystals are radiation hard and

provide a fast response time within the 25 ns bunch spacing of the LHC.

The scintillation light measured from the passage of electrons or photons in the medium is collected by a pair of avalanche photodiodes in the barrel and vacuum phototriodes in the endcaps, with the amount of light proportional to the energy of the incident particle. The light output is very temperature sensitive, and the crystals must be cooled within a small margin of $18 \pm 0.05^{\circ}$ C to maintain the desired energy resolution of the calorimeter.

3.2.5 Silicon Strip Tracker

The displaced vertex reconstruction at the crux of the search described in this dissertation relies heavily on the remaining subdetectors contained within the ECAL, which are both tracking detectors. The outermost tracker is the silicon strip tracker, 5.8 m in length and 2.5 m in diameter, covering $|\eta| < 2.5$, divided into four systems: the tracker inner barrel (TIB), the inner disks (TID), the outer barrel (TOB), and the endcaps (TEC), which can be seen in Figure 3.9. The TIB consists of four barrel layers capped with three disks on each end by the TID spanning a radius of $20\,\text{cm} < r < 55\,\text{cm}$. The TOB extends beyond the TIB with six additional barrel layers out to a radius of $r < 116\,\text{cm}$, while the TEC cover each end with nine disks. 9.3 million p-on-n silicon micro-strips are assembled on 15,148 strip modules distributed across the layers of the TIB/TOB and the disk of the TID/TEC, resulting in a total effective active area of $198\,\text{m}^{2}$, the largest silicon tracker in the world. While the silicon strip tracker provides a reasonable approximation of a charged particle trajectory, it is limited in its precision due to the coarse granularity of the strips and cannot operate alone in reconstructing a track without the final subdetector within CMS.

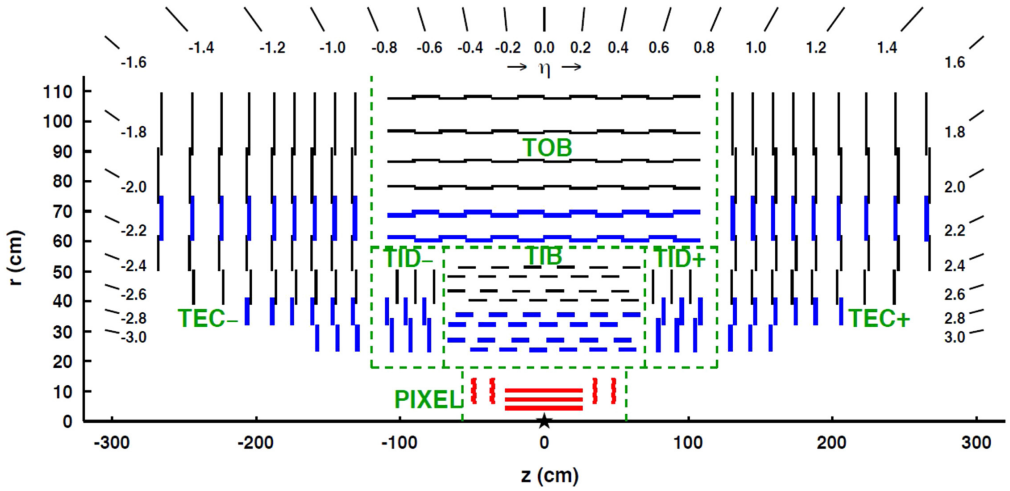

Figure 3.9: A schematic cross section in the *r*-*z* plane showing the different sections of the CMS tracker. [35]

3.2.6 Silicon Pixel Tracker

Enveloped within the strip tracker volume lies the silicon pixel tracker, the innermost detector of CMS, which provides much finer sensor granularity for higher precision track reconstruction. The discussion of the pixel tracker here is meaningful in several ways: the reconstruction of displaced vertices relies heavily on well-reconstructed tracks that are seeded from hit measurements of charged particle passing through the pixel detector; during the winter year-end technical stop between 2016 and 2017, the original (phase-0) pixel detector installed in 2009 was replaced with an upgraded system, referred to as the phase-1 pixel detector; the author of this dissertation had the privilege of working extensively on the pixel detector upgrade leading up to its installation, throughout the early commissioning stages, and first data-taking periods and experienced the difficulties that ensued. For these reasons, there will be a greater level of detail expounded on the pixel detector in this section than for previous subsystems.

At close proximity to the interaction region, the pixel tracker endures a harsh environment with significant radiation culminating in very high track density per event. By

design, the original phase-0 detector was able to handle a maximum instantaneous luminosity of $1 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ with a maximum number of 25 pileup interactions in each bunch crossing where the bunches are spaced 25 ns apart. As the LHC has undergone additional upgrades in the intervening years, these beam and collision parameters have been modified such that the maximum design instantaneous luminosity has effectively doubled and therefore has exceeded the original pixel detector's nominal design capability, resulting in diminished track reconstruction efficiency. Thus, in order to maintain the same level of performance while the LHC continues to increase the beam intensity, the phase-1 pixel detector was developed to deliver high quality tracking data for CMS while withstanding the cumulative radiation damage until the end of LHC Run 3, when it will be decommissioned and replaced by the phase-2 detector, in preparation for the High-Luminosity LHC.

The phase-1 detector is an evolutionary improvement upon the original detector in several key aspects. While occupying the same volume within the strip tracker as the original pixel detector, it includes an additional layer in both the endcap (FPix) and barrel (BPix) of the detector, while capitalizing on the recent installation of a narrower beam pipe of radius 23 mm by moving the innermost barrel layer closer to the interaction point. A schematic comparision of the previous detector with the upgrade is shown in Fig. 3.10. In order to manage the high particle flux from larger instantaneous luminosities and pileup interactions, each layer is equipped with new digital readout chips (ROCs) [54], which can transmit data faster than the previous analog-based ROCs to manage the increased hit rates. To stay within the material budget despite the addition of an entirely new layer of sensor modules, the detector is supported on a new lightweight carbon support structure and cooled using a new CO_2-based cooling system, replacing the C_6F_{14} system used previously. Finally, the backend data acquisition (DAQ) infrastructure was also improved with the adoption of the Micro Telecom Com-

puting Architecture (MicroTCA), transitioning from the previous Versa Module Europa (VME) system, permitting increased bandwidth capabilities in collecting data from the detector.

Figure 3.10: CAD schematic of the previous phase-0 pixel detector (top left), the upgrade phase-1 pixel detector (top right), and a photo of the volume in which they occupy within the CMS detector (bottom).

Front End

The front end refers to the components of the detector found within the volume of CMS, which not only include both BPix and FPix but also the auxiliary support cylinders

that house the different electronics needed to facilitate the readout and powering of the detector as shown in Fig. 3.11.

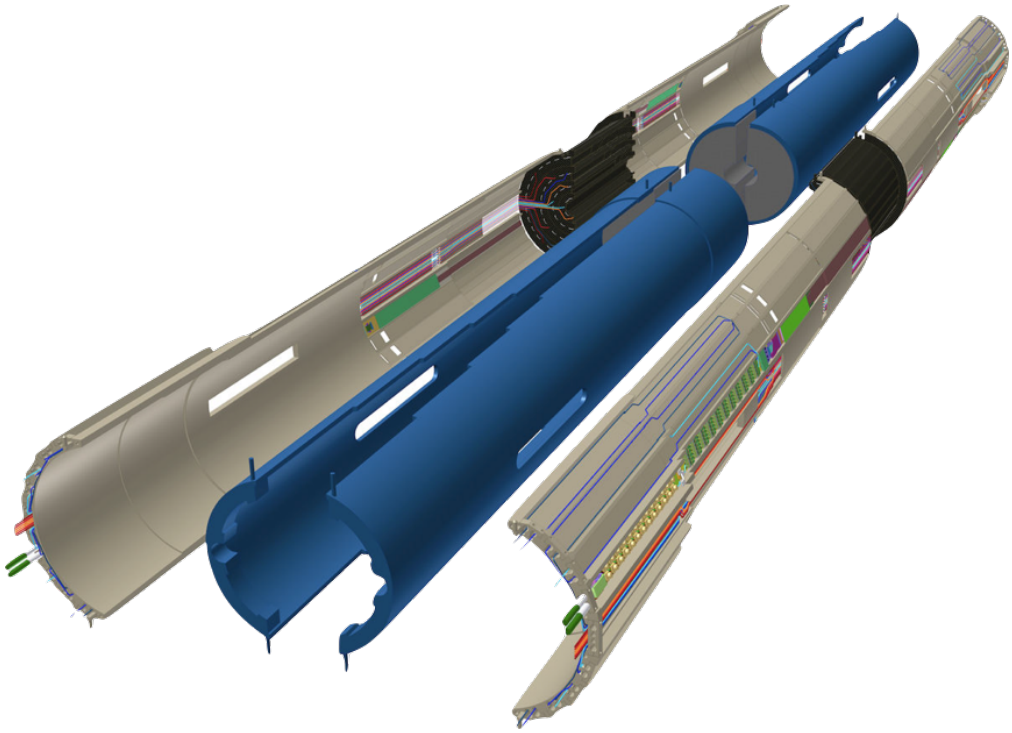

Figure 3.11: Diagram of the pixel detector including the service cylinders, pictured here in blue and gray, which contain the control and readout electronics, powering, and cooling lines.

The design layout of the detector was optimized to provide a four-hit coverage within $|\eta| < 2.5$, consisting of four concentric barrel layers 540 mm in length at radii of 29, 68, 109, and 160 mm, and three disks on each end at distances of 291, 396, and 516 mm from the geometric center of CMS with a disk inner radius of 45 mm and outer radius 161 mm. Figure 3.12 shows the layout of the phase-1 detector compared to the original.

The total resulting surface area coverage of silicon in the phase-1 pixel detector is 1.9 m^2 spread across 124 million individual readout channels. Each channel corresponds to a single n-in-n silicon pixel sensor 285 μm thick and $100 \times 150\,\mu\mathrm{m}^2$ in area. 4160 pixels, arranged into 26 double columns of 2×80 pixels are bump-bonded, using

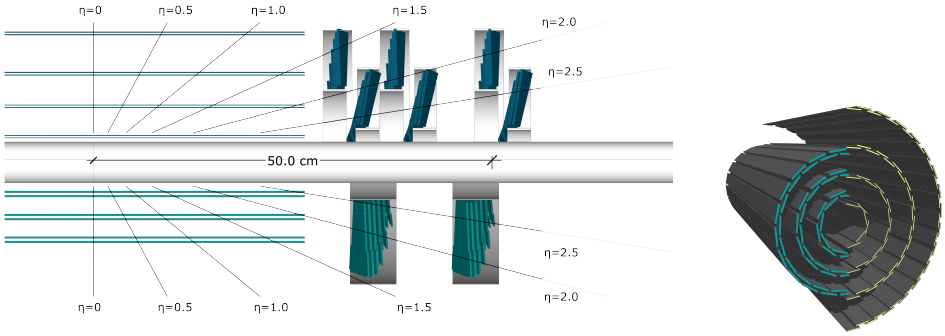

Figure 3.12: Comparison of the phase-1 upgrade pixel layout to the original pixel detector in the *r*-*z* plane (left) and a side by side comparison of barrel layers between the two pixel detectors (right).

a microscopic balls of tin-lead (SnPb) solder, onto a ROC. Two types of ROCs are used throughout the detector: the chip used in the outer three layers of BPix and in FPix is called the PSI46dig, while a specialized chip named the PROC600 is found only in the first barrel layer designed specifically to handle the expected 600 MHz/cm^2 particle fluence at its radial position. The two chip types are broadly similar except for additional functionality granted to the PROC600 to improve the data bandwidth. The chip architecture of the phase-1 ROC remains largely the same as its predecessor with a few crucial improvements. The chip periphery contains larger time stamp and data buffers compared to the original phase-0 chip. The time stamp buffer stores the bunch crossing number in which a charged particle signal was measured. The corresponding data is stored in the data buffer as an analog pulse height at a particular pixel address on the ROC. The larger buffer size accommodates the increased particle fluence as beam intensity increases, and more hit information must be stored while awaiting the trigger latency. Upon receipt of the trigger, the zero-suppressed data is read out for any pixels exceeding the programmed threshold of the chip. This readout has also improved in the upgraded ROCs, going from a 40 MHz analog signal to a 160 Mb/s digitized serial data stream provided by an analog-to-digital converter (ADC) in the chip. Figure 3.13 shows

a diagram of individual pixel cells and a ROC.

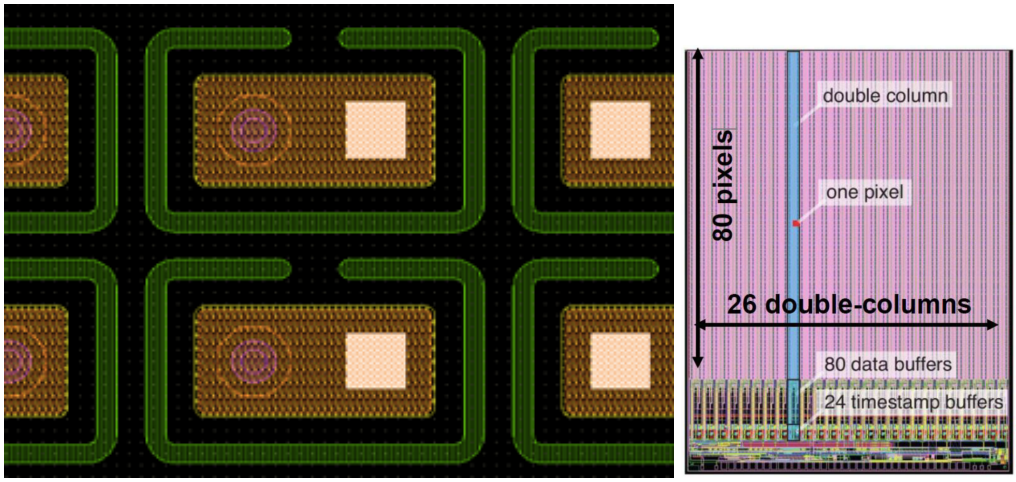

Figure 3.13: Schematic of two pixel cells in an FPix sensor (left) and a ROC schematic (right).

Sixteen ROCs are assembled into a single pixel module, along with one (two in the first BPix layer) token bit manager (TBM) chip. The TBM is a chip that synchronizes the readout of a group of ROCs on the module. It operates by issuing bit referred to as the *token bit* when an L1A trigger signal is received from the CMS back end DAQ. The token is passed sequentially through the ROCs in a chain to initiate the readout. Each ROC provides a ROC header and its pixel-level event data, which the TBM receives, encodes, and multiplexes using a 4b/5b encoding scheme. The TBM ultimately delivers a single 400 Mb/s data stream to the central CMS DAQ, formatted with a TBM header at the beginning that encodes the event number, then a series of ROC headers and associated ROC data, and terminated with a TBM trailer that also provides any error information. Figure 3.14 show the pixel modules used in the different layers of BPix and in FPix, which vary in order to optimize performance given the expected radiation conditions with differences primarily in the number and types of TBMs, the types of ROCs used, or the readout cable technology.

In total, 1856 pixel modules are installed on the layers and disks of the pixel detec-

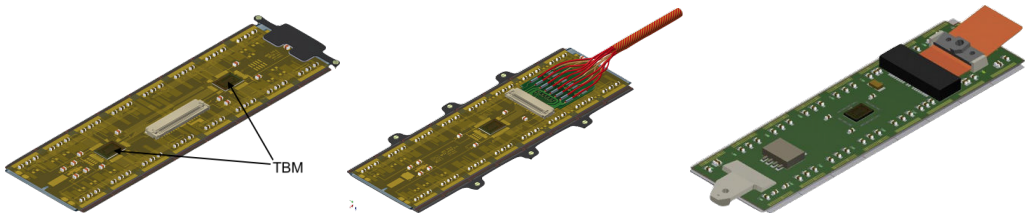

Figure 3.14: Diagrams of pixel modules designed for BPix layer 1 (left), BPix layers 2-4 (middle), and FPix (right).

tor, 1184 are found in BPix with the remaining 672 in FPix. In each layer of BPix, eight detector modules are mounted end-to-end, lengthwise, onto rows or *ladders*, which are arranged into alternating concentric, cylindrical layers. From the innermost layer to the outermost layer, there are 96, 224, 352, and 512 modules in each layer. Figure 3.15 shows a transverse cross section of a half cylinder of BPix, revealing the module arrangement structure for each layer. FPix is divided into 6 total disks, 3 capping each end of the cylinder. Each disk can be deconstructed into an inner and outer radial ring that both have a turbine-like mechanical support structure on which modules are mounted. The turbine geometry of the inner ring has a total of 22 *blades*, while the outer ring holds 34. Modules are mounted onto the front and back panels of the blades. Thus, each disk holds 112 total modules, 44 on the inner ring, and 68 on the outer. A schematic of a half disk of FPix can also be seen in Figure 3.15.

Additional auxiliary electronics are situated inside the service cylinders on either end past the FPix disks, which facilitate communication between the front end detector electronics and the back end DAQ. This includes providing the functionality to program the ROCs and TBMs with the appropriate calibrated configuration parameters; distributing the LHC clock signal and the central DAQ L1A triggers; adjusting timing and delays between clock, trigger, and data signals to maintain stable communication with the detector hardware; and reading in the digitized data from the modules as an electrical signal and converting them into optical signals to deliver to the DAQ system

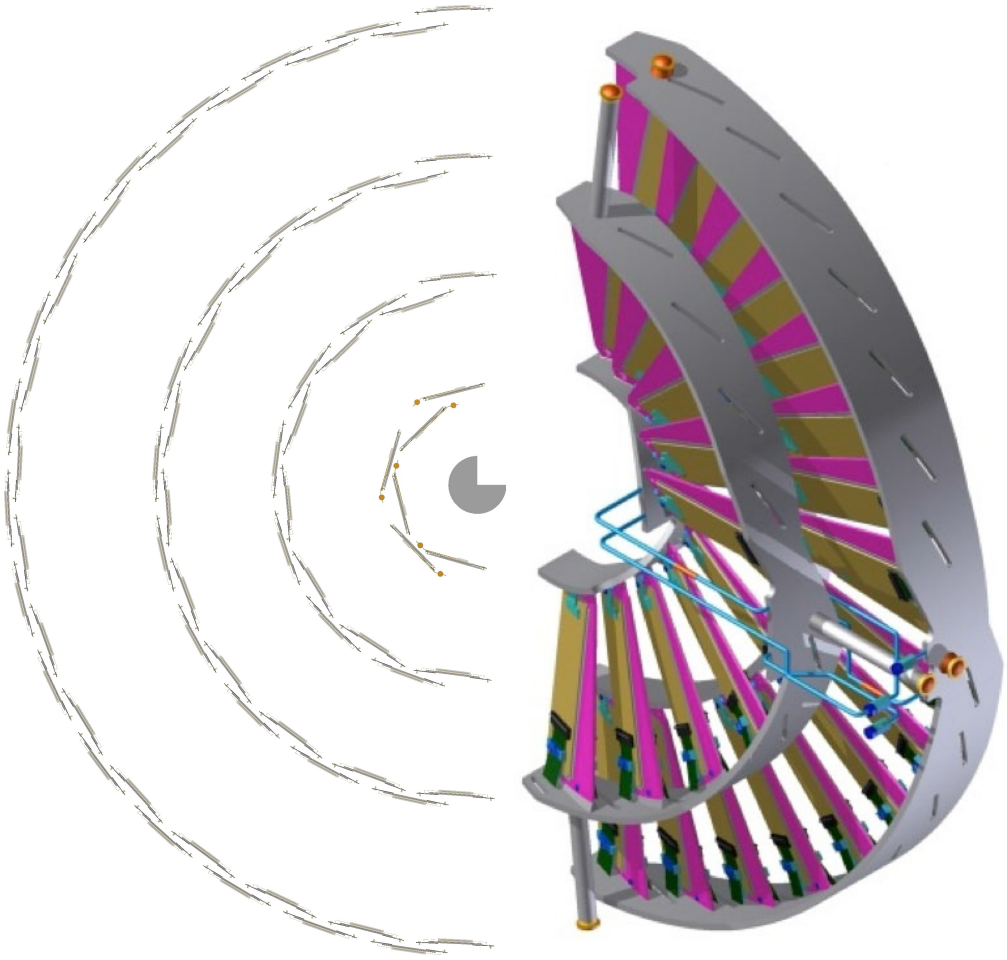

Figure 3.15: Schematic of BPix module arrangement in each half layer (left) and a CAD schematic of an FPix half disk (right).

via optical links, or *pixel opto-hybrids* (POH). Sending information from the detector to the back end occur through the POHs, however, a digital optical hybrid (DOH) is used to receive information on the detector from the back end. Power is distributed through cables lining the cylinder as well. The power supplies are situated on the balcony within the experimental cavern and are connected to the detector via $\sim 50\,\mathrm{m}$ long cables. The same powering infrastructure from the original detector is used for the upgrade, however with the increased number of chips to power and operate, the amount of supply current would effectively double if powering directly, leading to large ohmic losses in the cable. To avoid this, custom DC-DC converters, controlled using a FEAST2 ASIC, are used

to perform the necessary power conversion to provide adequate power to each of the modules without incurring significant energy loss in the cables. In total, 1216 DC-DC converters are installed within the detector in the service cylinders around a meter away from the detector modules, with each converter servicing anywhere from 1 to 4 modules, providing the low voltage (LV) necessary to power the chips. The high voltage (HV) bias powering, which provides the silicon depletion in the sensors, remains unchanged from the previous pixel detector. Lastly, the cooling pipes run through the service cylinders as well to regulate the thermal load generated from operating the detector, providing a CO_2 coolant temperature of around -22 °C.

Back End

Any communication with the on-board detector electronics are all performed through the off-detector DAQ and control infrastructure [9], which are located within the underground counting room. The phase-1 upgrade adopted the microTCA-based system using a combination of custom and commercial field-programmable gate array (FPGA) mezzanine cards (FMC). These FMCs can use different hardware and software to perform different functions, among which are the: Front-End Driver (FED), which receive and decode the pixel data; the Pixel Front-End Controller (FEC), which is used to send the module configuration parameters as well as distribute the clock and trigger signals and other fast commands; and the Tracker FEC, which is used to program and control the auxiliary electronics in the service cylinders. In total, 108 FEDs are used (28 for FPix and 80 for BPix), 16 pixel FECs (8 FPix and 8 BPix), and 3 tracker FECs (1 FPix and 2 BPix). These FMCs are housed within twelve microTCA crates. Each crate is also furnished with a custom CMS component, the AMC13, which receives the trigger and clock signals from the central DAQ and distributes them to the cards contained within

the crate via the connection to the backplane of the crate. Figure 3.16 shows a schematic of the whole pixel DAQ chain. Communication among components in the service cylinder or the detector occur via electrical signal transmission, while any communication from the front end to the back end is performed via optical transmission.

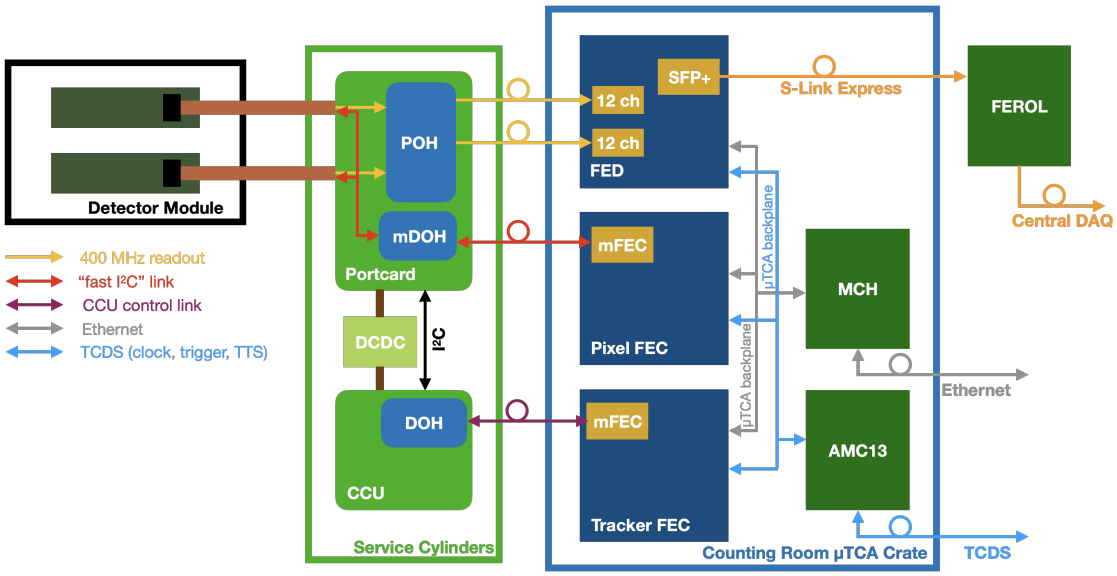

Figure 3.16: Schematic of the pixel data acquisition system, segmented into the back end and front end, the latter of which consists of the service cylinders and detector.

The back end DAQ is controlled via the pixel online software (POS), which not only oversees the data acquisition from the detector, but also contains the suite of calibration tools to determine optimal running configuration parameters for the detector components. POS is built on the CMS online software platform XDAQ, which provides a framework along with a set of tools and interfaces to develop custom data acquisition applications. A XDAQ application is used to connect to the FMCs within the microTCA crates using the IPBus protocol via the ethernet connection established by the backplane connection of the board. This allows for direct interaction through a web interface created when the application is instantiated.

The CMS DAQ system uses a finite state machine (FSM) computing model to define the behavior of the global detector system with a discrete, well-defined set of states

along with explicitly specified set of allowable transitions between the states. The pixel DAQ system has also adopted this FSM approach where the hardware state of the FMC is reflected in the software state of its corresponding XDAQ application, known as a *supervisor*. Each individual FMC is controlled by its own supervisor, with cross-communication between them provided by the simple object access protocol, or SOAP messages. During data-taking, the central CMS DAQ system controls a level-0 function manager (L0FM), which, in turn, drives a level-1 function manager (L1FM) for each of the different subdetectors. The L1FM for the pixel detector communicates directly to a top-level supervisor, which coordinates the activities of the other supervisors, particularly during the configuration stage or while running a calibration on the detector. Calibrations and other diagnostic tests are performed when the subdetector is running *locally* and is detached from the central DAQ during periods where no beam is circulating inside the LHC.

Commissioning and Operations

In order to better ensure high track data quality and stable data-taking, a host of calibrations are routinely performed to keep the pixel detector in an optimal condition. These calibrations are typically performed in a sequence, beginning first with calibrations of the auxiliary electronics, scanning over laser driver outputs in the opto-hybrids to establish both incoming and outgoing communication with the on-detector hardware. Next, a set of calibrations scan over timing and delay parameters to find the optimal setting for a reliable communication line between the DAQ and the detector with high data integrity in the expected data stream structure. Once a clear line of communication is established, the ROCs themselves are also programmable, with digital-to-analog converter (DAC) registers built into the ROC that can control different chip properties. These DAC set-

tings provide baseline ROC-wide parameters, which include performance-specific parameters such as the comparator threshold a charge collection pulse height must satisfy to be stored in a buffer, or settings that amplify or sculpt the shape of the charge pulse, as well as other calibration-specific parameters such as the magnitude of the per-pixel calibration charge injection mechanism and a parameter specifying a delay for the charge injection. Additional bits on the ROC provide an even higher granularity of control over individual pixels with fine adjustments of charge thresholds for single pixels or masking of individual noisy pixels that may have been damaged. These thresholds must be tuned and minimized to allow for a sensitivity to a wide range of input signals, while also sufficiently rejecting noise. Once the full readout chain is established, a *PixelAlive* test can be run, which tests each individual pixel with multiple calibrated charge injections above the pixel threshold and evaluates the efficiency in which the DAQ system is able to correctly ascertain the injected hit. An example output of a PixelAlive test is shown in Fig. 3.17, arranged in the shape of a module with sixteen ROCs. With a fully calibrated system, the pixel detector can be integrated with the rest of the CMS subdetectors to record collision data as a single cohesive unit.

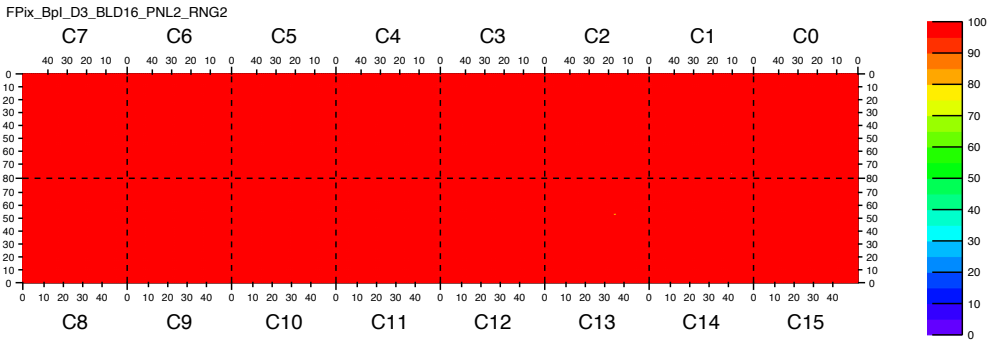

Figure 3.17: Result of a PixelAlive diagnostic, which shows the efficiency of individual pixel responses to a calibrated charge injection. This is a helpful test to diagnose issues with individual pixels, ROCs, modules, or entire readout groups when running the test on the entire detector to establish the location of a possible failure mode.

While each of the different components of the pixel detector received multiple exten-

sive testing campaigns to ensure the stability of operation within the extreme radiation environment produced by the LHC, there were still various surprises encountered during the early and mid stages of operation of the new detector, the most significant of which will be recounted here.

Single event upset (SEU) interactions are those in which a highly ionizing particle is able to change the state of a transistor in any of the microelectronics used within the detector. While SEUs can be relatively harmless and go unnoticed, other occurrences may present a much more adverse effect when sensitive circuitry in various parts of the detector are affected and corrupt the data or disrupt the readout capabilities of the system. Given the highly radiative environment near the interaction region, SEUs are frequent and expected occurrences such that software protocols are in place to detect these events and perform appropriate recovery measures by reprogramming or resetting the affected chips to return to standard operation. However, within weeks of initial collision data-taking, a design flaw in the TBM chip on the pixel modules was discovered in which an SEU-affected flip-flop bit could indefinitely halt the readout chain for a set of ROCs, but could not be reset by issuing any electrical signal. Instead, the readout could only be restored upon power cycling the chip. This meant that throughout the duration of an LHC fill, without any intervention, the pixel detector would slowly accumulate inactive regions that would no longer send any data. Perhaps seemingly serendipitously, the phase-1 upgrade of the pixel detector also renovated the powering system to include DC-DC converters within the service cylinder, which could be enabled and disabled remotely to cut and restore power to the modules connected to the DC-DC converter with finer granularity than is offered by the power supplies. This provided a mechanism in which to automate the restoration of the readout by integrating a targeted power cycling of modules after standard reset actions have failed during the SEU recovery protocol.

While this strategy worked smoothly for several months, on October 5, 2017, a few DC-DC converters had suddenly failed, no longer delivering power to the modules. The automated power cycling within the SEU recovery mechanism was halted immediately, and manual power cycles were performed between LHC fills, but more DC-DC converters were failing with each power cycle, reaching a final total of 5% of DC-DCs lost, 29 in FPix and 38 in BPix, by the end of the 2017 data-taking campaign. As the issue was unresolved, the pixel detector was extracted during the year-end technical stop to examine the failed DC-DC converters. An extensive investigation was launched and various irradiation tests were commissioned to understand the root cause of failure, which led to reproducible failures in a controlled lab setting consistent with those observed in the detector. But unfortunately, the failure mechanism of the DC-DC converter could not be determined before the end of the technical stop and thus, the detector was refurbished with new but otherwise identical DC-DC converters and reinstalled in CMS. However, before the LHC began circulating and colliding at full intensity, the problem was determined to be from the accumulation of radiation-induced current on a transistor, which normally flows to ground when the DC-DC is enabled, but will accumulate on a capacitor outside of its nominal tolerance when the DC-DC is disabled, thus causing permanent damage and rendering the DC-DC inoperable. The power cycling scheme to recover SEU-affected TBMs was modified accordingly, performing the power cycle with the power supplies directly, while leaving the DC-DCs enabled. This led to a successful 2018 run with no DC-DC losses. PixelAlive tests on the pixel detector revealed damage accrued in modules affected by the DC-DC failure due to application of the HV bias without the LV, which led to leakage currents flowing through unintended parts of the chip causing permanent damage. An example of the PixelAlive damage is shown in Fig. 3.18. Several damaged modules in the first BPix layer behind broken DC-DCs were replaced during the refurbishment between 2017 and 2018.

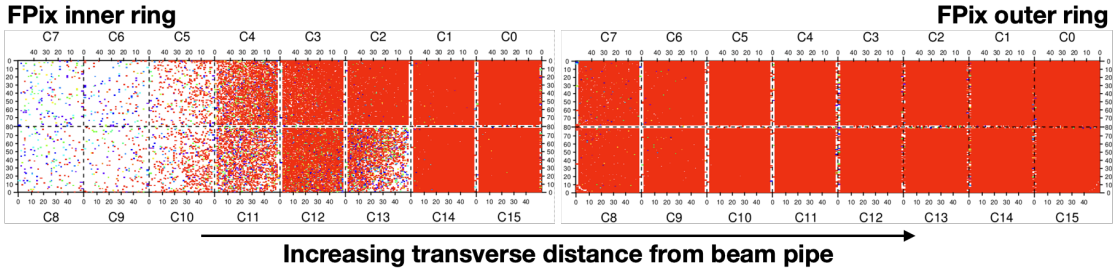

Figure 3.18: PixelAlive of damaged modules in FPix from operating with HV bias voltage on without LV due to DC-DC failure.

Detector Performance

Despite all of the difficulties encountered during data-taking, the pixel detector still maintained a peak active working fraction of 93.5% in BPix and 96.7% in FPix by the end of 2018 with periodic power cycles in between fills. Thanks to the tireless efforts of many of the operators and experts to maintain stable operation, the detector managed to deliver high quality data with both high hit efficiency and position resolution. The hit efficiencies across 2016 (with the previous detector), 2017, and 2018 are shown in Fig. 3.19, which demonstrate the crucial need for the upgrade to begin with, and the high performance that was attained with the upgraded detector. The position resolution is shown in Fig. 3.20, measured as the pixel hit residual in the r-ϕ and z directions in the third barrel layer. For nonisolated particles with transverse momentum p_{T} in the range $1 < p_{\mathrm{T}} < 10\,\mathrm{GeV}$ and $|\eta| < 1.4$, the track resolutions are typically 1.5% in p_{T} and 20–75 μm in the transverse impact parameter [6]. With the upgraded silicon pixel tracker used in this dataset, the track impact parameter resolution has improved by 25%.

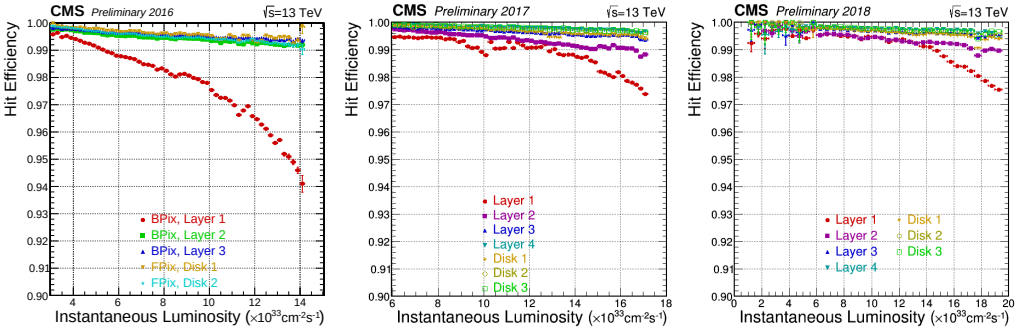

Figure 3.19: Cluster hit efficiencies as a function of instantaneous luminosity as measured in 2016 (left), 2017 (middle), and 2018 (right). Mind the x-axis ranges.

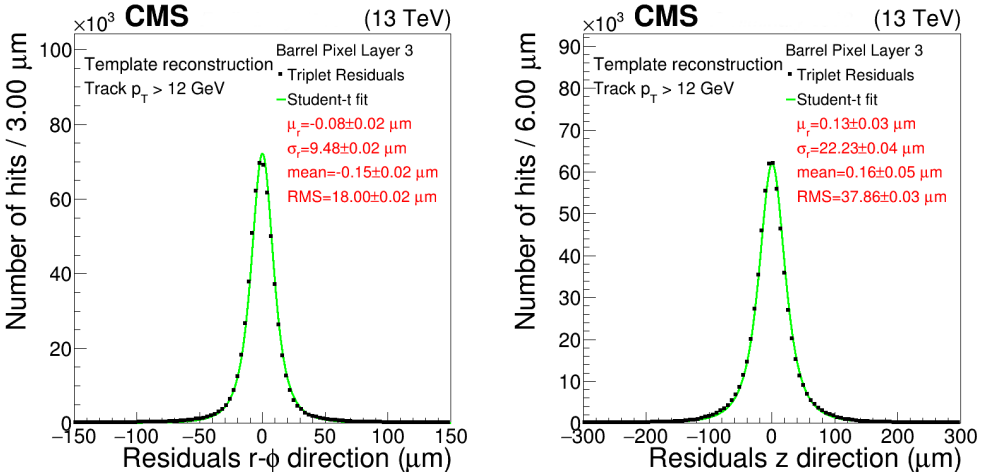

Figure 3.20: The fitted hit residuals in the third BPix layer in r-ϕ (left) and z (right). Residuals for the other layers and disks looks similar.

3.3 CMS Trigger

The LHC collides two proton bunches nearly once every 25 ns resulting in a collision rate of 40 MHz. A typical event size, corresponding to raw sensor information throughout the detector, amounts to around 1 MB of information. At the nominal collision rate, this easily exceeds 40 terabytes of data in a single second of data taking, or more than 10 petabytes of data in a typical LHC run length of several hours. This is unsustainable, and additionally, undesirable as very few collisions will actually contain a hard enough collision to produce an interesting interaction. Thus, a trigger system is used to filter out

the handful of interesting events for further analysis, while discarding the overwhelming majority to stay within the detector readout bandwidth and storage limitations.

The CMS trigger is a 2-tiered system, Level-1 (L1) and the High-Level Trigger (HLT). The L1 is the first level, which must make extremely quick decisions to cut down the 40 MHz rate. This level is a coarse filter based on simple criteria, accomplished using custom hardware processors that use information only from the calorimeters and muon chambers to quickly reconstruct crude event objects using lookup tables containing stored hit patterns. The system makes a decision whether to accept or reject the event on whether these objects pass a set of predefined criteria within $\sim 4\,\mu$s of the collision event. If the event is accepted, an L1A is issued to the entire detector system, which will then read out the associated data for that particular bunch crossing from the storage buffers on the readout electronics. As a result of this initial filtering, the L1 trigger reduces the event rate to $\sim$100 kHz and the detector information is sent to the HLT.

At the second level, the HLT is a software-based trigger employing a computing farm with over ten thousand processors to perform a more thorough reconstruction, but still less sophisticated than the full offline treatment with a time budget of a few hundred milliseconds to make a decision. The tracker is included at the HLT reconstruction level in addition to the muon and calorimeter systems. The HLT ultimately reduces event rate to $\sim$1 kHz where the accepted events are saved to tape and transferred for the full offline event reconstruction.

CHAPTER 4

EVENT RECONSTRUCTION

The previous chapter described the literal machinery used to collide beams of protons together and the detectors that aim to capture as much information in as much detail of the resulting collisions at every bunch crossing, referred to as an *event*. This chapter will be a discussion in taking the raw information of charge depositions littered throughout the detectors and transforming them into useful physics objects with which to analyze. This step is called the event reconstruction. The focus will be on the reconstruction of objects of particular interest to this analysis.

4.1 Charged particle tracks

The reconstruction of charged particle tracks [35] is the most difficult and computationally intensive procedure in the full event reconstruction chain. In every bunch crossing event, thousands of charged particles emerge from the hard scattering interaction and pass through the pixel detector and the strip tracker. The nearly uniform magnetic field produced by the solenoid guides the charged particles into helical trajectories. This presents the challenging problem of reconstructing, with high efficiency, the complicated path of the particles while suppressing fake tracks, which arise from misappropriating unrelated signals in the tracking layers with the track. The entire track reconstruction occurs in a multi-stage sequence, which will be described here.

In the wake of the spray of charged particles, they leave behind measurable traces of their passage in the form of freed electrons in the depleted silicon sensors.A *hit* in the pixel detector can be reconstructed from clusters of adjacent pixel sensors in which an electrical signal from these free electrons exceed the programmed threshold of the

readout electronics. The hit position and uncertainty is determined in a local two-dimensional coordinate system within the plane of the silicon sensor from a charge-weighted average of the pixel position. These local coordinates are converted into the global coordinate system with additional input from alignment calibrations and the geometry of the detector. A similar process using the strip detector provides a the position estimate of hits across the outer layers of the tracker. The trajectory or *track* of the charged particle can be reconstructed from these estimated hit locations throughout the tracking layers using an iterative track finding procedure. Each iteration consists of four main steps: (i) seed generation, (ii) track finding, (iii) track fitting, and (iv) track selection.

The seed generation step forms the initial track candidates from a set of hits across 3 layers alone or 2 layers with a beam spot constraint. This provides a preliminary estimate for the five track parameters and uncertainties required to describe a helical trajectory: the signed transverse impact parameter or x-y distance of closest approach d_0, the longitudinal impact parameter z_0, the azimuthal angle in the x-y plane ϕ, $cot\theta$, where θ is the polar angle, and the transverse momentum p_{T}.

Using the initial track parameter estimates from the seeds, tracks are extrapolated into subsequent layers of the detector. The Kalman filter method is used to build the track up layer by layer, searching for hits compatible with the extrapolated track in each layer, updating the track parameters as the hit is added to the track.

Once the algorithm has finished propagating the track to the outer layers and collecting hits compatible with the track, a final fit is applied to the track, using the full hit information and without the beam constraint. Any spurious hits are removed from the track if they are incompatible based on a χ^2 fit. The extrapolation of the trajectory takes into account the material scatter and inhomogeneities in the magnetic field to improve

the accuracy of the track reconstruction.

Finally, a set of quality requirements are applied to the track candidates to suppress the fake rate. These requirements are placed on the number of layer hits associated to the track, the χ^2/dof, the transverse and longitudinal impact parameter significances, among others.

This whole procedure is repeated six times. With each iteration, the hits associated with the selected tracks are removed from consideration to simplify the combinatorics in the subsequent iterations.

4.2 Primary vertices and beam spot

In 2017 and 2018, the mean number of simultaneous pp interactions occurring at each bunch crossing, or *pileup*, was approximately 32. Primary vertices correspond to points where the reconstructed charged particle tracks converge from an interaction. The reconstruction of a primary vertex [35] uses the reconstructed tracks, clustered based on the z coordinate at the point of closest approach to the beam line, to measure the location of the interaction vertex. A deterministic annealing algorithm is used to determine the primary vertex multiplicity in an event and how the tracks are assigned among them. The algorithm is an analogue to the actual physical process of annealing by heating a material and cooling it gradually in order to recrystallize the internal microstructure in the material into a refined state [30]. The algorithm performs this by tuning a parameter analogous to 'heating,' which assigns all tracks to a single primary vertex cluster at a high temperature. As the temperature parameter is lowered, the track set will 'recrystallize' as the system undergoes a phase transition and the tracks coalesce into the optimal configuration of independent clusters. The resulting track clusters are fitted using an

adaptive vertex fitter to obtain the position and other parameters of the primary vertices.

The beam spot represents the region in which the proton bunches are made to collide. The parameters of the spot are determined from an average across many events occurring within a specified time frame called a *luminosity section* (LS), which lasts 23 seconds. The position and shape of the beam spot is measured using the three-dimensional distribution of the primary vertices in a LS.

4.3 Particle-flow reconstruction

While many of the principal components of an event have been described, the full global event reconstruction aims to advance this further by reconstructing and identifying individual particles emerging from the interaction. This is done using the particle-flow (PF) algorithm [59], which cleverly combines pieces of information from each subdetector of CMS to classify the particles based on their associated interaction with the detector components.

Figure 4.1 shows the interactions of the detector with broad classes of particles described by the SM. The PF algorithm works by linking together individual PF elements, which are reconstructed particle tracks in the tracker, energy clusters in the calorimeters, or hits in the muon chambers. Tracks can be linked to energy clusters if the track extrapolation coincides with a measured cluster in either of the calorimeters. A *block* of linked elements can then be assigned a particle type based on which detector elements provided the PF element. The identification of particle type begins first with classifying blocks compatible with muons.

A muon type is assigned to a block consisting of hits in the muon system, but they

come in three varieties: standalone muons, which come from the reconstruction of tracks in the muon chambers; global muons match the standalone muon tracks to tracks reconstructed in the inner silicon tracker; and tracker muons, which are reconstructed beginning with a track reconstructed in the inner tracker and extrapolating it into the muon chambers. Tracker muons and global muons are complementary as the tracker muon reconstruction has a better efficiency for low-p_{T} muons that may be unable to penetrate enough layers in the muon system to reconstruct a standalone muon. The muon does not get stopped by either of the calorimeters, thus the energy of the muon is inferred from the curvature of its trajectory.

The next particles reconstructed are electrons, which are blocks identified by a track linked with an associated cluster in the ECAL. The remaining energy clusters in the ECAL that don't have a corresponding track linked from the tracker are identified as non-isolated photons. Similarly, charged hadrons are identified from HCAL and ECAL clusters linked to track elements, while neutral hadrons rely only on the HCAL and ECAL energy deposits.

4.4 Jets

The overwhelming majority of particles that emerge from the beam collision are composed of quarks and gluons. Due to color confinement, these are not observable independently, but instead hadronize into collimated showers of particles, collectively referred to as a *jet*. From data collected from the detector, jets are reconstructed from clusters of PF candidates using the anti-k_{T} algorithm [28] with a distance parameter R of 0.4. The momentum of the jet is calculated from the vector sum of momenta of the particle constituents in the jet. For the search described here, jets are used to trigger and select

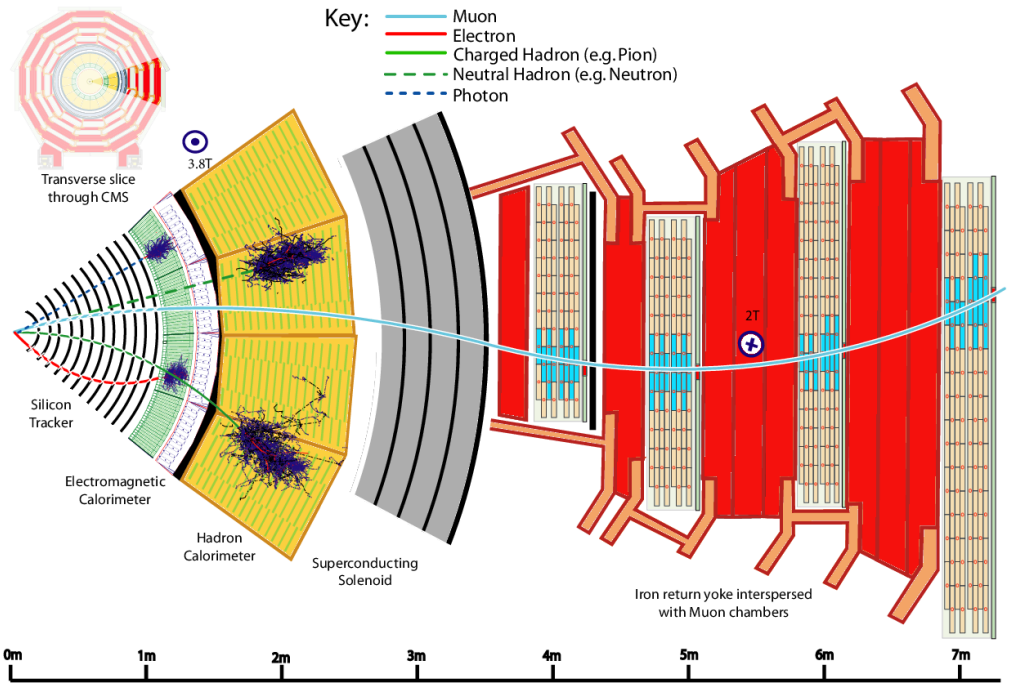

Figure 4.1: Visualization of different particle interactions in a transverse cross section of the CMS detector.

the initial events for further analysis.

As with any object reconstruction, the precision of the jet energy measurement is not always perfectly measured due to numerous challenges including the non-linear detector response to the particles, additional tracks from pileup, noise, or inefficiencies in the jet reconstruction algorithm itself. To account for these, various corrections and calibrations must be applied to recover the true particle or parton energy. First, any jet energy contributions from particles originating from other pileup events are discarded. Additional corrections are applied to restore a flat response as a function of the jet p_{T} and η. A final residual correction applied to the data corrects for the small remaining differences after the previous corrections between data and simulation. Various selection requirements can be applied to the jets to further refine the quality of the jet and reduce noise, whether the noise be electronic or misidentified leptons.

4.5 Heavy flavor tagging

Various studies throughout this search rely on identifying (*tagging*) jets that contain the decay of a b hadron, i.e. b jets. The b quark is unique for its relatively long lifetime of about 1.5 ps leading to displacements of a few mm or a cm, which can affect searches for other exotic long-lived particles and introduce correlations if not appropriately identified and handled. The effects of long-lived b quarks on this analysis will be discussed in future sections, whereas a description of the tagging algorithm will be provided here.

This analysis employs the DeepJet tagging algorithm [60], which is a deep neural network algorithm that exploits the properties of heavy-flavor jets, including impact parameter significances for jet constituents, reconstructed vertex mass, number of tracks in the vertex or jet, among many others. In total, the properties are encapsulated into 16 features for charged particles or 8 features for neutral particles for up to 25 jet constituents and 12 features for up to 4 secondary vertices to use as inputs into the network. These features are passed through a set of convolutional neural network layers to form a compressed representation of the features. This feeds into three recurrent layers or long short-term memories, which outputs to one dense fully-connected layer. The final output is in the form of a numerical *b discriminator* value, which represents the probability that the jet contains a b hadron. These discriminator values can be interpreted at various levels of *working points*, e.g. *loose*, *medium*, or *tight*, where the choice of working point is dictated by the level of efficiency required by the analysis and the level of tolerance for misidentification.

CHAPTER 5

SEARCHING FOR DISPLACED VERTEX PAIRS IN MULTIJET EVENTS

With the necessary framework established, this chapter will describe the details and methodology of the search for long-lived particles using the CMS detector by reconstructing displaced vertices representing the positions of the particle decay. The narrative will begin with a description of the datasets used, followed by a discussion of the basic trigger and event preselection requirements. Within the events satisfying these requirements, the techniques and algorithms used to reconstruct the displaced vertices will be described along with the overall strategy in which to discriminate signal events from background events. The background estimation strategy will be described along with the systematic uncertainties in this estimate and the signal yield estimate. These pieces will be integrated into the statistical interpretation of the unblinded data to extract the signal yield.

5.1 Datasets and samples

5.1.1 Data samples

For this search, we used data samples collected in 2017 and 2018 during proton-proton collisions at a center-of-mass energy of $\sqrt{s} = 13\,\text{TeV}$ at CMS, corresponding to a total integrated luminosity of $101\,\text{fb}^{-1}$. Figure 5.1 shows the integrated luminosity delivered by the LHC and how much was recorded by CMS [4] [5]. Discrepancies between these quantities can appear due to various inefficiencies in data taking. These can include *down time*, where the detector is down and not acquiring data for various reasons including when a subsystem undergoes hardware or software difficulties, or *dead time* where

the detector is actively taking data, but may be reaching the limitations of the data link bandwidth or buffer capacities due to transient issues such as exceedingly high instantaneous trigger rates overwhelming the data acquisition system. Lastly, although collision data may have been recorded by CMS, dedicated validation studies on the dataset are performed to ensure the quality of the data meets a reference standard in which the data can be reliably used to study physics. Such validation studies can reject small samples of data based on decisions from the different subdetector teams (e.g. pixels, HCAL) or offline reconstruction teams (e.g. tracking, jets) due to any potential problem within the detector that may have arisen during data taking. The final dataset upon completion of this review is then officially certified, and in 2017 and 2018, this corresponded to an integrated luminosity of 41.5 fb^{-1} and 59.7 fb^{-1}, respectively.

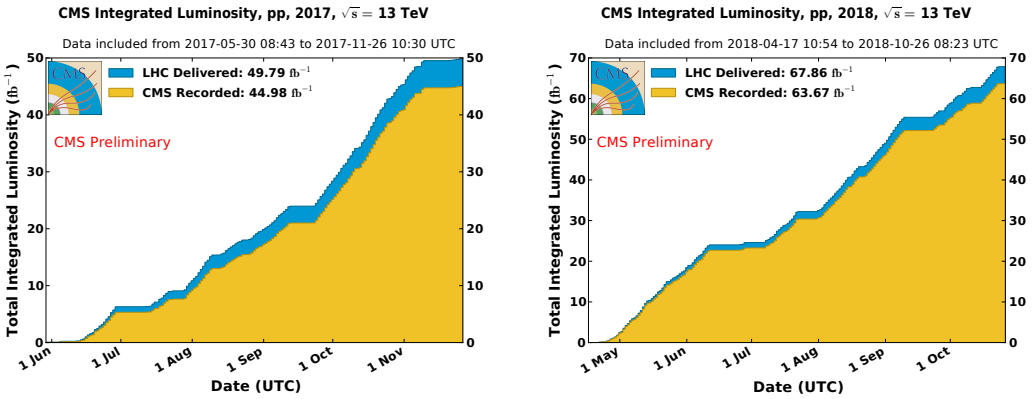

Figure 5.1: Integrated luminosity recorded per day in 2017 and 2018.

Table 5.1 lists the datasets used for each year of data taking, which are divided based on the HLT trigger that they pass. For this analysis, the events were selected using a jet H_{T} trigger, which will be discussed in Section 5.2.1. Each of the years are split into separate *eras*, which explicitly delineate periods of notable adjustments in the machine running conditions that warrant partitioning of the data.

Table 5.1: Datasets used in this analysis.

Era	Dataset	Run range	Int. lumi. (fb^{-1})
2017B	/JetHT/Run2017B-17Nov2017-v1	297047–299329	4.80
2017C	/JetHT/Run2017C-17Nov2017-v1	299368–302029	9.63
2017D	/JetHT/Run2017D-17Nov2017-v1	302031–302663	4.25
2017E	/JetHT/Run2017E-17Nov2017-v1	303824–304797	9.32
2017F	/JetHT/Run2017F-17Nov2017-v1	305040–306460	13.5
2018A	/JetHT/Run2018A-17Sep2018-v1	315257–316995	14.0
2018B	/JetHT/Run2018B-17Sep2018-v1	317080–319310	7.09
2018C	/JetHT/Run2018C-17Sep2018-v1	319337–320065	6.94
2018D	/JetHT/Run2018D-PromptReco-v2	320497–325175	31.9

5.1.2 Simulation samples

Monte Carlo (MC) simulations are used to model the RPV SUSY signal and relevant background processes. These simulations allow us to study the nature and composition of the signal and background events to construct the analysis and improve its sensitivity and also provide a way to validate the data-driven background estimation procedure.

Signal

Figure 5.2 shows a pair of diagrams that represent the signal models used to measure the efficiency and performance of this analysis, providing a benchmark assessment of the search methodology for a broader class of signal models with similar final states. The first model corresponds to the pair-production of long-lived neutralinos or gluinos, the lightest supersymmetric particles (LSP) of this model, where each LSP undergoes a three-body decay into top, bottom, and strange (anti-)quarks. This final state in which each long-lived particle decays into three or more jets will be referred to as the ”multi-

jet" final state throughout this search. To study models with lower track multiplicities, another RPV SUSY model in which the corresponding LSP is the top squark is considered, where each of the pair-produced top squarks decays into a "dijet" final state of two down anti-quarks. For various reasons that will be discussed in the full description of the analysis technique and strategy, this search is not sensitive to signal models decaying into a single jet, and are not used as a benchmark indicator. Signal events were simulated using PYTHIA 8.230 [67] with NNPDF3.1LO [20] providing the parton distribution functions. Simulations using the standard CP5 tune for the NNLO PDF led to non-neglible yields of events with negative weights, resulting in the need for higher MC generation in order to achieve the same effective statistics. To avoid this, the CP2 tune [63] is used instead to model the underlying event. The samples were produced centrally with 100,000 events each for a range of masses m ($400 \leq m \leq 3000\,\text{GeV}$) and mean proper decay lengths $c\tau$ ($0.1 \leq c\tau \leq 100\,\text{mm}$). These ranges correspond to the region of phase space to which this search is most sensitive.

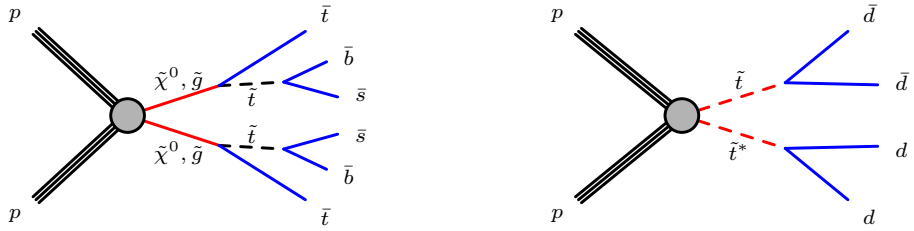

Figure 5.2: Diagrams of the signal processes for multijet(left) and dijet (right) events.

Background

Background events arising from SM processes come entirely from multijet (QCD) events and events with pair-produced top quarks ($t\bar{t}$) that contain enough jet activity to pass the H_T trigger. These background samples are simulated using MADGRAPH5_AMC@NLO 2.4.2 [12] with the NNPDF3.0 [19] parton distribution function set at leading order with

MLM merging [13]. PYTHIA 8.230 [67] was used to simulate the hadronization and showering with the CP5 tune [63] providing the underlying event model for the background simulation samples. Tables 5.2 and 5.3 provide information on the background samples used in this analysis.

Table 5.2: Statistics of 2017 background MC samples.

Process	Cross section (pb)	Number of events	Weight for 41.5 fb^{-1}
QCD, $700 < H_T < 1000$ GeV	6.4×10^3	48042655	5.49
QCD, $1000 < H_T < 1500$ GeV	1.1×10^3	16882838	2.70
QCD, $1500 < H_T < 2000$ GeV	99	11634434	0.353
QCD, $H_T > 2000$ GeV	20	5941306	0.141
$t\bar{t}$, $600 < H_T < 800$ GeV	1.8	81565576	9.25×10^{-4}
$t\bar{t}$, $800 < H_T < 1200$ GeV	0.75	40248127	7.76×10^{-4}
$t\bar{t}$, $1200 < H_T < 2500$ GeV	0.13	13214871	4.13×10^{-4}
$t\bar{t}$, $H_T > 2500$ GeV	1.41×10^{-3}	5155687	1.14×10^{-5}

Table 5.3: Statistics of 2018 background MC samples.

Process	Cross section (pb)	Number of events	Weight for 59.7 fb^{-1}
QCD, $700 < H_T < 1000$ GeV	6.4×10^3	43523821	8.74
QCD, $1000 < H_T < 1500$ GeV	1.1×10^3	15174716	4.33
QCD, $1500 < H_T < 2000$ GeV	99	11082955	0.535
QCD, $H_T > 2000$ GeV	20	5557453	0.218
$t\bar{t}$, $600 < H_T < 800$ GeV	1.8	14363689	7.6×10^{-3}
$t\bar{t}$, $800 < H_T < 1200$ GeV	0.75	10462756	4.3×10^{-3}
$t\bar{t}$, $1200 < H_T < 2500$ GeV	0.13	2897601	2.7×10^{-3}
$t\bar{t}$, $H_T > 2500$ GeV	1.41×10^{-3}	1451104	5.8×10^{-5}

Both the background and signal samples use a GEANT4-based [10] simulation of the CMS detector response to the simulated event.

5.2 Trigger and Preselection

5.2.1 Trigger Selection

Given the final state of signal models targeted by this search, each event is expected to contain at least four jets from the decay of the LSPs. Events in both data and simulation are selected using the HLT_PFHT series of triggers at HLT level, which uses chooses events based on the H_T of the event, where H_T is the scalar sum of the trigger-level jet p_T. In particular, the HLT_PFHT1050 is used in both 2017 and 2018, which requires $H_T > 1050\,\text{GeV}$ in each event for jets with $p_T > 40\,\text{GeV}$ and $|\eta| < 2.5$. The measured signal efficiency of this trigger is shown in Fig. 5.3 as a function of mass and lifetime for both final state topologies. For signal masses greater than 800 GeV, the overall trigger efficiency is greater than 95%, but for masses of 800 GeV or lower, the efficiency begins to diminish due to the H_T requirement.

A data-driven method is used to measure the efficiency of the trigger as function of event H_T, starting from muon-triggered events containing an isolated muon with $p_T > 27\,\text{GeV}$, namely the HLT_IsoMu27 trigger, and four or more jets reconstructed offline. Additional muon-enriched simulated datasets with the largest physics process cross sections are used to derive an estimate of the systematic differences in trigger efficiency between data and simulation for future systematic uncertainty estimates. The simulated datasets are listed in Table 5.4.

The denominator of the efficiency calculation is determined by the number of events that fire the HLT_IsoMu27 trigger and contain an offline reconstructed muon that satisfies the following requirements:

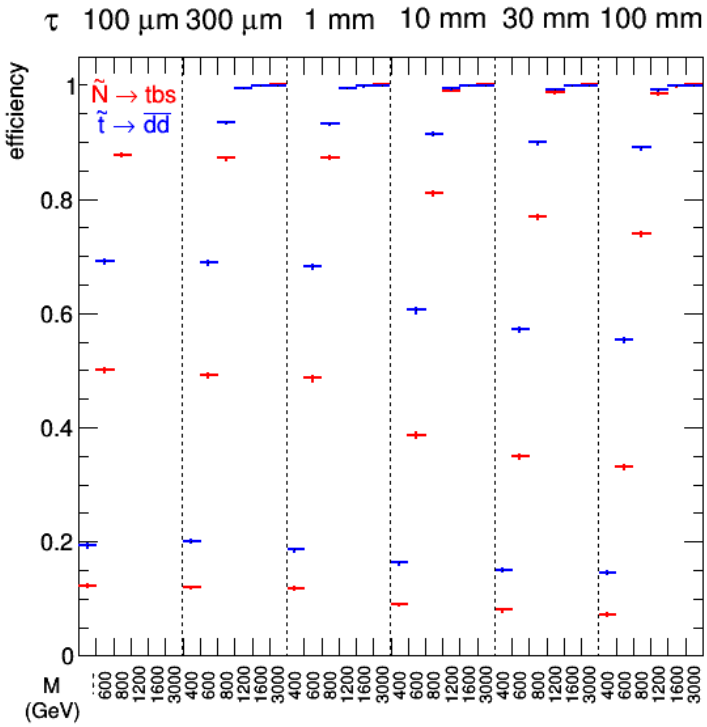

Figure 5.3: The efficiency of multijet (red) and dijet (blue) signal events to pass the `HLT_PFHT1050` trigger, as a function of mass and lifetime.

Table 5.4: Datasets used to study trigger efficiency in simulation.

Process	Dataset name
W+jets, $W \to \ell\nu$	/WJetsToLNu_TuneCP5_13TeV-madgraphMLM-pythia8
Z/γ^*+jets, $Z/\gamma^* \to \ell\ell$, gen. $10 < M_{\ell\ell} < 50$ GeV	/DYJetsToLL_M-10to50_TuneCP5_13TeV-madgraphMLM-pythia8
Z/γ^*+jets, $Z/\gamma^* \to \ell\ell$, gen. $M_{\ell\ell} > 50$ GeV	/DYJetsToLL_M-50_TuneCP5_13TeV-madgraphMLM-pythia8
$t\bar{t}$	/TTJets_TuneCP5_13TeV-amcatnloFXFX-pythia8
QCD, gen. $\hat{p_T} > 20$ GeV & μ $p_T > 15$ GeV	/QCD_Pt-20toInf_MuEnrichedPt15_TuneCP5_13TeV_pythia8

- $p_T > 27\,\text{GeV}$
- $|\eta| < 2.4$
- Global muon and PF muon
- Normalized global track $\chi^2 < 10$
- Number of inner tracker layers with hits > 5
- Number of pixel hits > 0
- Number of muon hits > 0 and number of matched muon stations > 1

The fraction of events satisfying the `HLT_PFHT1050` trigger is plotted as a function of the offline reconstructed jet H_T. The resulting curve is fitted with a function of the form:

$$A + 0.5B\left(1 + \text{erf}\,\frac{H_T - \mu_{on}}{\sigma_{on}}\right) \tag{5.1}$$

The weighted average of the efficiencies in the highest H_T bins (the plateau efficiency) is $A + B$, μ_{on} and σ_{on} correspond to the location and width of the turn-on, respectively. The turn-on curves for both 2017 and 2018 and the overall weighted sum of the simulated samples are shown in Fig. 5.4. The triggers are fully efficiency at offline $H_T > 1200\,\text{GeV}$, with a total efficiency of 0.984 ± 0.001 in the combined 2017 and 2018 data and 0.994 ± 0.002 in simulation, which gives a data to simulation ratio of 0.989 ± 0.002

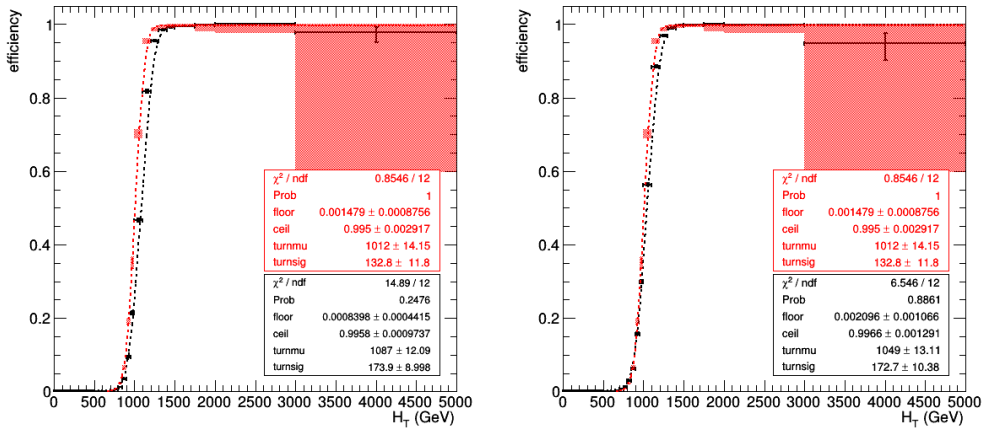

Figure 5.4: `HLT_PFHT1050` efficiency as a function of offline jet H_T in muon-triggered events in data (black) and simulation (red shaded bands), for 2017 (left) and 2018 (right).

5.2.2 Event preselection

To ensure the trigger behavior is well understood throughout the rest of the analysis, we impose an offline requirement that $H_T > 1200\,\text{GeV}$, where H_T is the scalar sum of offline reconstructed p_T of jets in the event with $p_T > 40\,\text{GeV}$ and $|\eta| < 2.5$ to match the trigger definition. Furthermore, the event must contain at least four jets each with $p_T > 20\,\text{GeV}$ and $|\eta| < 2.5$. This substantially reduces the unnecessary processing of background events, while negligibly affecting the efficiency for benchmark signal models in which the final state includes at least four quarks. Figure 5.5 shows distributions of these two variables used for event preselection. A total of 35.8 million events satisfy these requirements in data.

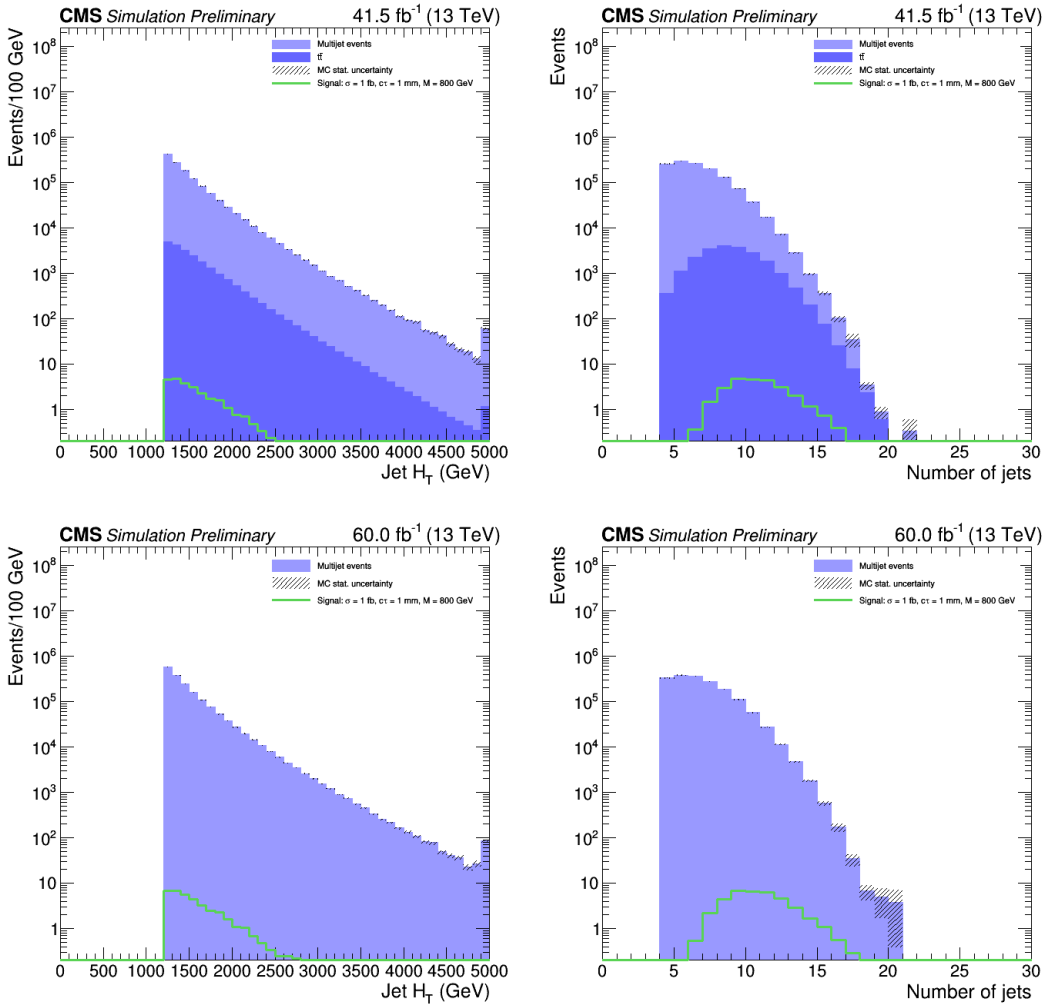

Figure 5.5: Distributions of variables used for event preselection in simulated background and a simulated multijet signal with LSP $c\tau = 1\,\text{mm}$, $M = 800\,\text{GeV}$, and production cross section 1 fb for 2017 QCD plus $\text{t}\bar{\text{t}}$ MC (top) and 2018 QCD MC (bottom). Event preselection criteria have been applied. H_T, the scalar sum of the p_T of jets in the event with $p_\text{T} > 40\,\text{GeV}$, is required to be at least 1200 GeV (left) and the number of jets in the event is required to be at least four (right).

5.3 Displaced vertex reconstruction

This section describes the displaced vertex reconstruction in this analysis, which is accomplished in several steps beginning first with a selection of well-reconstructed tracks. The tracks reconstructed from the raw detector information in data have a higher assumed hit resolution compared to the tracks reconstructed from simulated events passing through a simulation of the detector. To account for this difference, which can lead to an overestimation of the signal efficiency, the track covariance matrix elements for tracks in simulation are rescaled to reflect the same behavior in data. The selected tracks are then passed through a vertex reconstruction algorithm to merge the tracks into vertices.

5.3.1 Track selection

In each event, the tracks reconstructed from recorded hits of charged particles traversing through the silicon tracker layers as described in Sec. 4.1 form the basic elements from which secondary displaced vertices are reconstructed. During the extended year-end shutdown between 2016 and 2017, the new phase-1 pixel detector upgrade was installed in CMS in place of the original phase-0 pixel detector, as mentioned in Sec. 3.2.6. The upgrade includes an additional barrel and endcap layer allowing for increased tracking and vertexing performance. One outcome of this is a reduction in uncertainty of the measured d_{xy} as in Fig. 5.6, which shows the $\sigma_{d_{\mathrm{xy}}}$ distribution in 2016 and 2017. With the upgraded silicon pixel tracker used in this dataset, the track impact parameter resolution has improved by approximately 25%. This is especially important to this analysis, where accurate track reconstruction is needed to determine a reliable measurement of a displaced vertex.

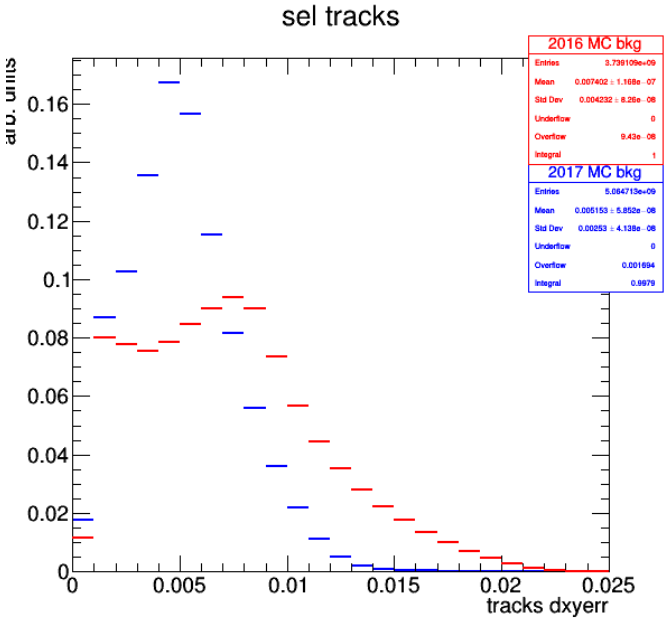

Figure 5.6: Comparison of $\sigma_{d_{xy}}$ between 2016 (red) and 2017 (blue) background MC. Both plots have nominal track quality selection from 2016.

To further suppress background, which arises from vertices that include one or more poorly reconstructed tracks with a large transverse impact parameter d_{xy}, we impose a set of track quality cuts. We optimized these track quality requirements to reduce the transverse track impact parameter uncertainty, $\sigma_{d_{xy}}$, while maintaining a sufficient yield of tracks to form vertices, especially in signal. We require tracks to have p_T of at least 1 GeV; recorded hits in at least two layers of the silicon pixel detector, where one of the layers must be the innermost layer; and recorded hits in at least six layers of the silicon strip detector. Cumulatively, these track requirements reduce the mean uncertainty in the transverse impact parameter of the tracks by approximately 17%. One final requirement is imposed on the significance of the transverse impact parameter, which is the magnitude of the transverse impact parameter divided by its uncertainty, denoted $|d_{xy}|/\sigma_{d_{xy}}$, to be at least 4. This condition favors tracks with larger impact parameters, thereby suppressing the SM background and selecting displaced tracks to seed the reconstruction of displaced vertices. Figure 5.7 shows the distributions for each of these

selection variables in 2017 simulation with the respective selection requirement for each indicated by the vertical red line. The distributions are similar in the 2018 simulation.

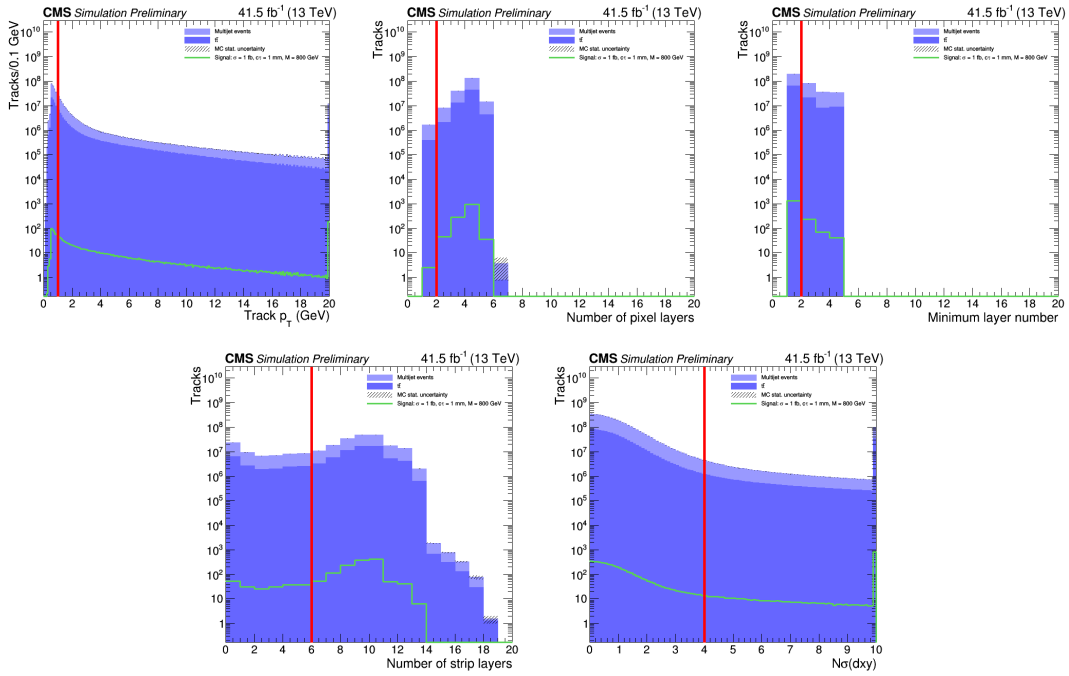

Figure 5.7: Distributions of variables used for track selection in 2017 MC. These are "n-1" plots where all event preselection and track selection criteria have been applied, except the one related to the variable shown. The vertical red line delineates the selection requirement on that particular variable, where in all but the minimum layer number variable, entries to the right of the line are kept. The p_{T} of the track is required to be at least 1 GeV (top left); the number of pixel layers is required to be at least 2 (top middle); the minimum layer number, i.e. the inner most layer with a recorded hit, is required to be 1, note: entries to the left of the vertical red line are kept for this plot (top right); the number of strip layers is required to be at least 6 (bottom left); the transverse impact parameter significance $|d_{\mathrm{xy}}|/\sigma_{d_{\mathrm{xy}}}$ is required to be at least 4 (bottom right).

5.3.2 Systematic differences in track resolution between data and simulation

In our assessment of the systematic uncertainties in signal efficiency from the vertex reconstruction procedure, we observed differences between data and simulation that war-

ranted deeper investigation.

Figure 5.8 shows the difference in the shape of the distributions of the track impact parameter uncertainty between simulation and data for 2017 and 2018. The noticeable discrepancy occurs at low values of uncertainty where the simulation events retain a greater number of tracks with small impact parameter uncertainty compared to the data after applying the standard track quality selection requirements (excluding the track impact parameter significance requirement). The difference occurs in both years, but is slightly worse in 2017 than it is in 2018.

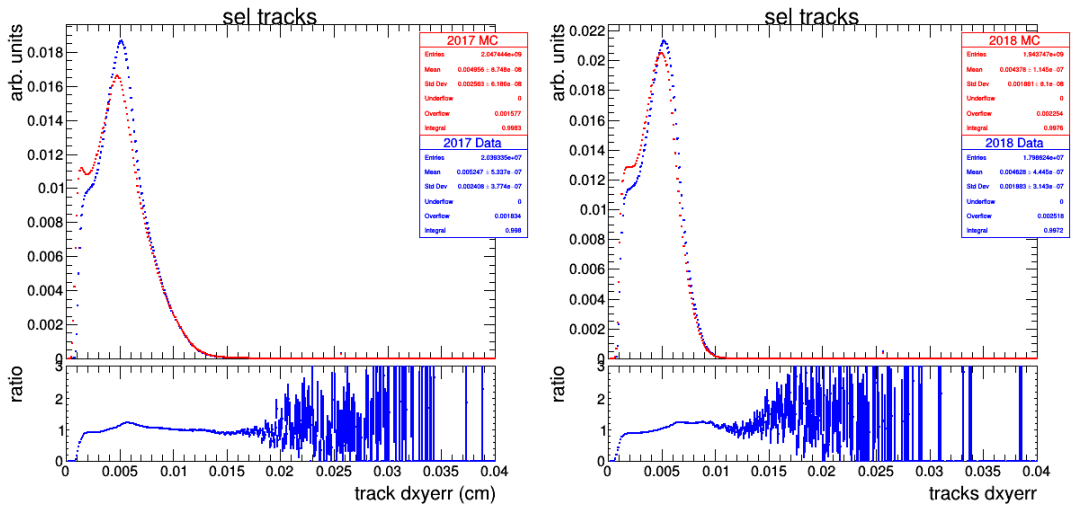

Figure 5.8: Comparison of the track impact parameter uncertainty distribution for 2017 (left) and 2018 (right) between MC simulation (red) and data (blue) tracks satisfying standard track quality requirements.

The discrepancy has many potential explanations, especially in a sample of tracks that integrate across a wide range of many different variables, e.g. p_{T}, η, or the number of hits in different layers across the tracking detector. To disentangle these confounding effects, we apply a highly constraining set of requirements to the tracks to obtain an "ultra clean" sample of tracks. We proceed with our investigation on this controlled sample of tracks to isolate the cause of the data/simulation differences. The ultra clean requirements are as follows:

- $|\eta| < 0.8$;
- $|z_0| < 10\,\text{cm}$;
- Hit measurements in four pixel layers;
- Hit measurements in at least six strip layers.

The η and z_0 requirements restrict our attention to tracks in the barrel to remove effects of differences in impact parameter uncertainties between forward and barrel tracks. The tracking hit requirements similarly maneuver around imperfections in the detector caused by transient inactive channels that affect the uncertainty. With these restrictions applied, Figure 5.9 shows the resulting distributions of the track impact parameter uncertainty within narrow 1 GeV p_{T} slices, showing the persistence of this problem in this restrictive sample of tracks across the p_{T} spectrum.

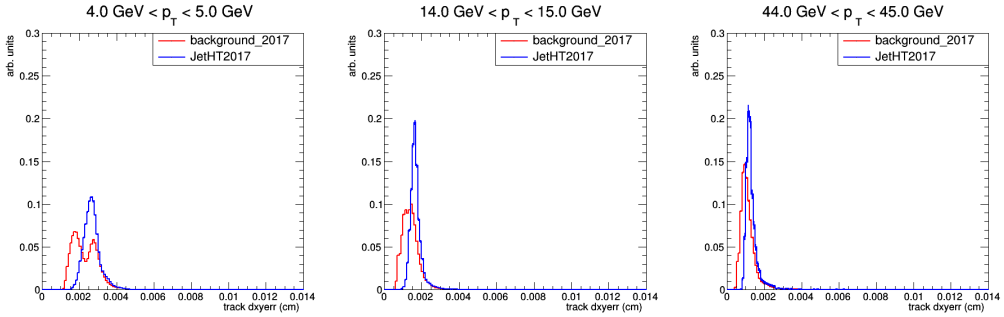

Figure 5.9: Comparison of the track impact parameter uncertainty distribution between 2017 MC simulation (red) and 2017 data (blue) for tracks satisfying the ultra clean track requirements in different narrow p_{T}: tracks with $4 < p_{\text{T}} < 5\,\text{GeV}$ (left), tracks with $14 < p_{\text{T}} < 15\,\text{GeV}$ (middle), and tracks with $44 < p_{\text{T}} < 45\,\text{GeV}$ (right).

The behavior as a function of p_{T} can best be captured with a profiling of the distribution. Figure 5.10 shows the mean impact parameter uncertainty as a function of p_{T}, demonstrating the systematic difference between the background simulation and data across the range of p_{T}. This difference also shifts variably between the different run conditions demarcated by the eras of the 2017 data-taking period. The extreme limits

of this profile have different expected dependencies. At low values of p_{T}, the dominant contributor to the impact parameter uncertainty should come from the multiple scattering of particles with the detector material. This, in turn, worsens the resolution resulting in the large uncertainty observed as p_{T} approaches smaller values. Since this effect scales as $1/p$, its effect on the impact parameter uncertainty diminishes as the p_{T} of the track increases. As a result, for large p_{T} tracks, the principal effect on the uncertainty comes from the assumed hit resolution in the track fit and various geometric effects. At this large p_{T} limit, we observe a systematic shift between simulation and data as well as variations between the different eras of the 2017 data, suggesting differences in the initial assumptions when processing the simulation tracks compared to those in data.

Figure 5.10 summarizes the differences between data and simulation as a function of p_{T} with a ratio of the mean impact parameter uncertainty evaluated at each p_{T} bin. This reveals up to 30-40% differences between simulation and data for the impact parameter uncertainty. Such differences can have consequences downstream following the vertex reconstruction where differences in the track resolution results in differences in the vertex uncertainty in d_{BV} between simulation and data. These differences between the uncertainties lead to differences in our evaluation of the vertex reconstruction efficiency after applying the vertex quality selection for simulation and data. Ultimately, this efficiency difference will affect the evaluation of the signal systematic uncertainty from vertex reconstruction.

To counteract this data/simulation difference and obtain a more reliable assessment of the systematic uncertainty, we can rescale the parameters of the track covariance matrix in MC events to more closely reflect the distributions in data. With a clearer understanding of the differences using the “ultra clean” track sample, the strict requirements can be relaxed separately in succession to recover the same analysis track quality

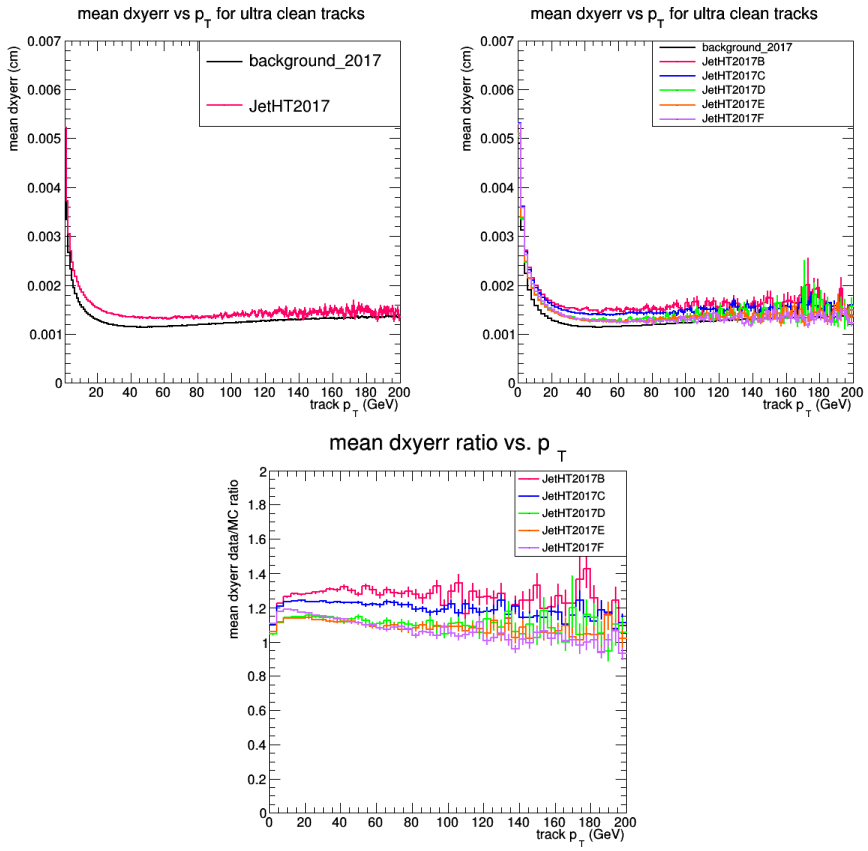

Figure 5.10: Comparison of mean track impact parameter uncertainty as a function of p_{T} between 2017 MC simulation (black) and 2017 data comparing the MC simulation with the whole 2017 run period combined (top left), comparing the MC simulation with the 2017 data decomposed into its run period eras (top right), and taking the ratio of the each of the data era curves to the MC curve (bottom)

selection requirements as described in Section 5.3. The additional tracks using the relaxed cuts exhibit the same behavior with minimal deviation from the curve derived from tracks satisfying the stronger "ultra clean" selection. There does appear to be a difference between tracks in the barrel compared to those reconstructed in the endcap, motivating separate treatment in two η bins divided at η of 1.5. Figures 5.11 and 5.12 show the $\sigma_{d_{\mathrm{xy}}}$ profiles and data/MC ratios after loosening the requirements to the standard analysis track selection and splitting the tracks in η for both 2017 and 2018. Due to the differences between the eras, each will be handled separately in this mitigation procedure, except in 2017, the behavior in eras D and E were similar enough to combine

into a single set. Splicing together functions to fit the ratio curves in either Figure 5.11 or 5.12 will provide the p_{T}-dependent scale factor to adjust the MC track covariance matrix elements as shown in Figure 5.13.

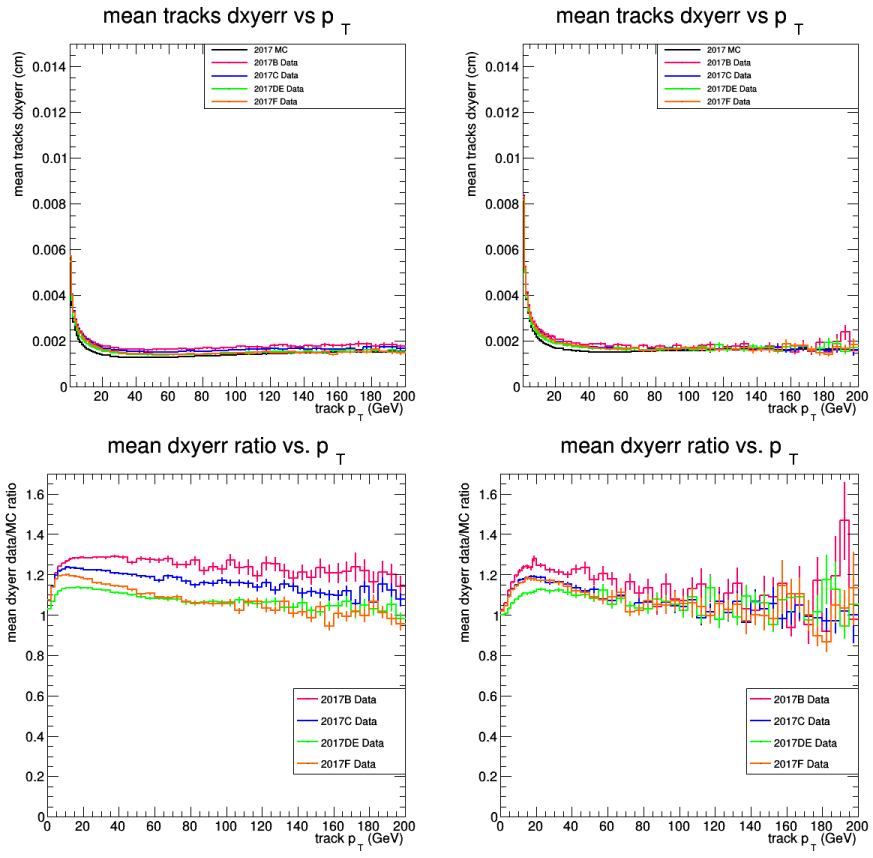

Figure 5.11: Mean $\sigma_{d_{\mathrm{xy}}}$ profiles (top) and mean $\sigma_{d_{\mathrm{xy}}}$ data/MC ratios (bottom) for $\eta < 1.5$ (left) and $\eta > 1.5$ (right) in 2017.

The transverse impact parameter uncertainty appears in the calculation of the uncertainty in the transverse vertex distance d_{BV}, while other elements of the covariance matrix do not contribute directly to this uncertainty. However, the other covariance matrix elements can still indirectly affect the reconstructed vertex during the reconstruction process in the arbitration step concerning the assignment of tracks to a given vertex, which can create shifts in the vertex variable distributions. Thus, differences between data and MC in the other covariance matrix elements are still worth consideration to cor-

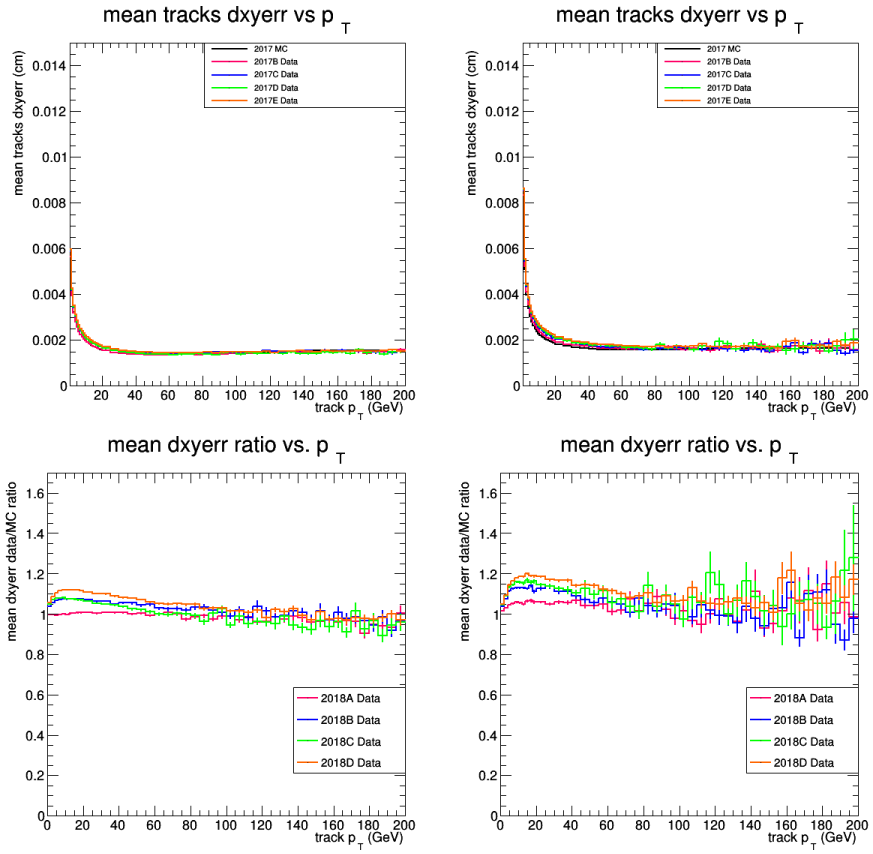

Figure 5.12: Mean $\sigma_{d_{xy}}$ profiles (top) and mean $\sigma_{d_{xy}}$ data/MC ratios (bottom) for $\eta < 1.5$ (left) and $\eta > 1.5$ (right) in 2018.

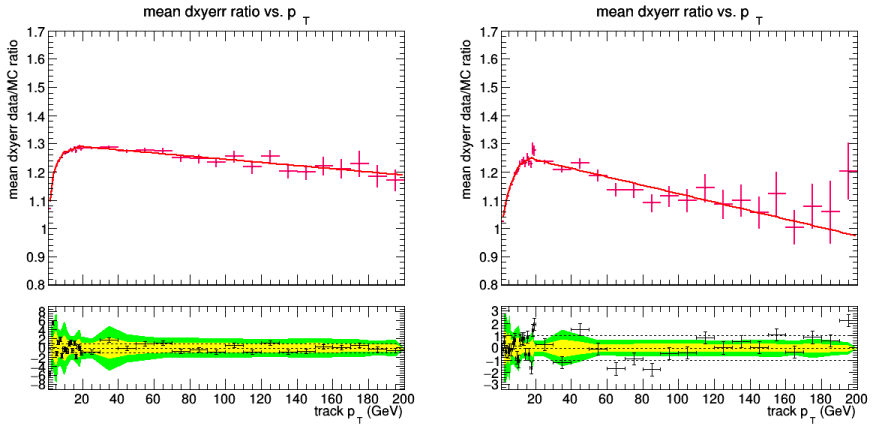

Figure 5.13: Fit applied to the 2017B era data/MC ratio vs. p_{T} curve for $|\eta| < 1.5$ (left) and $|\eta| > 1.5$ (right). The graphs below the fits show the residuals of the actual values to fitted curve, i.e. their difference divided by the uncertainty.

rect in a similar procedure as outlined for the transverse impact parameter uncertainty. Figure 5.14 shows the distributions of the other relevant parameters of the covariance matrix, namely the uncertainty in d_{sz} and the cross-term d_{xy}-d_{sz}.

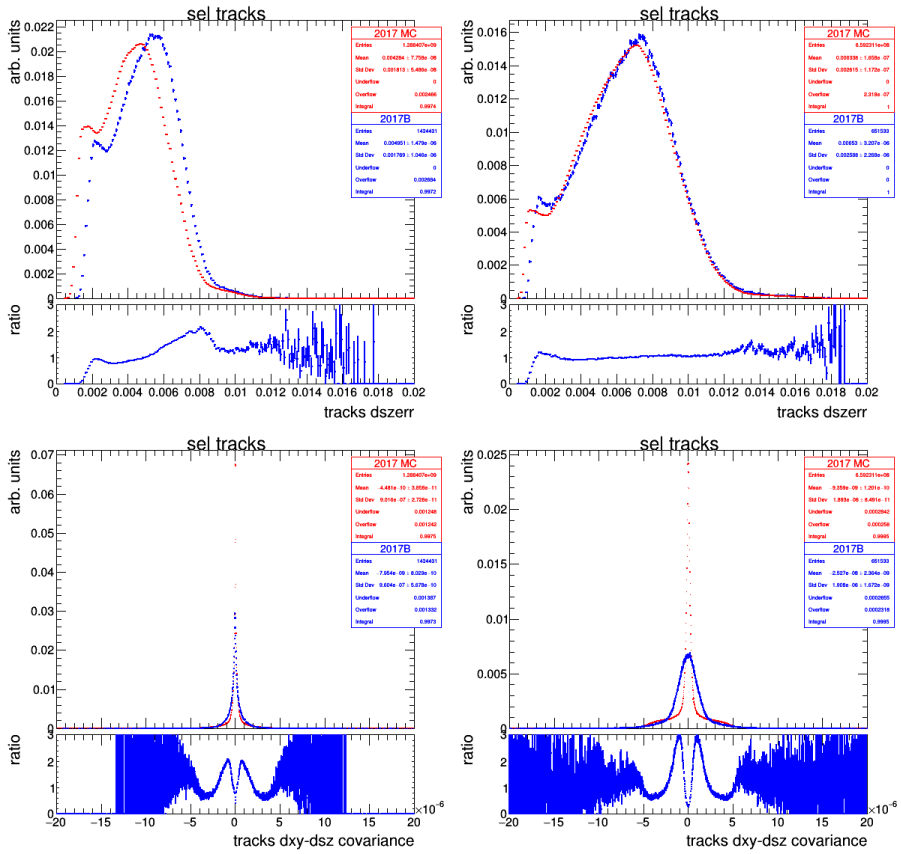

Figure 5.14: Comparing the uncertainty in d_{sz} (top) and the d_{xy}-d_{sz} cross term (bottom) for $|\eta| < 1.5$ (left) and $|\eta| > 1.5$ (right) between 2017 MC and 2017B data.

Implementing the same procedure for the uncertainty on d_{sz}, we can obtain the mean profile as a function of the track p_T as shown in Figure 5.15. From these curves, another correctional fit function can be derived from the ratios of data to MC to rescale the d_{sz} uncertainty in MC to match the mean uncertainties across the different data eras as shown in Figures 5.16 and 5.17.

As a final refinement of the overall track impact parameter uncertainty, the d_{xy}-d_{sz} covariance term, having been rescaled along with the individual variances throughout

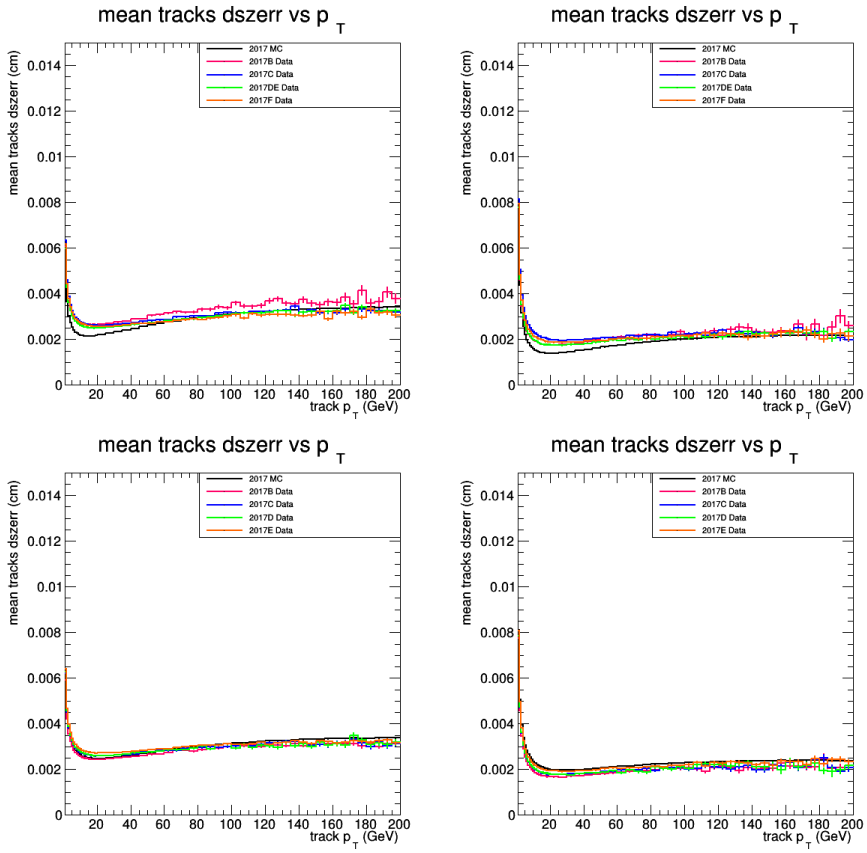

Figure 5.15: Comparison of mean track $\sigma_{d_{sz}}$ as a function of p_T between MC simulation (black) and data broken up by era for $|\eta| < 1.5$ (left) and $|\eta| > 1.5$ (right) in 2017 (top) and 2018 (bottom).

this procedure, can receive its own separate rescaling. The parameter can assume both negative and positive values as shown in Figure 5.14. The asymmetries in this shape are small enough that following the same prescription using the absolute value of this distribution is sufficient to derive a rescaling fit. Figure 5.18 shows the mean profile of the absolute value of this parameter as a function of the track p_T, comparing the unscaled MC to the data along with a curve representing the rescaled MC with the previous rescaling fits applied to match the data. Figure 5.19 shows the final fit for the ratio of the data profile curve to the rescaled MC curve.

To summarize, for each era of data collection in 2017 and 2018, a curve is fitted to

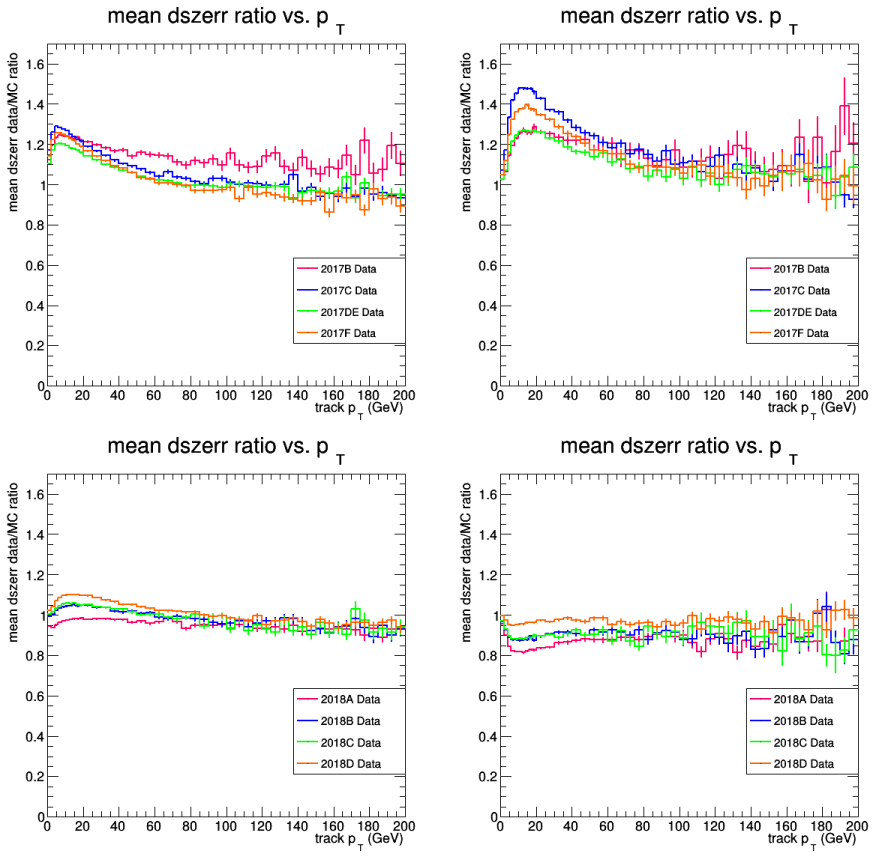

Figure 5.16: Ratio of the mean track d_{sz} uncertainty in data (broken up by era) and MC for $|\eta| < 1.5$ (left) and $|\eta| > 1.5$ (right) in 2017 (top) and 2018 (bottom).

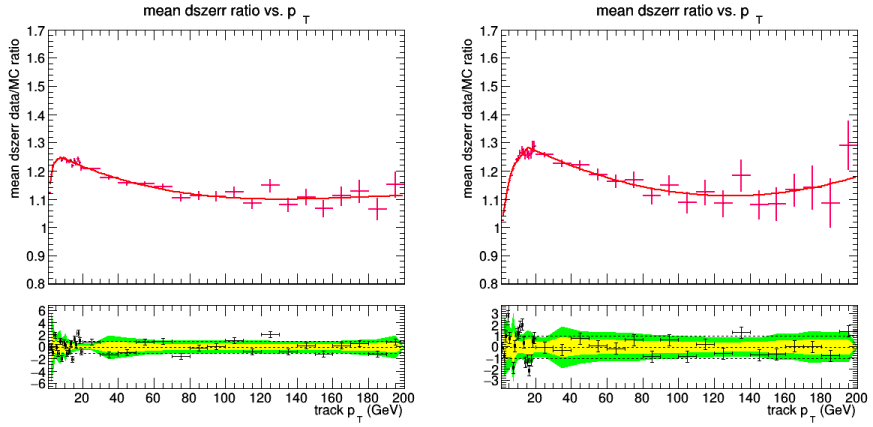

Figure 5.17: Fit applied to the 2017B era data/MC d_{sz} uncertainty ratio vs. p_T curve for $|\eta| < 1.5$ (left) and $|\eta| > 1.5$ (right).

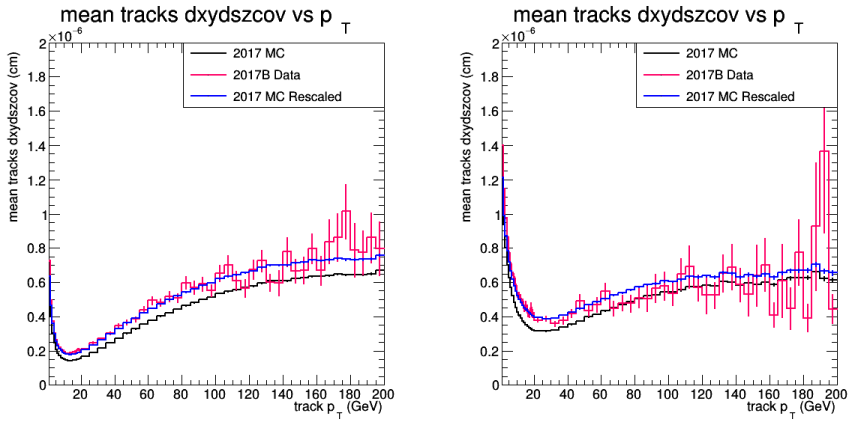

Figure 5.18: Comparison of mean track $|d_{xy}$-$d_{sz}|$ covariance term as a function of p_T between 2017 MC simulation (black), 2017B data (pink), and the 2017 MC simulation with d_{xy} and d_{sz} rescaled (blue) for $|\eta| < 1.5$ (left) and $|\eta| > 1.5$ (right).

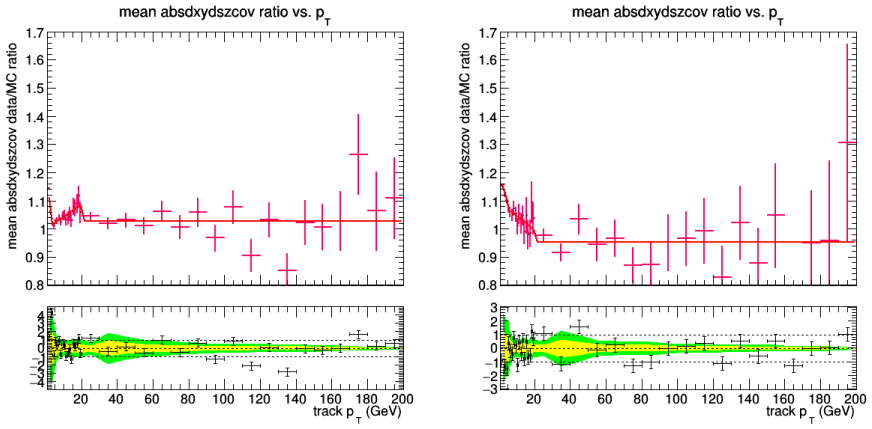

Figure 5.19: Fit applied to the ratio of the 2017B data to the rescaled MC $|d_{xy}$-$d_{sz}|$ covariance term vs. p_T curve for $|\eta| < 1.5$ (left) and $|\eta| > 1.5$ (right).

the ratio of data to simulation for the components of the track covariance matrix elements d_{xy}, d_{sz}, and the d_{xy}-d_{sz} covariance term. Using these curves, the corresponding component in signal simulation is rescaled in order to better reflect the data in any studies of the signal efficiency. Figure 5.20 shows the ratio of the signal efficiencies with and without applying the track covariance matrix element rescaling. The effect is most prominent for smaller lifetimes as expected, as the upward scaling of the covariance terms will make it more difficult for less displaced tracks from short lifetime LSPs to satisfy the impact parameter significant requirement. Additionally, as will be discussed

in the vertex reconstruction and selection section, the covariance matrix elements are used to compute uncertainties in the vertex position, a variable on which a requirement is placed.

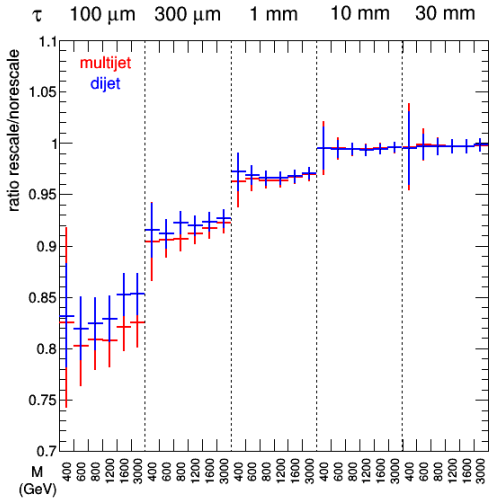

Figure 5.20: Ratio of rescaled track covariance matrix signal efficiencies to the non-rescaled signal efficiencies as a function of signal mass and lifetime for dijet and multijet signal models.

Cross-check using displaced tracks in K^0_S

To validate this procedure, we study the displaced tracks from K^0_S candidates. K^0_S candidates are identified by examining the invariant mass of all pairs of tracks that pass our track selection requirements. We vertex all oppositely-charged pairs of selected tracks and require it to satisfy $\chi^2/\text{dof} < 7$ and have the vertex positioned within the beam pipe. Other requirements such as $c\tau > 268\,\mu$m, $\rho > 0.268$ cm, and requiring the cosine of the 2D angle between fitted candidate momentum and flight direction from primary vertex to be larger than 0.99975 improve the purity of the K^0_S sample. The K^0_S post-fit candidate mass is shown in Figure 5.21. K^0_S candidates are taken from a mass window around the nominal K^0_S mass (490–505 MeV), while low-mass (420–475 MeV) and high-mass (525–580 MeV) sidebands are used to estimate the background.

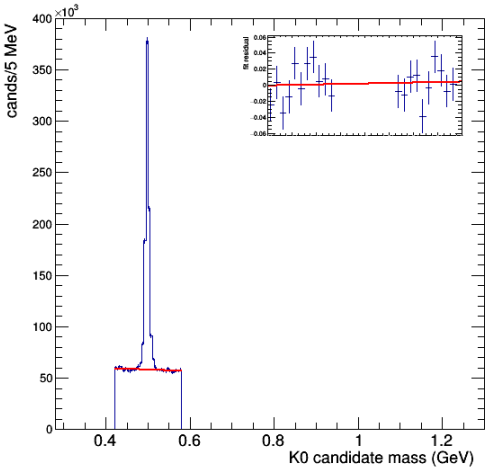

Figure 5.21: Distribution of the post-vertex fit invariant mass of K^0_S candidates in 2017 simulation scaled to 41.5 fb^{-1}. The candidates are pairs of tracks that pass our track selection requirements and form a vertex with $\chi^2/\text{dof} < 7$. The candidates in the K^0_S mass window (490–505 MeV) are chosen, and low-mass (420–475 MeV) and high-mass (525–580 MeV) sidebands are fit to estimate the background. The inset shows the fit residuals for the low-mass and high-mass sidebands.

Using the 2017B dataset, in the on-peak mass window there are 59164 candidates, 18661 of which are estimated to be background, giving a purity of 76%. For the 2017 MC simulation scaled to the 4.8 fb^{-1} of this dataset, there are 69892 candidates with 20170 estimated background giving a 78% purity.

Figure 5.22 shows the resulting $\sigma_{d_{xy}}$ profile for the displaced tracks of the K^0_S candidates after following the recipe described above. The "inclusive" set of curves correspond to the entire set of tracks in each corresponding dataset. The data to simulation ratios of the two sets of tracks is shown in the ratio plot along the bottom with a blue curve overlaid representing the fitted curved (which was constructed from the inclusive curve, and should thus align with the gray ratio curve.) The black ratio curve, corresponding to the data to simulation ratio of the mean $\sigma_{d_{xy}}$ in displaced tracks shows agreement with the blue and gray curve, providing validation that the rescaling performs its intended function on genuine displaced tracks.

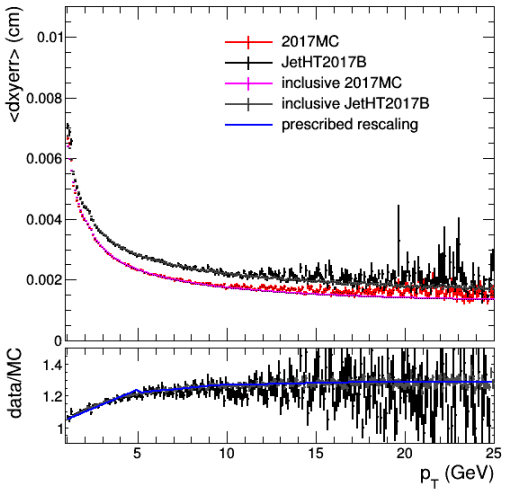

Figure 5.22: Mean $\sigma_{d_{xy}}$ as a function of p_T for the displaced tracks of the K^0_S candidates in MC background simulation (red) and 2017B data (black). The pink and grey lines, denoted "inclusive" correspond to the full dataset for the MC simulation and 2017B datasets respectively.

5.3.3 Vertex reconstruction

The next step in the vertex reconstruction procedure is to generate seed vertices from all pairs of tracks that satisfy the track selection criteria. The Kalman filter method is used to form a vertex from two or more tracks. The vertex will only be considered valid if its χ^2 per degree of freedom is less than 5. If two vertices share a track and the three-dimensional distance between the vertex pair is less than 4 times the uncertainty in that distance, these are considered compatible and a vertex fit is applied to the complete set of tracks from both vertices. If the resulting fit satisfies the χ^2 requirement, the two vertices are replaced by one single merged vertex. Otherwise, the two vertices remain separated, requiring a track arbitration step to decide which vertex is assigned the shared track. The track arbitration depends on the value of the track's three-dimensional impact parameter significance with respect to each of the vertices. If both impact parameters are within 1.5 standard deviations of both vertices, the shared track is assigned to the vertex with the larger number of tracks already; if the track has an impact parameter

that is more than 5 standard deviations from either vertex, the shared track is removed from that vertex; otherwise, the shared track is assigned to the vertex to which it has the smaller impact parameter significance. When a track is dropped from a vertex, that vertex is refit with its remaining tracks and replaced with a new vertex if the fit satisfies the χ^2 requirement, otherwise the vertex is erased entirely. Pairs of vertices are merged iteratively following this algorithm until no two vertices share a track.

Occasionally, a vertex with more than two tracks is formed from the accidental intersection of tracks that originate from separate pileup vertices, an example is shown in Fig. 5.23. As a final step in the reconstruction to suppress these, we remove any track from a vertex whose position along the beam axis shifts appreciably when the track is dropped from the vertex fit. The threshold for removal is a 50 μm shift. This additional procedure is a new refinement in this iteration of the analysis and removes a significant number of background vertices that have large displacements with minimal impact on signal efficiency.

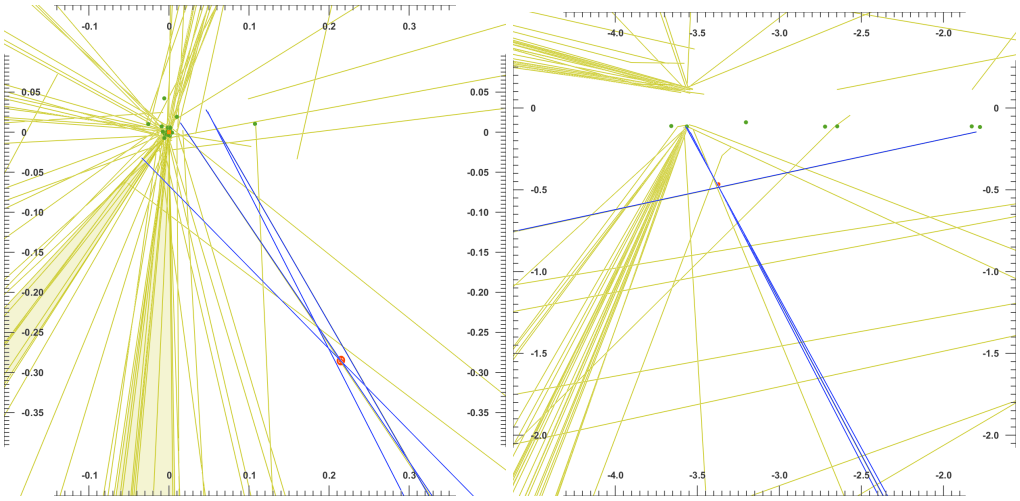

Figure 5.23: Event display for (run, lumi, event) number (322431, 170, 289311315) in 2018 in the ρ-ϕ plane (left) and ρ-z plane (right). The vertex has a displacement of approximately 0.36 cm from the beam axis, which is the point of intersection of the three tracks from the primary vertex and one additional track from a pileup vertex.

5.3.4 Vertex selection

We select vertices with features consistent with a signal vertex by requiring the vertices to: have at least five tracks; have an x-y distance from the beam pipe center of less than 20.9 mm to suppress background vertices arising from interactions of particles with the beam pipe or detector material; have an x-y displacement from the beam axis, defined as d_{BV}, of at least 100 μm to suppress background from displaced primary vertices; and have an uncertainty in d_{BV} of less than 25 μm to select vertices formed from high quality tracks with large opening angle, which also suppresses vertices from b jets that tend to have narrow opening angles.

The fiducial requirement received a minor adjustment in this iteration of the analysis. In 2016, the selection was made on the distance of the vertex to the geometric detector origin, rather than the beam pipe center, and required to be within a 20 mm radius. This sufficiently removed vertices produced from interactions with the beam pipe and detector components. However, small shifts in the beam pipe in both 2017 and 2018 have led to many vertices in data that occupy a small ϕ sector where the beam pipe intersected with this fiducial cut. We have therefore modified the fiducial requirement to cut directly on the vertex distance from the beam pipe center instead, located (0.113, -0.180) cm relative to the geometric center in 2017 and (0.171, -0.175) cm relative to the geometric center in 2018. Figure 5.25 shows the new fiducial requirement.

This search focuses on pair-produced signal models, which we impose by requiring events to have two reconstructed vertices satisfying the requirements listed above. Few events in background contain even one reconstructed displaced vertex, with occurrences of higher vertex multiplicity events even rarer. Simulations of background have predicted less than one event in the two-vertex search region for 101 fb^{-1} of data, which can be seen in Table 5.5, which lists the yields for background simulation and various

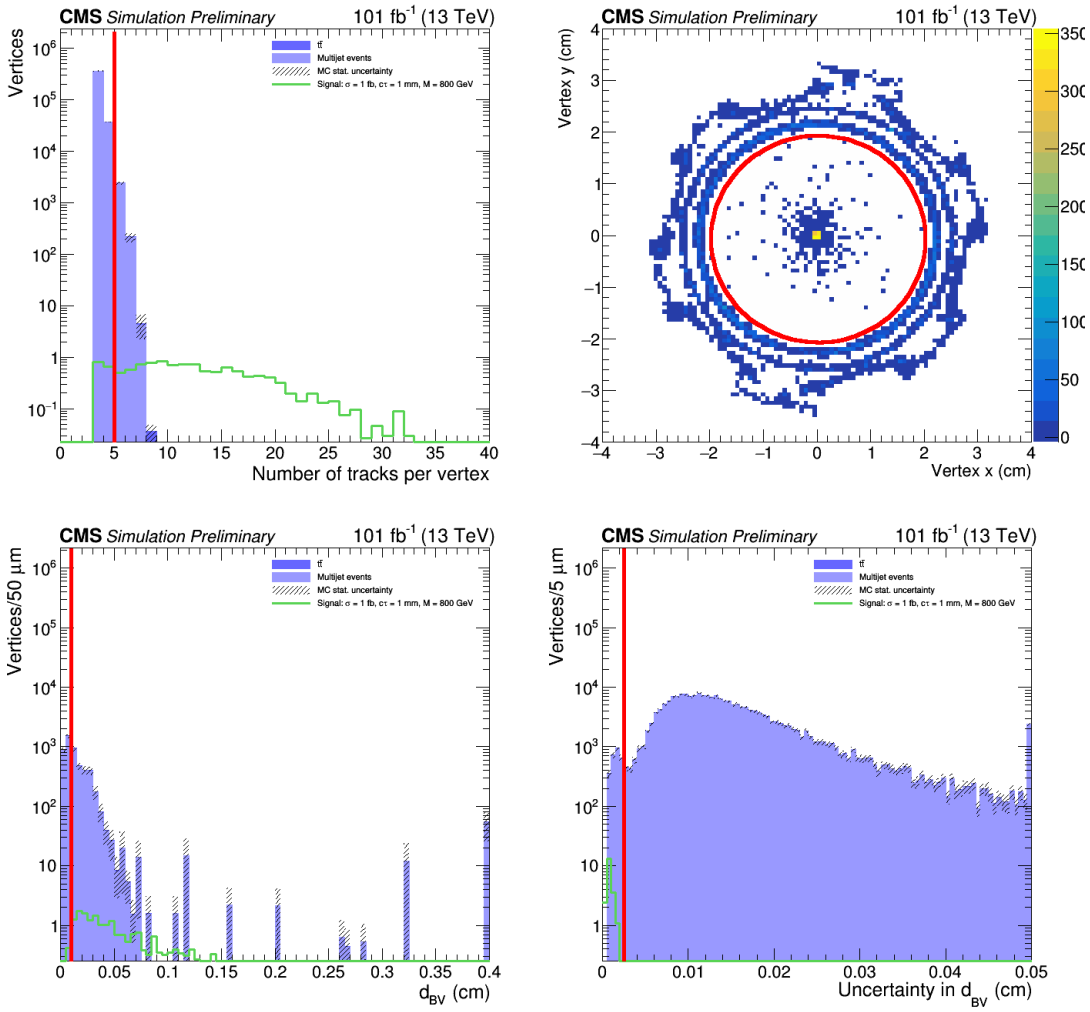

Figure 5.24: Distributions of variables used for vertex selection in the 2017 and 2018 background MC as listed in Tables 5.2and 5.3. These are "n-1" plots: all event preselection and vertex selection criteria have been applied, except for the one related to the variable shown. The number of tracks per vertex is required to be at least five (top left); the x-y distance of the vertex from the beam pipe center is required to be less than 20.9 mm (top right); the x-y distance of the vertex from the beam axis, d_{BV}, is required to be at least 100 μm (bottom left); the uncertainty in d_{BV} is required to be less than 25 μm (bottom right).

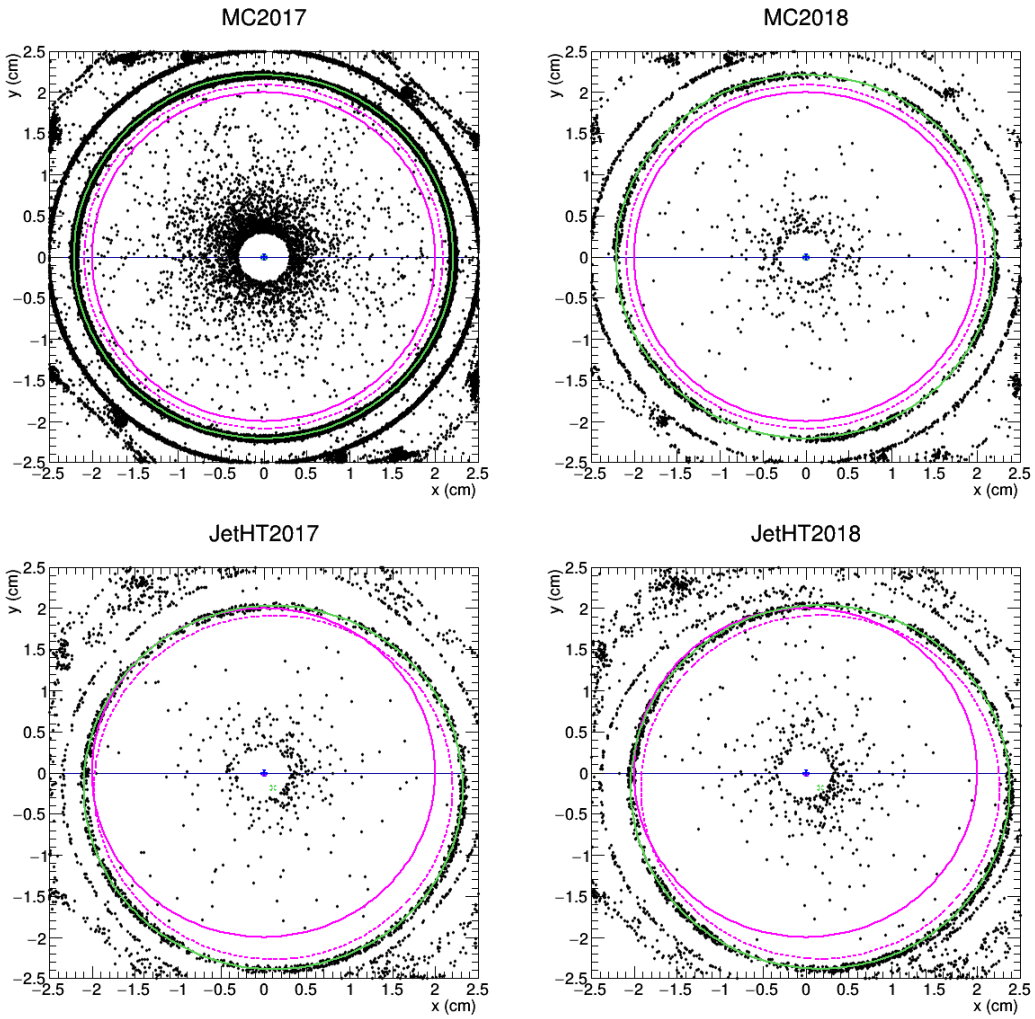

Figure 5.25: 3-track 1-vertex positions in the *x*-*y* plane for MC (top) and data (bottom) in both 2017 (left) and 2018 (left). The solid pink circle with the pink marker at the center corresponds to the previous fiducial requirement from 2016. The solid green circle and marker represent the beam pipe, located at (0.113, -0.180) cm relative to the geometric center in 2017 and (0.171, -0.175) cm in 2018. The dotted pink line is the new fiducial requirement. The blue marker marks the location of the beamspot.

signal points. However, a reliable measurement of the signal strength in data requires a more precise estimation of background, which we derive using a data-driven method.

Table 5.5: Combined 2017 and 2018 background simulation yields for ≥5-track vertices, representing 101 fb^{-1} and 2017 signal simulation sample yields for ≥5-track vertices.

MC Sample	one-vertex	two-vertex
QCD, $700 < H_\text{T} < 1000\,\text{GeV}$	9 ± 9	0
QCD, $1000 < H_\text{T} < 1500\,\text{GeV}$	275 ± 46	0
QCD, $1500 < H_\text{T} < 2000\,\text{GeV}$	111 ± 11	0
QCD, $H_\text{T} > 2000\,\text{GeV}$	34 ± 4	0
$\text{t}\bar{\text{t}}$, $600 < H_\text{T} < 800\,\text{GeV}$	4.8 ± 1.2	0
$\text{t}\bar{\text{t}}$, $800 < H_\text{T} < 1200\,\text{GeV}$	7.3 ± 1.4	0
$\text{t}\bar{\text{t}}$, $1200 < H_\text{T} < 2500\,\text{GeV}$	2.5 ± 0.8	0
$\text{t}\bar{\text{t}}$, $H_\text{T} > 2500\,\text{GeV}$	0.06 ± 0.06	0
Total background	444 ± 73	0
Multijet signals: $\sigma = 1$ fb		
$c\tau = 100\,\mu\text{m}$, $M = 800\,\text{GeV}$	5.40 ± 0.21	1.05 ± 0.09
$c\tau = 300\,\mu\text{m}$, $M = 800\,\text{GeV}$	11.06 ± 0.30	6.07 ± 0.22
$c\tau = 1\,\text{mm}$, $M = 800\,\text{GeV}$	9.18 ± 0.28	15.35 ± 0.36
$c\tau = 10\,\text{mm}$, $M = 800\,\text{GeV}$	5.32 ± 0.21	22.95 ± 0.44

While we select vertices with at least five tracks for further analysis, vertices composed of at least three tracks function as useful control samples to validate the background estimation method. Events with 3-track and 4-track vertices occur a factor of ten to a hundred times more often than events with ≥5-track vertices; moreover, the large yield dwarfs any potential contamination by signal, so they provide a nearly pure background sample. Distributions of event-level variables (e.g. H_T, jet multiplicity) and vertex-level variables (e.g. d_BV, uncertainty in d_BV) in simulation are similar for events with 3, 4, and ≥5-track vertices, ensuring that the vertices of lower track multiplicity

provide a reliable approximation for the 5-track vertex event sample. Table 5.6 shows the event yields in the control regions for simulation.

Table 5.6: Combined 2017 and 2018 background simulation yields for 3-track and 4-track vertices, representing 101 fb^{-1} and 2017 signal MC sample yields for 3-track and 4-track vertices.

	3-track		4-track × 3-track	4-track	
MC sample	one-vertex	two-vertex	two-vertex	one-vertex	two-vertex
QCD, $700 < H_\text{T} < 1000\,\text{GeV}$	607 ± 98	0	0	74 ± 35	0
QCD, $1000 < H_\text{T} < 1500\,\text{GeV}$	15971 ± 348	43 ± 16	15 ± 10	2592 ± 142	0
QCD, $1500 < H_\text{T} < 2000\,\text{GeV}$	6378 ± 78	14 ± 4	5 ± 2	1150 ± 33	0.89 ± 0.63
QCD, $H_\text{T} > 2000\,\text{GeV}$	1768 ± 26	6 ± 2	1.6 ± 0.7	339 ± 11	0.26 ± 0.19
$\text{t}\bar{\text{t}}$, $600 < H_\text{T} < 800\,\text{GeV}$	409 ± 11	2.2 ± 0.8	1.1 ± 0.5	67 ± 4	0.16 ± 0.16
$\text{t}\bar{\text{t}}$, $800 < H_\text{T} < 1200\,\text{GeV}$	417 ± 11	2.3 ± 0.7	1.1 ± 0.5	64 ± 4	0
$\text{t}\bar{\text{t}}$, $1200 < H_\text{T} < 2500\,\text{GeV}$	108 ± 5	0.9 ± 0.5	0.1 ± 0.1	17 ± 2	0
$\text{t}\bar{\text{t}}$, $H_\text{T} > 2500\,\text{GeV}$	1.9 ± 0.3	0.04 ± 0.04	0.02 ± 0.01	0.2 ± 0.1	0
Total background	25660 ± 577	68 ± 24	24 ± 13	4305 ± 150	1.3 ± 0.8
Multijet signals: $\sigma = 1$ fb					
$c\tau = 100\,\mu\text{m}$, $M = 800\,\text{GeV}$	5.19 ± 0.21	0.80 ± 0.08	1.67 ± 0.12	2.51 ± 0.14	0.18 ± 0.04
$c\tau = 300\,\mu\text{m}$, $M = 800\,\text{GeV}$	6.80 ± 0.24	1.77 ± 0.12	3.20 ± 0.16	4.10 ± 0.18	0.47 ± 0.06
$c\tau = 1\,\text{mm}$, $M = 800\,\text{GeV}$	8.04 ± 0.26	2.53 ± 0.14	4.84 ± 0.20	5.32 ± 0.21	0.71 ± 0.08
$c\tau = 10\,\text{mm}$, $M = 800\,\text{GeV}$	7.25 ± 0.25	1.75 ± 0.12	3.52 ± 0.17	1.57 ± 0.11	0.49 ± 0.06

Finally, Table 5.7 shows the event yields in each of the control regions in the 2017 and 2018 data.

5.4 Search strategy

We select events that contain at least two vertices each with 5 or more tracks to search for pair-produced long-lived particles. We use the distance between the vertices in the x-y plane, defined as d_VV, as the main discriminating variable between signal and the SM background. In signal events, the pair-produced long-lived particles tend to be emitted

Table 5.7: Number of events in the full 2017 and 2018 data, representing 41.5 fb^{-1} and 59.7 fb^{-1}, respectively.

Control sample	Number of 2017 events	Number of 2018 events
3-track one-vertex	32152	29666
3-track two-vertex	113	72
4-track one-vertex	7838	6892
4-track × 3-track two-vertex	59	42
4-track two-vertex	9	3
5-track one-vertex	1303	908

back-to-back, resulting in larger vertex separations than in background where d_{VV} tends to be small. In events with more than two vertices, the two vertices with the highest number of tracks are chosen for the d_{VV} calculation. If there is a tie in the number of tracks, a mass value is assigned to the vertex, reconstructed from the track momenta associated with the vertex, and the one with the higher mass is chosen. However, in the 2017 and 2018 data, we observe no events with more than two reconstructed vertices.

The d_{VV} variable cannot be reliably ascertained in background simulations with fewer than one event in the search region. Additionally, background vertices are sensitive to the misreconstruction of tracks, which is difficult to faithfully replicate in simulation. Thus we construct a background template distribution of d_{VV} using one-vertex events in data, as described in Section 5.6. Figure 5.26 compares the d_{VV} distributions for simulated multijet signals of various mean proper decay lengths, mass of 800 GeV, and production cross section of 0.3 fb overlaid with the background template derived from data. The background peaks near 0.3 mm with low probability above 0.7 mm where most signal events are found.

Ultimately, the background and signal templates are fit to the d_{VV} distribution observed in data in order to extract the final signal yield. When fitting the d_{VV} distribution,

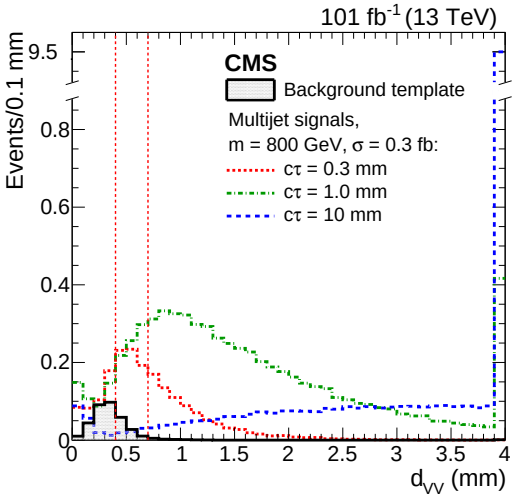

Figure 5.26: The distribution of distances between vertices in the x-y plane, d_{VV}, for simulated multijet signals with mass of 800 GeV, production cross section of 0.3 fb, and $c\tau$ = 0.3, 1.0, and 10 mm. The background template distribution derived from data is overlaid. All vertex and event selection criteria have been applied. The last bin includes the overflow events. The two vertical red dashed lines delineate the boundaries of the bins used in the fit.

we divide it into three bins: 0-0.4 mm, 0.4-0.7 mm, and 0.7-40 mm. The two bins with $d_{\mathrm{VV}} > 0.4$ mm have the least background. This binning scheme optimizes the signal significance in models with lifetimes in the 0.1 to 100 mm range.

5.5 Signal efficiency measurement

To study the vertex reconstruction efficiency, we apply the reconstruction procedure to track collections that have been manually displaced from the primary vertex. This produces artificial signal-like vertices in data and simulation. Starting from events satisfying the trigger and offline preselection requirements that also have a well-reconstructed primary vertex, we randomly select light parton or b quark jets that have $p_{\mathrm{T}} > 50$ GeV and at least four matched particle-flow candidate tracks. The jets are identified as a light or b quark jet based on their heavy flavor jet tagger discriminator value. The tracks asso-

ciated with the selected jets are then displaced in the direction defined by the vector sum of the selected jet momenta. The displacement vector direction is smeared to capture the mismeasurement from tracking inefficiency or missing neutral particles. The displacement magnitude is sampled from an exponential distribution with a configurable scale parameter $c\tau$ and is restricted to values between 0.3 and 20 mm, which is characteristic of vertices that would enter the third bin in d_{VV}. The track impact parameter resolutions in simulation are scaled to match data as a function of p_{T} and η. After track selection, vertex reconstruction, and vertex selection, we compute the fraction of events passing the event selection criteria that also contain a vertex reconstructed within 84 μm (i.e. the measured resolution of the vertex reconstruction algorihtm) of the expected location. This efficiency carries the track reconstruction and selection efficiencies in addition to the one from vertex reconstruction. Figure 5.27 shows the efficiencies as functions of variables relevant to the artificial vertex or the overall event when moving $n_l = 2$ light jets and $n_b = 1$ b jets.

5.5.1 Dijet signal

For dijet signals, we displace the tracks associated with two light jets, $(n_l, n_b) = (2, 0)$. The observed energy and angular separation of the artificially displaced jets are different from the jets in the dijet signal simulation. We therefore apply a set of requirements to the moved jets in order to more faithfully reproduce the signal jet characteristics. When moving two jets, we require that the softer jet have at least $p_{\mathrm{T}} > 125$ GeV and the angular distance between the two jets must satisfy $\Delta R < 1.0$. We order and index all jets in the event by decreasing p_{T} and require that the harder moved jet has an index no more than three less than the softer moved jet. This latter requirement models the tendency of jets within the dijet signal to have similar p_{T}.

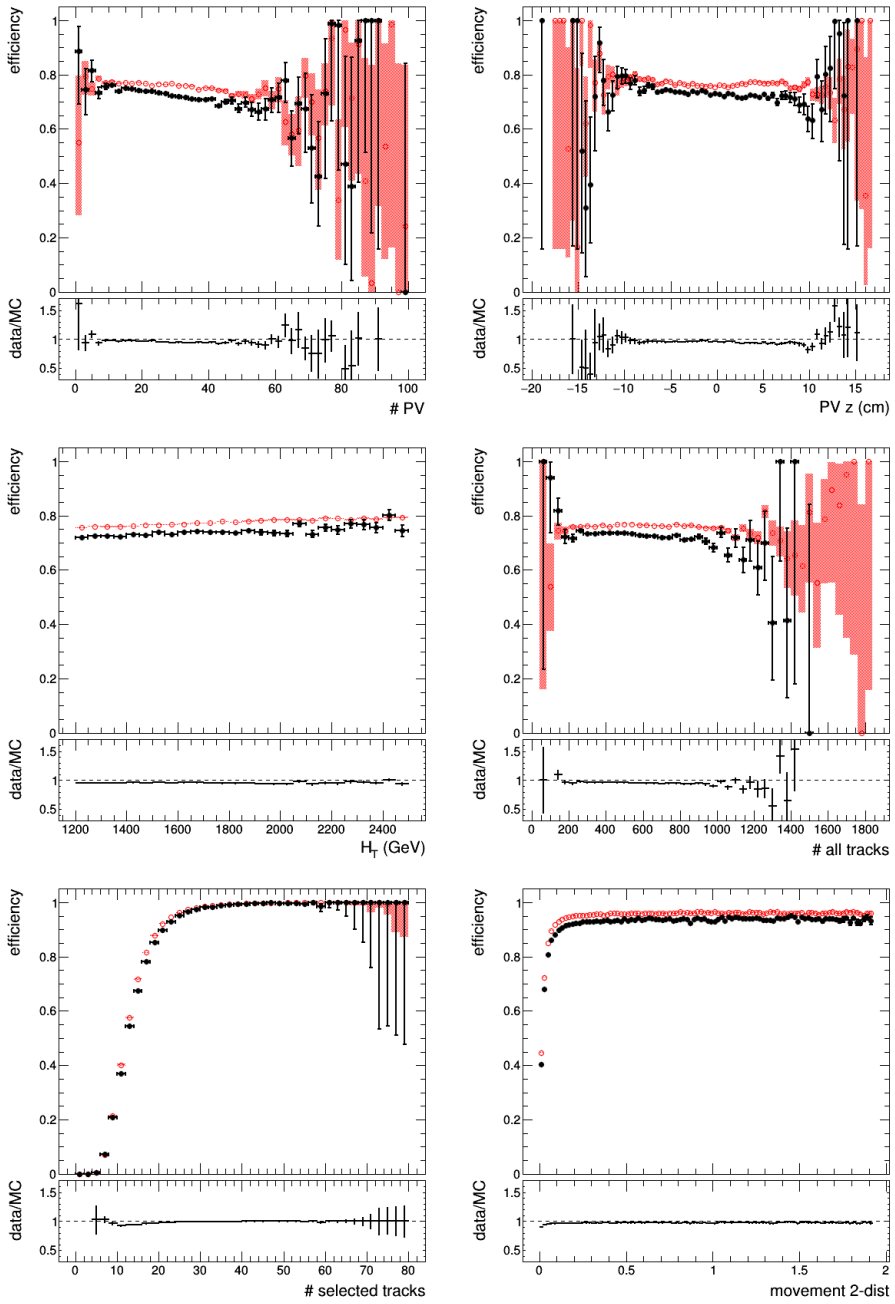

Figure 5.27: Vertex reconstruction efficiency measured as a function of several variables for artificially displaced vertices in the data (black) and simulation (red) for the $(n_l, n_b) = (2, 1)$ configuration. The variables are: the number of primary vertices in the event (top left); the primary vertex z-coordinate relative to the beam spot (top right); the H_T of jets in the event (middle left); the total number of tracks in the event (middle right); the selected number of tracks displaced to the artificial vertex (bottom left); the artificial flight distance in the transverse plane (bottom right).

In the dijet signal, vertex reconstruction efficiency significantly depends on the separation in ϕ of the associated jets, the absolute value of which we label $|\Delta\phi_{\mathrm{JJ}}|$. In particular, the efficiency is suppressed when the two jet momentum vectors point in opposite directions, i.e. large $|\Delta\phi_{\mathrm{JJ}}|$, as shown in Fig. 5.28. This is because of the large resulting uncertainty in the vertex position when reconstructing a vertex from anti-parallel or back-to-back tracks and the vertex displacement vector is aligned with the jet axes. This can happen often in this procedure when selecting two back-to-back jets whose momenta are imbalanced, and the resulting artificial displacement is in the same direction as one of the jet momenta. Quantitatively, Fig. 5.29 shows how the vertex position uncertainty differs for large and small $|\Delta\phi_{\mathrm{JJ}}|$. We therefore reweight efficiencies of these large-angle jet-pair events and the rest of the events based on the proportion of large-angle jet-pair vertices observed in the signal simulation. We define $|\Delta\phi_{\mathrm{JJ}}| > 2.7$ as the cutoff for large-angle jet-pairs. The fraction events in dijet signal simulation with large angle separations between the jets in the LSP decay is less than 35%, which is shown in Fig. 5.30. To account for these vertices, we combine the relative difference in efficiency between data and MC, weighted by the fraction of vertices with $|\Delta\phi_{\mathrm{JJ}}| > 2.7$ in signal Monte Carlo. This fraction, f_{B2B}, is less than 35% for each dijet signal point and is shown in Figure 5.30.

The result of this procedure for the data and background simulation are shown in Fig. 5.31. These efficiencies are used to rescale the signal simulation yields to reflect the efficiency one can expect to observe in data. A systematic uncertainty is assigned to this rescaling procedure using the relative difference between the rescaled signal and the nominal signal efficiency, which can be seen in Fig. 5.32. The vertex reconstruction efficiency differences range from 5% for the longest lifetimes to 16% for the shortest.

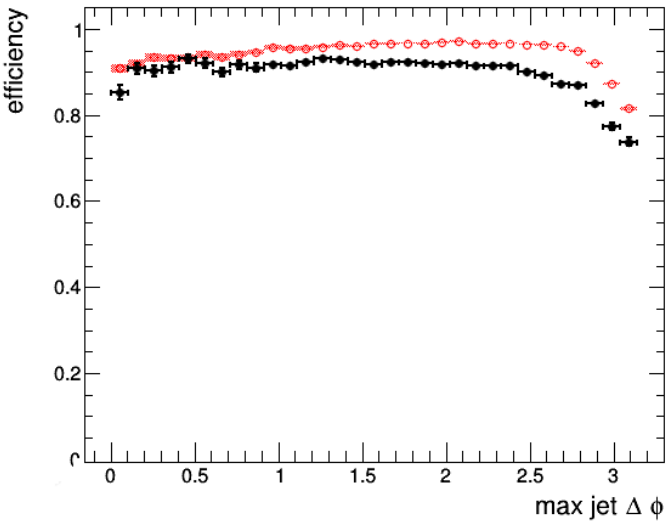

Figure 5.28: Vertexing efficiency measured as a function of $|\Delta\phi_{JJ}|$ for the dijet signal simulation, where the 600 GeV, 10 mm signal point is shown in black and the 1200 GeV, 10 mm signal point is shown in red. All vertex quality cuts are applied.

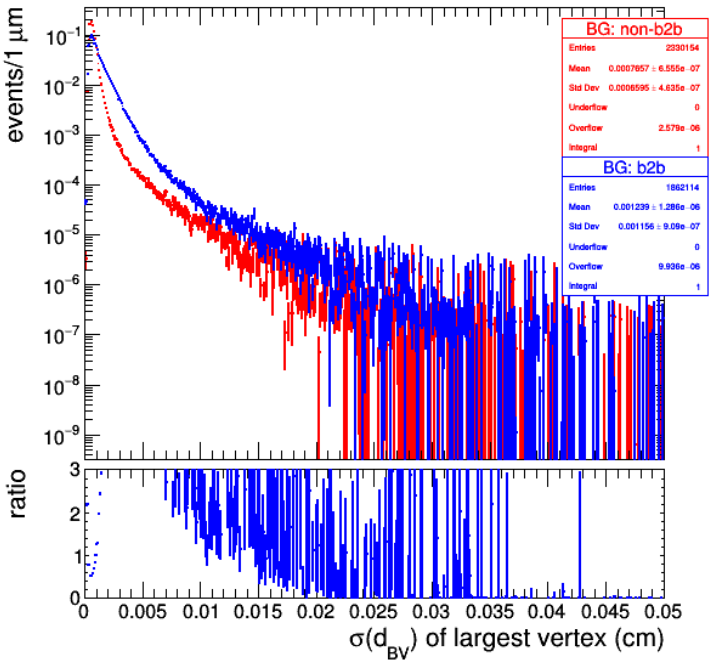

Figure 5.29: Uncertainty in distance between beamspot and artificial vertex in 2017 background data for the $(n_l, n_b) = (2, 0)$ case for a $|\Delta\phi_{JJ}| \leq 2.7$ slice (in red) and a $|\Delta\phi_{JJ}| > 2.7$ slice (in blue). No vertex quality cuts are applied, and distributions have been normalized to unity. Vertices with $\sigma(d_{BV}) > 0.0025$ cm are removed.

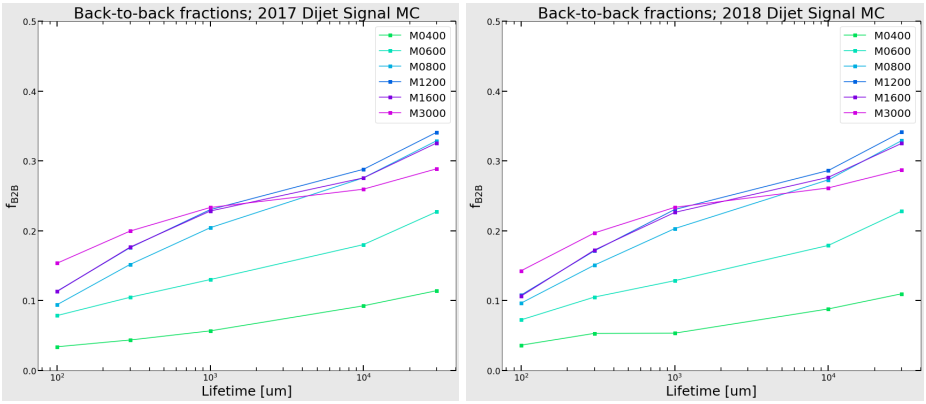

Figure 5.30: The fraction of vertices corresponding to jets with $|\Delta\phi_{JJ}| > 2.7$ across the 2017 (left) and 2018 (right) dijet signal samples.

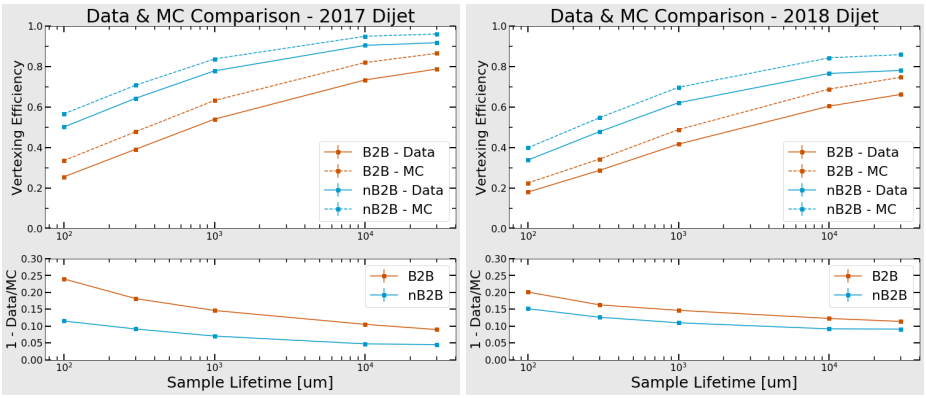

Figure 5.31: Comparison of overall vertex reconstruction efficiencies obtained for artificially displaced vertices formed from 2017 (left) and 2018 (right) data and simulation for two moved light jets with all vertex quality cuts applied. Vertices corresponding to jets with $|\Delta\phi_{JJ}| > 2.7$ are shown in orange, and those corresponding to jets with $|\Delta\phi_{JJ}| \leq 2.7$ are shown in blue. Relative differences between data and MC are shown in the bottom subplot.

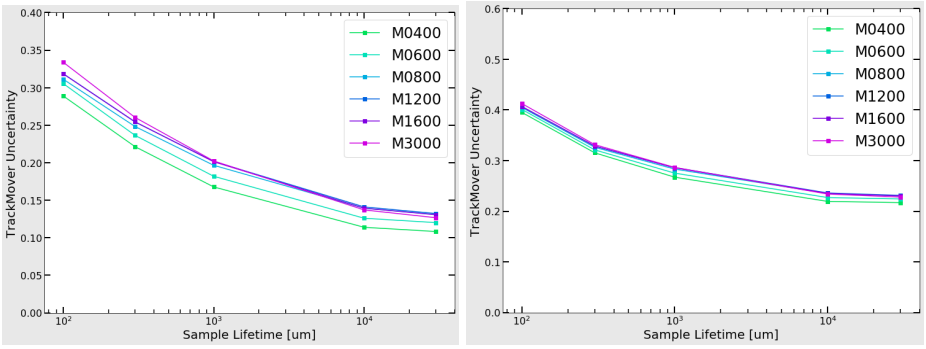

Figure 5.32: Fractional uncertainty in vertex reconstruction for 2017 (left) and 2018 (right) dijet signal samples for varying signal masses and lifetimes.

5.5.2 Multijet signal

Initial studies into the multijet signal samples suggest that signal-mocking cuts analogous to those of the dijet signal samples are not needed. Furthermore, since there are more moved jets per artificial vertex when $n_l + n_b \geq 3$, the vertexing efficiency does not depend on $|\Delta\phi_{JJ}|$ as strongly as the dijet artificial vertices. Thus, a slicing-and-weighting scheme is not applied in the multijet case.

However, the multijet study can be deconstructed in another way to assess the differences between data and simulation. For multijet artificial vertices, we select and displace tracks associated to a $n_l = 3$ and $n_b = 2$ configuration of jets. Each event is reweighted based on the d_{BV} value of the artificial vertex to match the d_{BV} distribution of vertices in signal simulation that fall into the third d_{VV} bin with the lowest background. The d_{BV} distribution is shown in Fig. 5.33. In the multijet scenario, a descriptive indicator for the vertex reconstruction efficiency is the total number of nearby high quality displaced tracks, i.e. tracks passing the track selection criteria, that pass close to the location of the artificially displaced vertex, which we will refer to as the number of *close seed tracks* per vertex. We quantify the distance or proximity of the track relative to the artificial displacement location using the impact parameter significance, where we identify tracks with a significance less than 5 as being a close seed track. Figure 5.34 shows the distribution of close seed tracks per vertex in data and simulation. For a given lifetime point, the mean of the close seed track distributions for artificial data and simulation differ by 2 to 3 tracks. Figure 5.35 shows the same distribution for signal samples of three different masses. Figure 5.36 shows the vertex reconstruction efficiency as a function of this variable in data and simulation. The efficiency turn-on curves for the artificial vertices in data and simulation differ by 0.35 tracks in 2017 and 1.0 track in 2018.

We capture the cumulative effects of these differences between data and simulation,

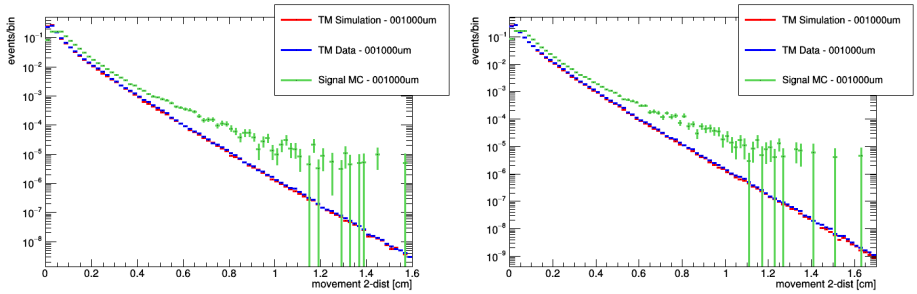

Figure 5.33: d_{BV} distribution for the vertices in signal simulation and the artificial vertices in data and simulation for 2017 (left) and 2018 (right).

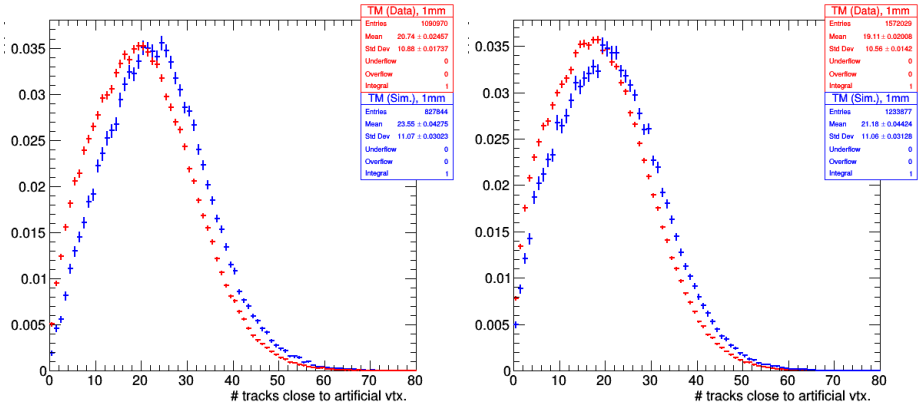

Figure 5.34: Distribution of close seed tracks near the artificial displaced vertex of 1 mm lifetime in data (red) and simulation (blue) in 2017 (left) and 2018 (right).

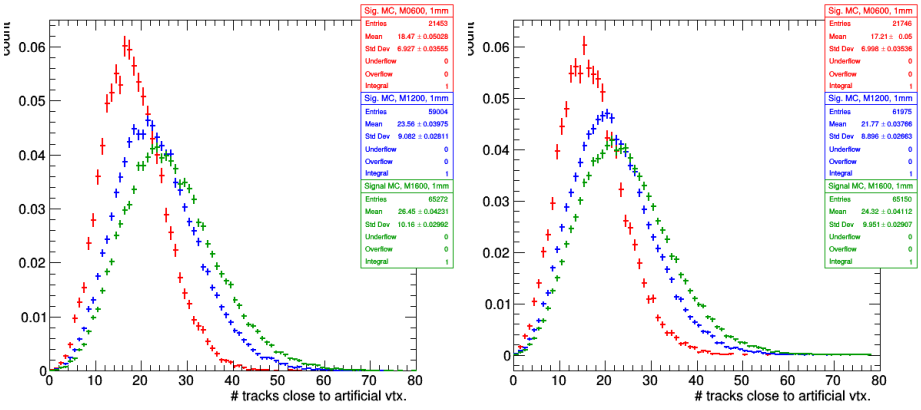

Figure 5.35: Distribution of close seed tracks near the displaced vertex in signal simulation of 1 mm lifetime for an LSP of mass 600 GeV (red), 1200 GeV (blue), and 1600 GeV (green) in 2017 (left) and 2018 (right).

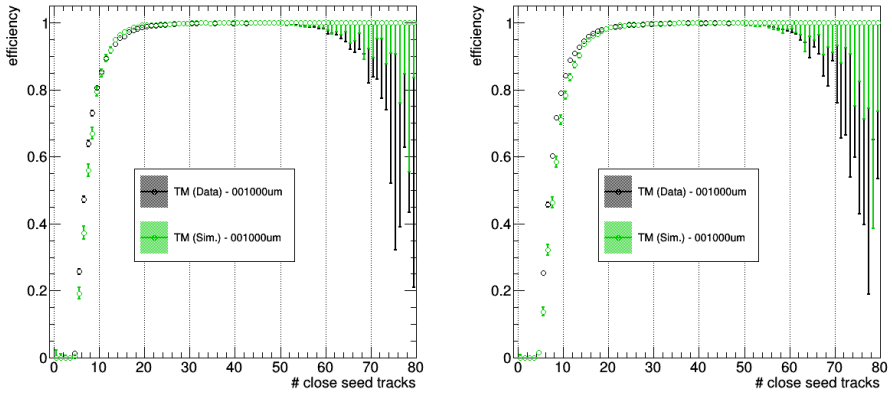

Figure 5.36: Vertex reconstruction efficiency for a 1 mm lifetime artificially displaced vertex as a function of the number of close seed tracks in data (black) and simulation (green) for 2017 (left) and 2018 (right).

namely the difference in the distributions and the vertex reconstruction efficiency turn-on curves, by modifying the corresponding attributes in the signal simulation with a set of transformations. This involves a sequence of steps that are visually summarized in Fig. 5.37 and will be explained in further detail in the following paragraphs.

First, the efficiency turn-on curve in signal simulation is shifted accordingly for each year with a reallocation of bin contents to an adjacent bin, using a linear interpolation between bins for non-integer shifts, to reflect the appropriate turn-on behavior in data compared to simulation. The result of this transformation is an effective *pseudo-data* representation of the signal simulation turn-on curve. We observe different means in the close seed track distributions for artificial vertices in data and simulation. We therefore perform a rescaling of the same distribution in signal simulation using scale factors obtained from the ratio of the data to simulation distributions, providing what we denote the *rescaled pseudo-data distribution*. Figure 5.38 shows example scale factors as a function of close seed tracks for a 1 mm lifetime artificial vertex. As the scale factors tend to be large for low numbers of close seed tracks, and rely on the tails of the close seed track distributions, an alternative transformation was also devised to cover potential

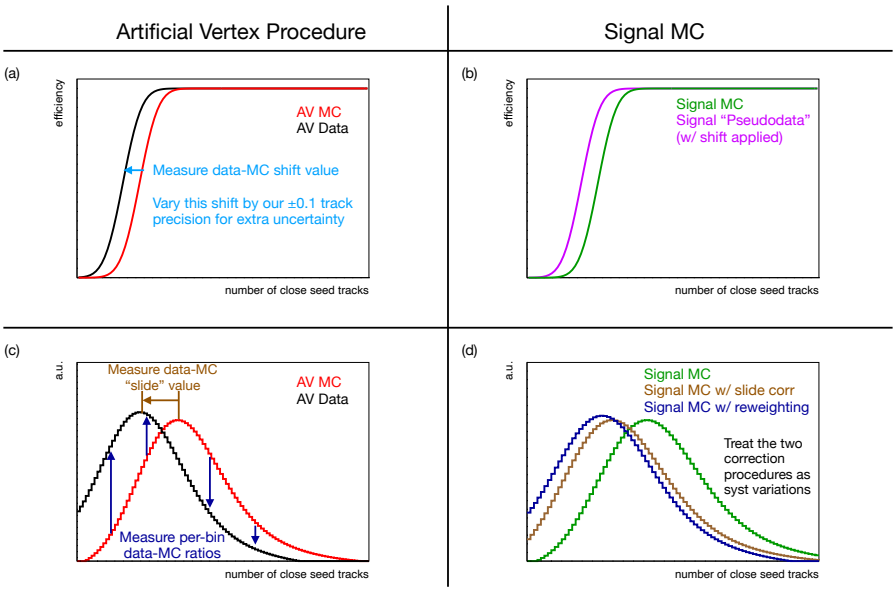

Figure 5.37: Illustrative diagram to describe the systematic uncertainty estimation procedure for the multijet signal vertex reconstruction efficiency. The artificial vertex effcency turn-on curve (top left) shift value corresponding to the difference between the data and simulation is applied to the signal turn-on curve (top right) to produce a pseudo-data turn-on curve. The differences in the close seed tracks distribution for artificial vertices in data and simulation (bottom left) provide two transformation methods to apply to the signal distribution (bottom right). These transformations are rescaling or sliding.

systematic differences between data and simulation. In this alternative transformation, we perform a shift (or *sliding* to avoid confusion with the shift of the efficiency turn on curve position) of the signal simulation close seed tracks distribution to account for the difference between the data and simulation close seed track distributions for artificial vertices, where there are 1.85 fewer tracks in 2017 data and 2.0 fewer tracks in 2018 data compared to simulation across all lifetimes. The sliding procedure is performed similar to the shifting of the efficiency turn-on curve and the resulting distribution is referred to as the *slid pseudo-data distribution*. These two transformations to the close seed tracks distribution are treated as systematic variations of the procedure. Figure 5.39 shows the rescaled and slid distributions compared to the original un-transformed signal distribution for the 1 mm lifetime samples.

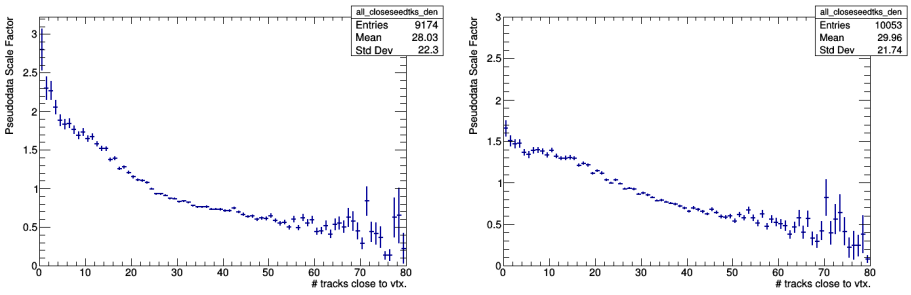

Figure 5.38: Scale factors obtained from dividing the artificial vertex close seed tracks distribution in data by the close seed tracks distribution in simulation for 2017 (left) and 2018 (right) for the 1 mm lifetime.

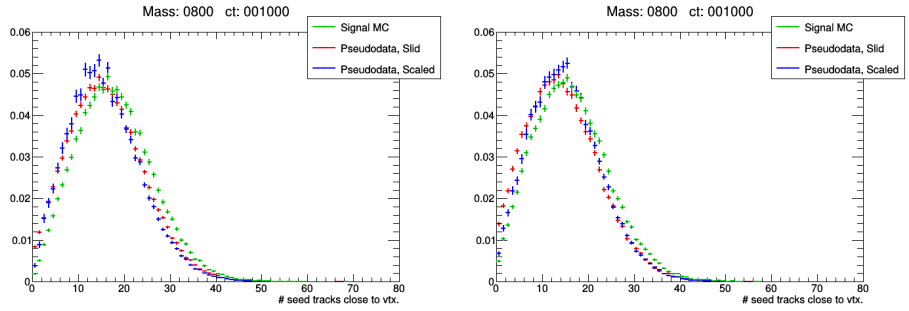

Figure 5.39: Comparison of the original 1 mm lifetime signal simulation close seed tracks distribution (green), rescaled pseudo-data distribution (blue), and the slid pseudo-data distribution (red) for 2017 (left) and 2018 (right).

Using these two variants of the pseudo-data distributions and the pseudo-data turn-on curve, we can now measure the vertex reconstruction efficiency between data and simulation by multiplying the distributions with the turn-on curve. This provides an estimate of the number of vertices successfully reconstructed, which we can determine in the pseudo-data signals as well as the unmodified signal. The ratio of the resulting vertex yield in the pseudo-data signals and the unmodified signal provides a measure of the vertex reconstruction efficiency differences between data and simulation using the two different pseudo-data represenations (rescaled and slid). The efficiencies between these two methods are averaged to obtain the final vertex reconstruction efficiency difference for multijet. On top of the uncertainty from the 100% difference between the corrected and uncorrected signal simulation efficiency, three additional uncertainties are included

when using this method: the statistical uncertainty, the location of the turn-on curve, and the difference between the rescaled and sliding methods. The 100% uncertainty due to the size of the correction factor and the difference between rescaled and sliding tend to be the dominant uncertainties for most lifetimes and masses. The final vertex reconstruction efficiency correction factors as a function of signal mass and lifetime is shown in Fig. 5.40, where the three additional uncertainties are added in quadrature and reflected in the error bars for each point. Ultimately, we measure a vertex reconstruction efficiency difference between the data and simulation to be within 0.1% and 14%, with better agreement at longer lifetimes.

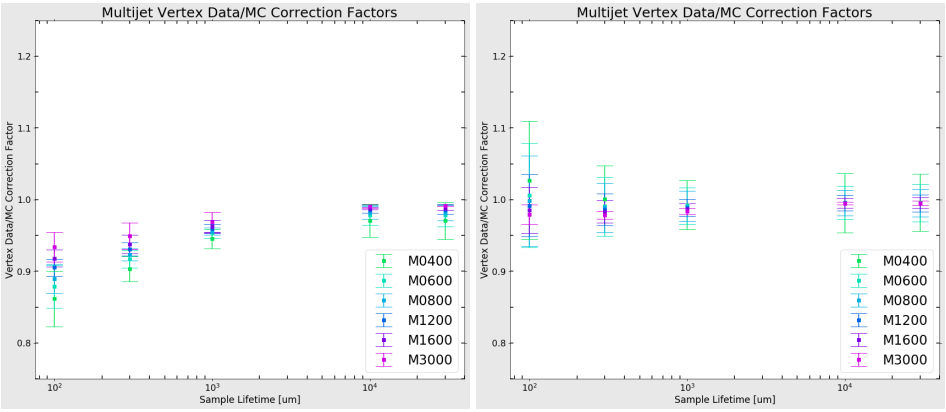

Figure 5.40: Comparison of the vertex data to simulation correction factor as a function of mass and lifetime for 2017 (left) and 2018 (right). The error bars shown correspond to the three additional uncertainties when following this procedure.

The final differences in the vertex reconstruction efficiencies between data and simulation in both multijet and dijet signals are used to correct the signal simulation yields. The correction is applied twice to account for each displaced vertex in signal. Figure 5.41 shows the signal efficiency in both multijet and dijet signals after applying all event and vertex requirements for events containing a pair of vertices with further corrections provided by the procedure described. The increase in efficiency with mass comes from the higher probability of satisfying the trigger and offline H_{T} requirements. Initially, the efficiency increases with lifetime while moving away from the prompt re-

gion, but decreases for large lifetimes because of the requirement that vertices lie within the beam pipe.

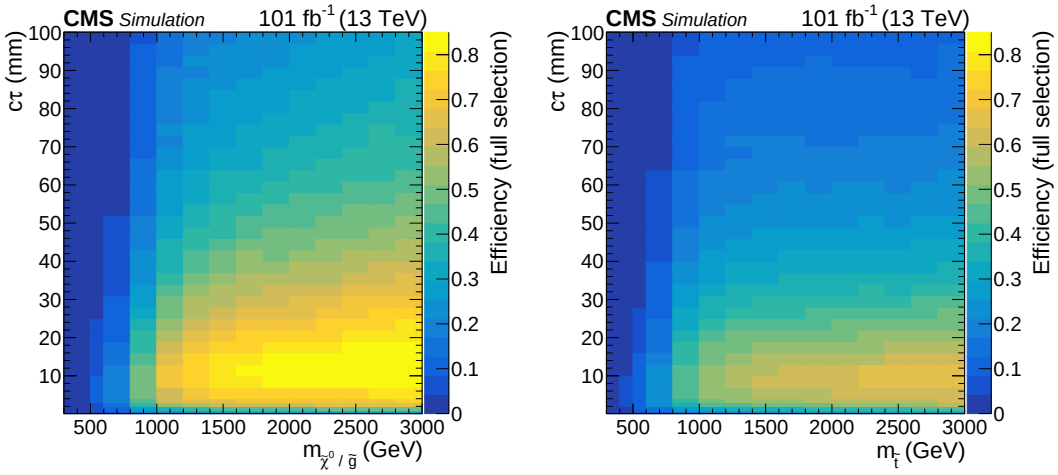

Figure 5.41: Multijet (left) and dijet (right) signal efficiencies as a function of the signal mass and lifetime for events satisfying all event and vertex requirements with corrections based on systematic differences in the vertex reconstruction efficiency between data and simulation.

5.6 Background template

In most background events, each displaced vertex arises from the inclusion of at least one misreconstructed track. This track misreconstruction endows the vertex with a displacement away from the interaction point. Background events contain two misreconstructed vertices whose displacements are independent of one another except for correlations due to events with b quarks. The independence of the two vertex displacements is a crucial feature as it offers a method with which to predict the shape of the search variable distribution, d_{VV}, in two-vertex events using information from events containing only one vertex. The constructed template, denoted as $d_{\mathrm{VV}}^{\mathrm{C}}$, provides the predicted two-vertex yields in each of the three d_{VV} bins. Events with two or more vertices are rarer than one vertex events by a factor of 100 to 1000 as shown in Table 5.7. This abundance

of one-vertex events offers greater statistical precision in the template construction.

Constructing a single value of $d_{\mathrm{VV}}^{\mathrm{C}}$ requires the following elements: two values of d_{BV}, one for each vertex, and one $\Delta\phi_{\mathrm{VV}}$ value to specify the azimuthal angle between the two vertices. Repeated sampling of d_{BV} and $\Delta\phi_{\mathrm{VV}}$ forms the preliminary shape template $d_{\mathrm{VV}}^{\mathrm{C}}$, subject to additional corrections. The sampling repeats until the total number of entries in the final $d_{\mathrm{VV}}^{\mathrm{C}}$ template is equal to 20 times the number of one-vertex events in data. The oversampling reduces the statistical uncertainty and increases the probability of adequately probing the tail of the d_{BV} distribution. The statistical uncertainty is computed for each bin, determined by the root-mean-square of yields in an ensemble of simulated pseudodata sets. These fractional statistical uncertainties are listed for each control sample and $d_{\mathrm{VV}}^{\mathrm{C}}$ bin in Tables 5.8 and 5.9. The details of the input variables to the $d_{\mathrm{VV}}^{\mathrm{C}}$ template along with corrections will be described in the paragraphs that follow.

Table 5.8: Fractional statistical uncertainties in background yield in each $d_{\mathrm{VV}}^{\mathrm{C}}$ bin arising from the limited number of one-vertex events for 2017.

Control sample	0–400 μm	400–700 μm	700–40000 μm
3-track	0.004	0.007	0.029
4-track	0.007	0.015	0.079
≥5-track	0.015	0.043	0.155

Table 5.9: Fractional statistical uncertainties in background yield in each $d_{\mathrm{VV}}^{\mathrm{C}}$ bin arising from the limited number of one-vertex events for 2018.

Control sample	0–400 μm	400–700 μm	700–40000 μm
3-track	0.004	0.007	0.030
4-track	0.008	0.016	0.083
≥5-track	0.018	0.051	0.185

The distribution of d_{BV} in ≥5-track one-vertex events is shown in Fig. 5.42, which also overlays the ≥5-track one-vertex d_{BV} distributions for simulated signal samples

of varying lifetimes. The effects of signal contamination are negligible from existing upper limits on the signal cross sections of 0.3 fb due to the much larger one-vertex background at low d_{BV}. There is a slightly higher mean d_{BV} in events with b quarks compared to events without, leading to correlations. We account for this correlation with separate templates constructed for events with and without b tags.

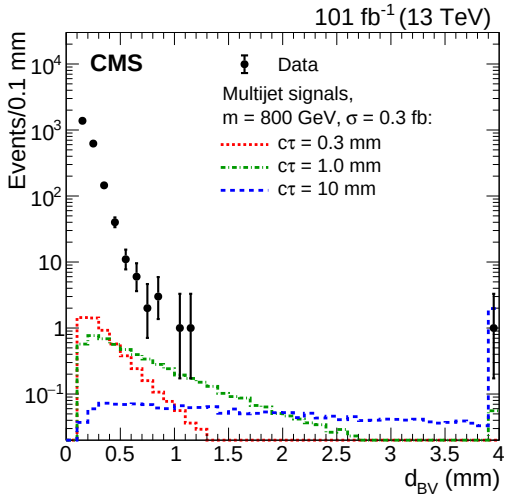

Figure 5.42: The distribution of d_{BV} for ≥5-track one-vertex events in data and simulated multijet signal samples with mass of 800 GeV, production cross section of 0.3 fb, and $c\tau$ = 0.3, 1.0, and 10 mm. All vertex and event selection criteria have been applied. The last bin includes the overflow events.

The angular separation between the vertices, $\Delta\phi_{\mathrm{VV}}$, is drawn from the distribution of angular separations between pairs of jets. The jets comprise the tracks that form the reconstructed vertices. The distribution of azimuthal angles between all possible pairs of jets in an event, denoted as $\Delta\phi_{\mathrm{JJ}}$, is close to uniform, but has small structures that are also present in the distribution of $\Delta\phi_{\mathrm{VV}}$ in events with low-track-multiplicity vertices. The $\Delta\phi_{\mathrm{JJ}}$ distribution is consistent across events containing vertices with different track multiplicities, so the $\Delta\phi_{\mathrm{JJ}}$ distribution from the high statistics 3-track one-vertex events is used to sample a $\Delta\phi_{\mathrm{VV}}$ angle for the $d_{\mathrm{VV}}^{\mathrm{C}}$ template construction.

The procedure occasionally results in small separations between the vertices. How-

ever, small $d_{\mathrm{VV}}^{\mathrm{C}}$ values are suppressed by the vertex reconstruction algorithm, which merges nearby vertices. We therefore correct the $d_{\mathrm{VV}}^{\mathrm{C}}$ template using the survival efficiency of vertex pairs as a function of their separation to capture this behavior. This efficiency is estimated in data by comparing the number of vertex pairs that remain after each merging iteration with the initial number of vertices for 3-track one-vertex events. This correction suppresses small $d_{\mathrm{VV}}^{\mathrm{C}}$ separations in the final template.

5.6.1 Template construction using b-tagged and non-b-tagged events

The production of b-quarks in pairs introduces a correlation between the vertex distances in two-vertex events that is not accounted for when single vertices are paired at random. Single b jet vertices rarely satisfy the requirement on the d_{BV} uncertainty because the narrow collimation of tracks from the b jet results in poor d_{BV} resolution. However, events with b quarks are four times more likely to have a displaced vertex than those without because the additional displaced b jet tracks satisfy the $|d_{\mathrm{xy}}|/\sigma_{d_{\mathrm{xy}}}$ requirement more often. Figure 5.43 compares the distributions of d_{BV} in one-vertex events with and without generator-level b-quarks. In simulation, while the shape of the d_{BV} distribution in events with b quarks is very similar to the shape in events without, the b quark events have larger d_{BV} on average by $39\pm2\,\mu$m in 3-track vertex events, $37\pm6\,\mu$m for 4-track vertex events, and $21\pm17\,\mu$m in ≥5-track vertex events. The fractions of events with b-quarks are consistent in 3-track, 4-track, and ≥5-track vertices: ∼50% of one-vertex events have b-quarks, and ∼79% of two-vertex events have b-quarks.

Although the b-quark correction used in the previous version of the analysis [61] was successful, there are two primary reasons to modify this part of the procedure for this analysis:

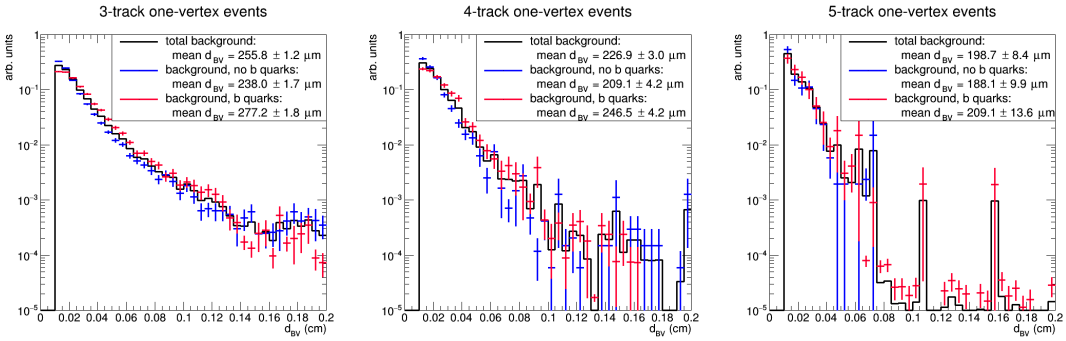

Figure 5.43: Distributions of the x-y distance from the beam axis to the vertex, d_{BV}, in one-vertex events with and without b-quarks with 2017 and 2018 MC.

1. The b-quark correction was the only part of the background template that relied entirely on simulated events.

2. The uncertainties on the b-quark correction were found to be large, particularly those uncertainties associated with the statistics available for simulated background events with ≥5-track vertices.

An alternative approach to the background template construction is considered, which relies on events classified by whether or not a b-tagged jet is present. Since this does not rely on generator-level b-quarks, this can be performed almost entirely using events in data instead of simulation. To determine whether events contain a b-tag, the "Tight" working point of the DeepJet tagger is used. Figure 5.44 demonstrates the event-level b-tagging signal efficiency and fake rate in simulated events for various working points and requirements on the number of b-tagged jets; a requirement of at least one Tight b-tag was chosen due to the high signal efficiency and background rejection that it provides.

Once events have been sorted into those with at least one Tight b-tag and those without, one can compare the d_{BV} distributions between the b-tag sorted and b-quark sorted events. Figure 5.45 shows that the events with and without a b-tagged jet reproduce

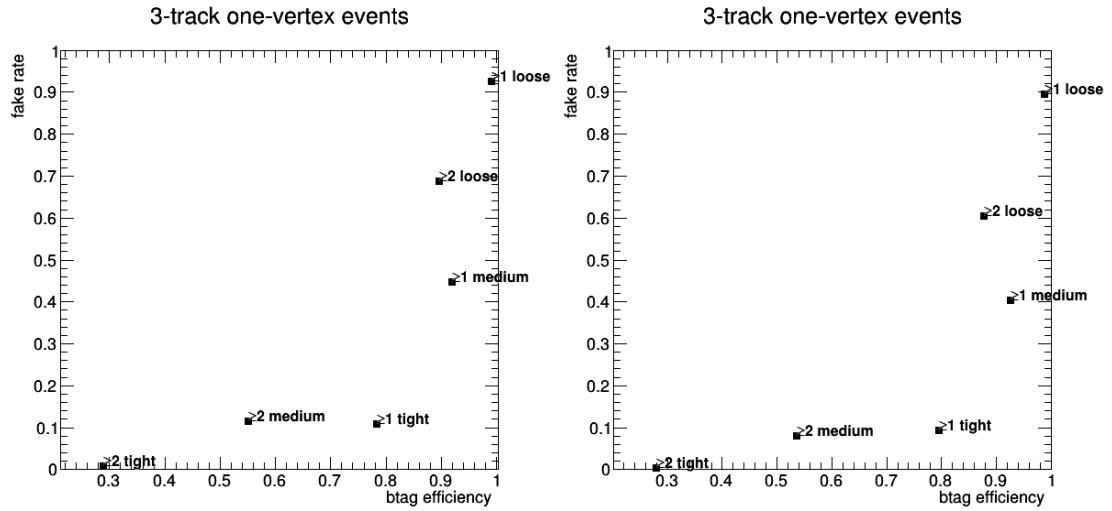

Figure 5.44: The event-level signal efficiency and fake rate for various b-tagging working points of the DeepJet tagger are shown for events containing a 3-track vertex, as measured in simulation for 2017 (left) and 2018 (right). The samples of simulated events for 2017 and 2018 include QCD and $\bar{\text{t}}\text{t}$.

the d_{BV} shape for events with and without a generator-level b-quark. As a result, the b-tag sorted events can then be directly used to construct the d_{VV}^{C} background template, after properly normalizing the distributions (as opposed to the "old" method, where the generator-level b-quark sorted events were used to determine a correction factor to apply to the background template).

A detailed writeup of the calculation of weights to combine the events with at least one b-tag and those without is provided in Appendix A. In 2017 plus 2018 simulation, this procedure was used to obtain $F_{\text{2-vertex, b-quark}} = 0.95$ for ≥5-track two-vertex events. As a result, the component of the d_{VV}^{C} template consisting of events with b-tagged jets is scaled by $F_{\text{2-vertex, b-quark}} = 0.95$ while the component consisting of events without b-tagged jets is scaled by $(1 - F_{\text{2-vertex, b-quark}}) = 0.05$. These two components are then added together to create the combined d_{VV}^{C} background template.

Separate d_{VV}^{C} templates are constructed in which the events have been sorted into two sets: events that contain at least one tagged b jet and events in which no jet was

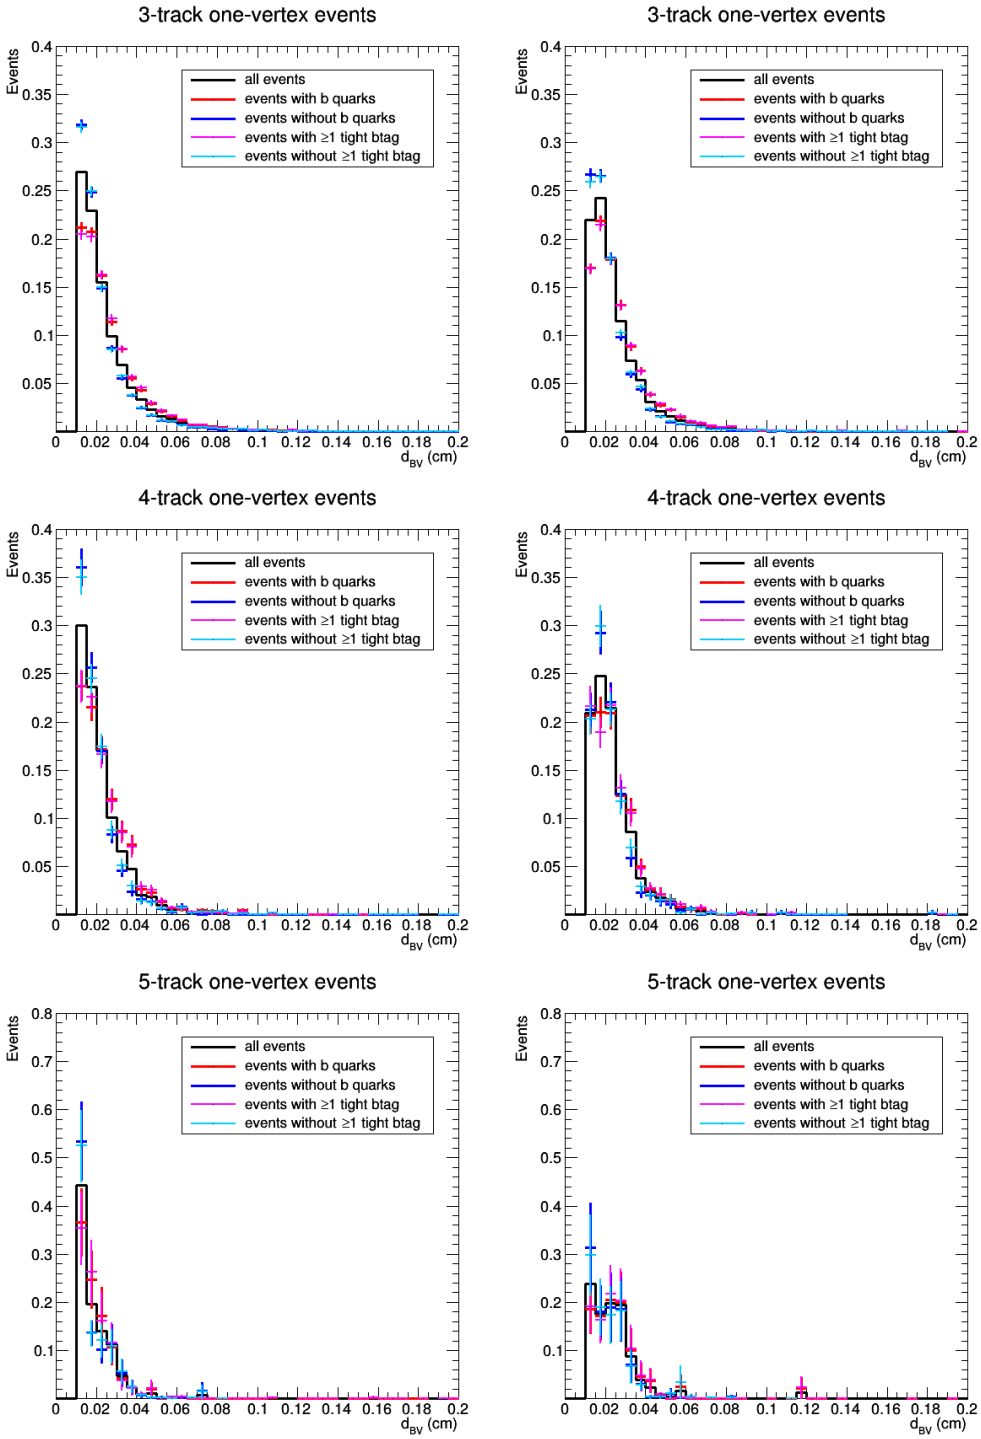

Figure 5.45: Distributions of d_{BV} for events classified by either the presence of a generator-level b-quark or the presence of at least one Tight b-tagged jet. These distributions are shown for events with one 3-track (top), 4-track (middle), or ≥5-track (bottom) vertex, as determined using simulated events for 2017 (left) and 2018 (right). The samples of simulated events for 2017 and 2018 include QCD and $\mathrm{t\bar{t}}$. Good agreement between the two methods is observed.

identified as a b jet by the tagging algorithm. These templates are combined into a single template by weighting them according to the expected fraction of two-vertex events with and without b quarks. The percentage of b quark events is determined by unfolding the b jet identification efficiencies and misidentification probabilities to relate b tagged events to b quark events and is 85% in 3-track two-vertex events, 89% in 4-track two-vertex events, and 95% in 5-track two-vertex events. This procedure leads to a 53% enhancement in the third $d_{\mathrm{VV}}^{\mathrm{C}}$ bin.

In the background-only fit, the template is normalized to the total two-vertex event yield observed in data. In situations where no two-vertex events are observed, the template is normalized using the squared vertex reconstruction efficiency for events with b quarks and events without b quarks, corrected for the survival efficiency of vertex pairs. We validate the use of the latter normalization in the 3-track two-vertex control sample in which the observed yield agrees with the predicted yield with a ratio of 1.02 ± 0.07.

Figure 5.46 compares the background templates to the observed two-vertex d_{VV} distributions. The yields in each of the three $d_{\mathrm{VV}}^{\mathrm{C}}$ bins in data are consistent with predictions from the template. With the background template constructed from data and the signal template from simulation of events with long-lived particles, we extract the signal yield from a fit of the two templates to the observed d_{VV} distribution.

Tables 5.10 and 5.11 show the predictions made by the template construction for each of the control samples in each $d_{\mathrm{VV}}^{\mathrm{C}}$ bin and the corresponding yield in the two-vertex data observed in each bin.

Since the background estimation uses one-vertex events in data to construct the template, the concern for potential signal contamination is a valid one as it could lead to overestimation of the background in the signal region and therefore weaken the back-

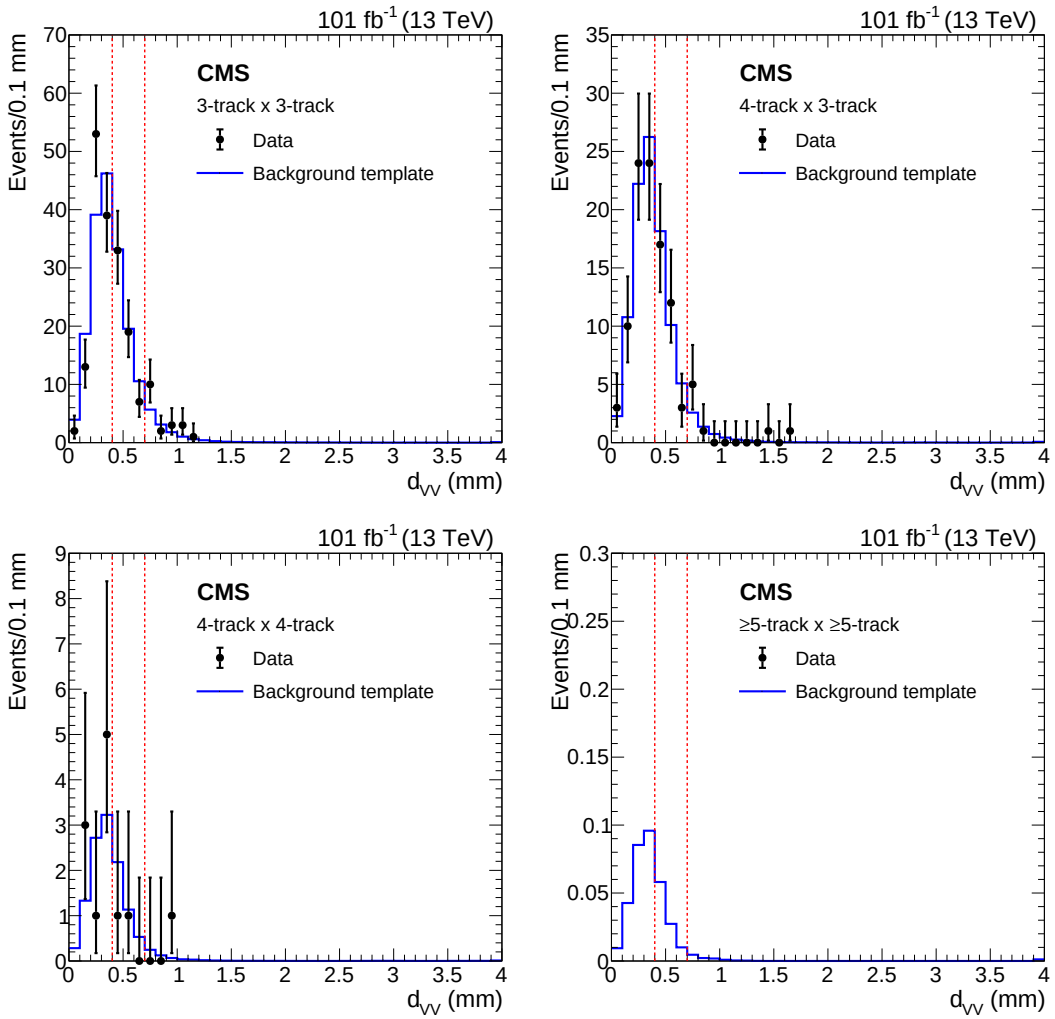

Figure 5.46: Distribution of the x-y distances between vertices, d_{VV}, for 2017 and 2018 data, overlaid on the background template $d_{\mathrm{VV}}^{\mathrm{C}}$ constructed from one-vertex events in data normalized to the two-vertex data for events with 3-track vertices (top left), events with exactly one 4-track vertex and one 3-track vertex (top right), events with 4-track vertices (bottom left), and events with 5-track vertices (bottom right). No 5-track two-vertex events were observed. The two vertical red dashed lines delineate the three d_{VV} bins.

Table 5.10: Two-vertex event prediction from background template and actual yield in 2017 data.

Control sample	0–400 μm	400–700 μm	700–40000 μm
3-track prediction	68.8 ± 0.3	36.6 ± 0.3	7.6 ± 0.2
3-track yield	67	37	9
4-tk × 3-tk prediction	37.4 ± 0.3	18.3 ± 0.3	3.4 ± 0.3
4-tk × 3-tk yield	36	20	3
4-track prediction	5.90 ± 0.04	2.67 ± 0.04	0.43 ± 0.03
4-track yield	8	1	0

Table 5.11: Two-vertex event prediction from background template and actual yield in 2018 data.

Control sample	0–400 μm	400–700 μm	700–40000 μm
3-track prediction	40.2 ± 0.2	25.8 ± 0.2	5.9 ± 0.2
3-track yield	40	22	10
4-tk × 3-tk prediction	24.5 ± 0.2	14.7 ± 0.2	2.8 ± 0.2
4-tk × 3-tk yield	25	12	5
4-track prediction	1.81 ± 0.01	1.04 ± 0.02	0.15 ± 0.01
4-track yield	1	1	1

ground template's discriminatory power from the signal templates. From the d_{BV} distribution shown in Fig. 5.26 and the yields provided in Table 5.5, while the ≥5-track one-vertex events in signal may contribute a small number vertices with large displacements, the effect of its few events is subdued by the much larger number of background events that contain vertices with much small displacements.

To validate this with a quantitative assessment of the effects from signal contamination on the background template, we injected signal into the simulated one-vertex background d_{BV} distribution and constructed the $d_{\mathrm{VV}}^{\mathrm{C}}$ template. Figure 5.47 compares the resulting $d_{\mathrm{VV}}^{\mathrm{C}}$ distributions with and without the injected signal for different life-

times. In the most signal-sensitive bin with $d_{VV}^{C} > 700\,\mu$m, the predicted fraction of events is 0.158 ± 0.001. For a 1 fb 1 mm 800 GeV signal, the fraction is 0.178 ± 0.001, with a 1 fb 10 mm 800 GeV signal, it is 0.187 ± 0.001, and with a 1 fb 10 mm 400 GeV signal, it is 0.165 ± 0.001. The magnitude of the change in predicted yield is much smaller compared to the size of the uncertainties assigned to the bin, and thus, the effect of signal contamination is negligible.

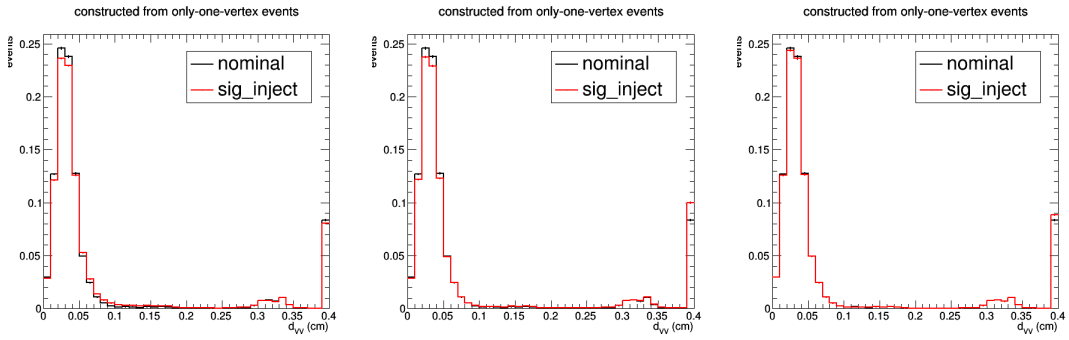

Figure 5.47: Background d_{VV}^{C} template constructed from simulated ≥5-track one-vertex background events, with and without signal injected. The injected signals are multijet signal samples with $c\tau = 1$ mm, $M = 800$ GeV, at a production cross section 1 fb (left), $c\tau = 10$ mm, $M = 800$ GeV, at a production cross section 1 fb (middle), and $c\tau = 10$ mm, $M = 400$ GeV, at a production cross section 1 fb (right).

5.7 Systematic Uncertainties

With the background template constructed from data and the signal template from simulation of events with long-lived particles, we extract the signal yield from a fit of the two templates to the d_{VV} distribution observed in the data signal region. The overall normalization of the signal and background template are free parameters of the fit under the constraint that their total integrated yield matches the yield observed in the two-vertex event data. The results obtained from the fit has some dependence on the relative yields of the three d_{VV} bins in each of the templates, but is otherwise insensitive to differences

in the finer details within the distribution. The associated systematic uncertainty in the signal and background templates are discussed here where the systematic uncertainty is evaluated for each individual bin.

5.7.1 Systematic uncertainties related to signal

Since the fit uses signal d_{VV} templates from simulation, potential differences between data and simulation give rise to systematic uncertainties. The dominant uncertainties come from vertex reconstruction efficiency and the parton distribution function (PDF) uncertainty in the simulation, with other effects such as pileup, jet energy resolution and scaling, integrated luminosity, trigger efficiency, and run conditions affecting jet efficiency providing smaller contributions.

Vertex reconstruction efficiency

We assign a systematic uncertainty equal to the size of the correction of the signal vertex reconstruction efficiency described in Sec. 5.5, along with additional uncertainties associated with variations to the procedure. The systematic uncertainty assigned for each signal point then falls within the range of 11% to 41% for dijet signals and 1% to 36% for multijet signals. In general, the greater uncertainty in dijet signals comes from their reduced efficiency due to fewer tracks from the decay point.

Impact of PDF on signal acceptance

The impact of the PDF uncertainty on the signal reconstruction efficiency is estimated by generating simulation samples reweighting 100 NNPDF replica sets [27]. Generating

100 separate sets of events with each respective PDF set is very computationally intensive and unnecessary, instead, we implement this by privately generating signal MC samples that contain 100 event weights that correctly weight the events to reproduce the distributions as if generating the event with the respective PDF replica set using the UncertaintyBands parameter in PYTHIA. We can derive the resulting signal acceptance with each of these weights. While we could take the RMS deviation of these acceptances for the uncertainty, this is sensitive to outliers and instead, we order the 100 values for the acceptances A in ascending order such that:

$$A^{(1)} \leq A^{(2)} \leq ... \leq A^{(99)} \leq A^{(100)} \tag{5.2}$$

The relative uncertainty is obtained from the 68% interval of the signal acceptance of these 100 replica sets:

$$\frac{A^{(84)} - A^{(16)}}{2} \tag{5.3}$$

this is compared to the value at the midpoint of the interval:

$$\frac{A^{(84)} + A^{(16)}}{2} \tag{5.4}$$

Thus, we obtain the following for the relative uncertainty in the acceptance:

$$\frac{A^{(84)} - A^{(16)}}{A^{(84)} + A^{(16)}} \tag{5.5}$$

Figure 5.48 shows the distribution of this value as a function of mass for three different signal lifetimes. They uncertainty ranges between 1% and 8%, depending primarily

on the signal point mass due to the underlying uncertainty in the parton luminosities that vary with the mass of the particle [20]. The drawn red line corresponds to the corresponding uncertainty values used in the limit setting at the shown lifetimes; for other lifetimes, linear interpolation between the drawn curves is used. The uncertainties correspond to the statistical uncertainty observed in the original acceptance.

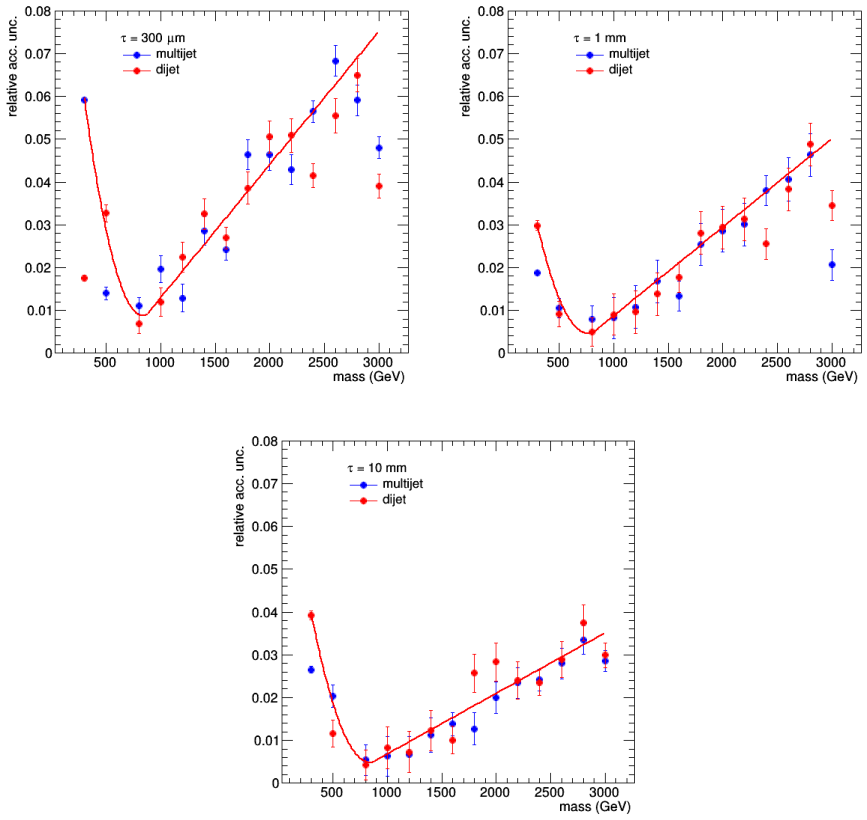

Figure 5.48: Relative uncertainty in acceptance as a function of signal mass for 300 μm (top left), 1 mm (top right), and 10 mm (bottom) signal lifetimes.

Jet energy scaling and resolution

As described in Sec. 4.4, various corrections are needed to stabilize the non-linear detector response to jets and account for pileup contributions or noise to provide an accurate measurement of the jet energy. Variations in these corrections would change the scaling

and resolution of the jet energies. Since this analysis uses jets at the initial trigger and event preselection stage, differences in jet energy measurements could affect whether an event passes or fails the requirement. To evaluate the systematic uncertainty due to the jet energy corrections, we vary the jet energy scale and resolution both upward and downward and evaluate the resulting signal efficiency with the modified jet energies. The uncertainty is taken as the difference from unity of the ratio of the signal efficiency in the variation to the nominal signal efficiency. Figure 5.49 shows the effects on the signal efficiency due to the jet energy scale and resolution as a function of signal lifetime and mass. The effect due to variations in the jet energy scale is ~5% or less for all signal samples, and the effect due to jet energy resolution is ~2% or less.

Pileup

To study signal efficiency sensitivity to the pileup distribution, we vary the pp cross section for minimum bias events used in the pileup weighting by ±5%. We take the ratios of the resulting signal efficiencies to the nominal signal efficiency, and average the magnitudes of the differences of the ratios from 1 for the + and − variations. Figure 5.50 shows the effect on the signal efficiency due to pileup as a function of signal lifetime and mass. The effect is ~2% or less for all signal samples, we assign this 2% as a systematic uncertainty in the signal efficiency.

Issues in Run 2 data

Unforeseen circumstances that occurred during the Run 2 data collection period affected parts of the data set, which can subsequently affect signal efficiencies that are quantified here.

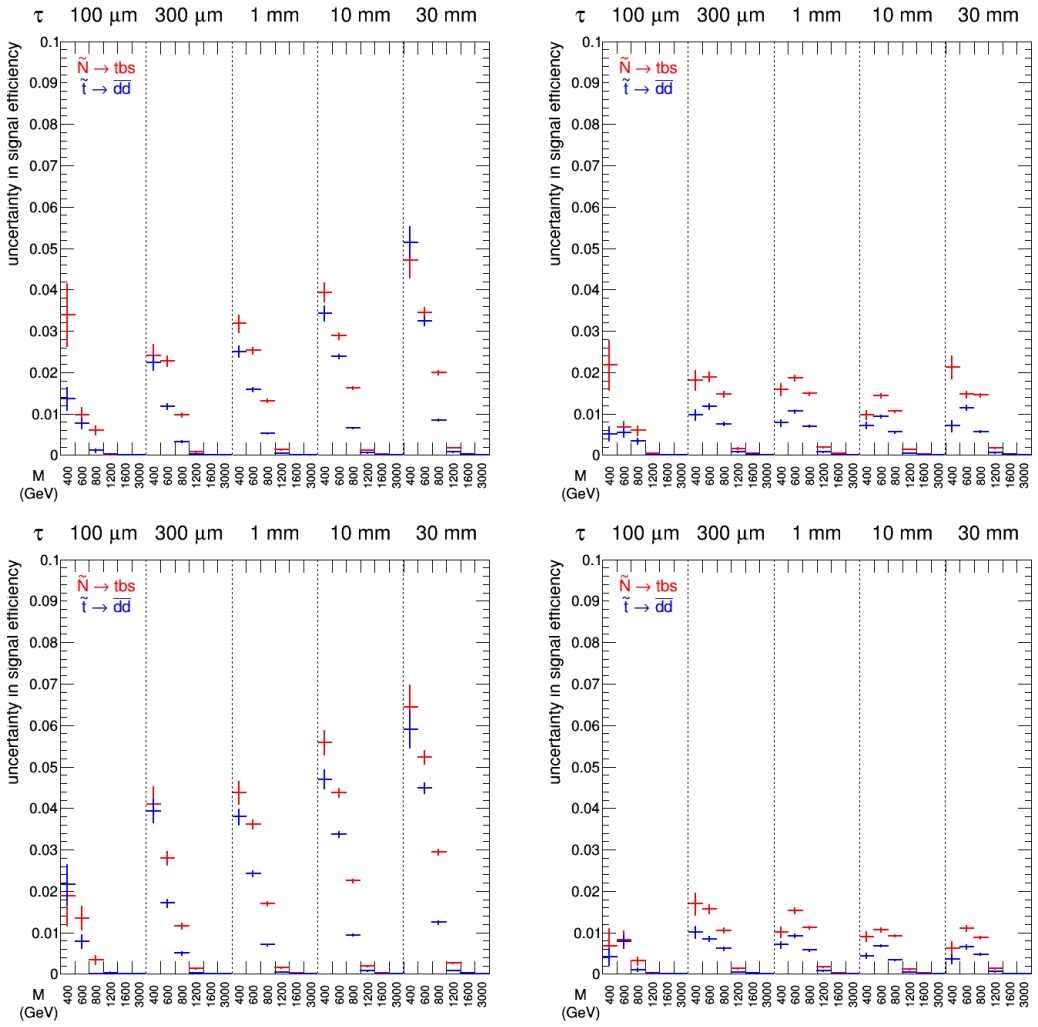

Figure 5.49: Fractional change in signal efficiency due to variations in the jet energy scale (left) and the jet energy resolution (right), as a function of signal mass and lifetime for 2017 signal samples (top) and 2018 signal samples (bottom).

In 2016 and 2017, a mistiming in the ECAL endcaps at the L1 trigger level resulted in the improper assignment of trigger primitives to the previous bunch crossing, which could lead to unintentional event vetos due to trigger rules. We evaluate the potential effect of this issue on the signal MC by removing events containing jets that have both $p_{\mathrm{T}} > 100\,\mathrm{GeV}$ and $2.25 < \eta < 3.0$. Figure 5.51 plots the ratio of signal efficiencies following this recipe to the nominal efficiencies as a function of signal mass and lifetime for the two benchmark signal models.

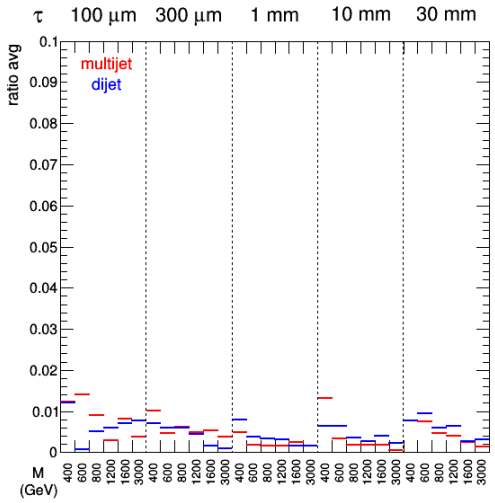

Figure 5.50: Fractional change in signal efficiency due to a variation in the pileup distribution, as a function of signal mass and lifetime.

In 2018, a power interruption triggered by a fire alarm resulted in inoperable sectors in the negative endcap of the HCAL, namely HEM15 and HEM16, resulting in effectively a 40 degree section ($-3.0 < \eta < -1.3$ and $-1.57 < \phi < -0.87$) of the HCAL turned off. This affects around 64% of the 2018 dataset. Similar to the treatment of the ECAL issue, we evaluate the effects of this incident on the signal efficiency by omitting jets in the affected ϕ and η region and discarding any events that subsequently fail the event preselection requirements on the number of jets in the event or H_{T}. Figure 5.52 shows the ratio of signal efficiencies following this recipe to the nominal efficiences as a function of signal mass and lifetime for the two benchmark signal models.

The signal efficiency is rescaled to account for these issues. Additionally, the differences in efficiencies is around 2% for each of these effects, but since they only affect a subset of the full dataset themselves, we will assign a flat 1% uncertainty for each issue.

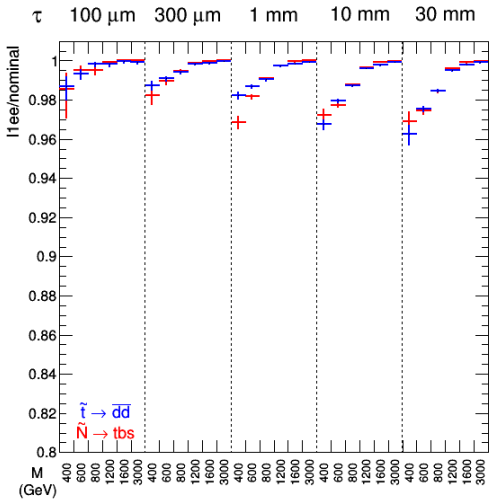

Figure 5.51: Ratio of L1EE prefiring-affected efficiences to the nominal efficiencies as a function of signal mass and lifetime for dijet and multijet signal models.

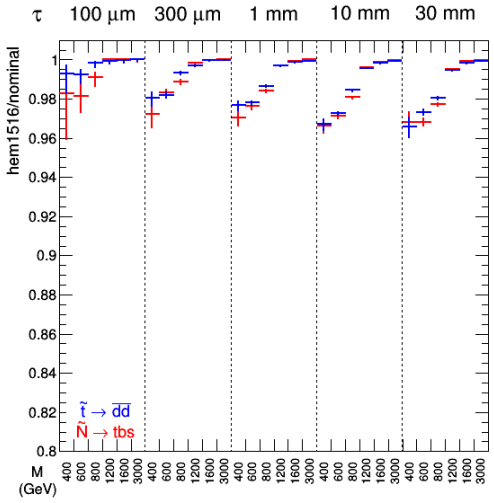

Figure 5.52: Ratio of HEM15/16-affected efficiences to the nominal efficiencies as a function of signal mass and lifetime for dijet and multijet signal models.

Trigger efficiency

From the trigger efficiency study presented in Sec. 5.2.1, the ratio of the measured trigger efficiency between data and simulation is 0.99 for $H_{\mathrm{T}} > 1200\,\mathrm{GeV}$. We scale the signal efficiency by this ratio and assign a related uncertainty of 1%.

Instantaneous luminosity

The uncertainty in the integrated luminosity is 2.3% in 2017 [4] and 2.5% in 2018 [5].

Summary of signal systematic uncertainty

Table 5.12 summarizes the systematic uncertainties related to the signal models. We assume no correlations between the different contributions and obtain an overall systematic uncertainty by adding each value in quadrature.

Table 5.12: Systematic uncertainties in signal for dijet and multijet signal models. The overall uncertainty is the sum in quadrature of the individual uncertainties.

Systematic effect	Dijet uncertainty (%)	Multijet uncertainty (%)
Vertex reconstruction	11-41	1-36
PDF uncertainty	1-8	1-8
Integrated luminosity	2.3-2.5	2.3-2.5
Jet energy scale	5	5
Jet energy resolution	2	2
Pileup	2	2
Trigger efficiency	1	1
Changes in run conditions	1	1
Overall	13-42	7-36

5.7.2 Systematic uncertainties in background templates

Systematic uncertainties in the background template come from effects that modify the shape of the constructed $d_{\mathrm{VV}}^{\mathrm{C}}$ distribution away from the shape of the true two-vertex d_{VV} distribution. The 3-track vertex control sample provides a way to reliably assess these differences with statistical precision. Thus, within each of the three bins in the $d_{\mathrm{VV}}^{\mathrm{C}}$ template in the 3-track vertex control sample, we evaluate the ratio of the yield predicted by the template to the true observed two-vertex yield in data, referred to as the closure, and take the deviation from unity as a measure of the systematic uncertainty for each $d_{\mathrm{VV}}^{\mathrm{C}}$ bin. We find the $d_{\mathrm{VV}}/d_{\mathrm{VV}}^{\mathrm{C}}$ ratio is 0.99 ± 0.10 in the 0-0.4 mm bin, 0.93 ± 0.12 in the 0.4-0.7 mm bin, and 1.38 ± 0.32 in the 0.7-40 mm bin. Additional systematic uncertainties are obtained by measuring the difference from unity of the ratio of template yields in each of the $d_{\mathrm{VV}}^{\mathrm{C}}$ bins for variations in the template construction compared to the nominal template. More concretely, since the 3-track control sample is used to validate the ≥5-track template with higher statistical precision, variations in the input parameters of the template are taken to capture potential differences between 3-track events and ≥5-track events within each bin. An uncertainty in the ≥5-track template normalization is also computed for the estimation of the absolute yield of background events in ≥5-track two-vertex events.

Vertex pair survival efficiency

The template shape is sensitive to the vertex pair survival efficiency correction, which uses the d_{VV}-dependent efficiency for vertex pairs to survive. The efficiency curve is extracted directly from the reconstruction algorithm, which counts every pair of vertices and their separation distances at each iteration to assess the number of merges and pair survivals. Under normal circumstances, the vertex reconstruction algorithm initiates the

iterative merging starting from seed vertices formed from all combinatorial track doublets. To vary this procedure and derive an alternative efficiency curve, we consider seed vertices formed from all possible combinatorial track quintuplets, i.e. five tracks. The reconstruction algorithm then proceeds normally with the iterative merging of these high-track-multiplicity seed vertices with shared tracks. In this way, the variation attempts to capture the scenario being modeled in the background template construction and determine the rate at which two fully formed vertices are merged or kept as separate vertices as a function of their separation distance. We construct the $d_{\mathrm{VV}}^{\mathrm{C}}$ template with the resulting efficiency curve using this variation and evaluate the fractional change per bin of the template and assign this as the systematic uncertainty. Figure 5.53 shows the efficiency curves from the two methods along with the corresponding ratio in the predicted yields for each of the $d_{\mathrm{VV}}^{\mathrm{C}}$ bins.

Systematic uncertainty from template normalization

The normalization of the background template was calculated following the same principle as the template itself. Thus, the same variations are taken to assess the sensitivity of the normalization factor to determine a systematic uncertainty. The dominant contributor driving the size of this uncertainty is the vertex pair survival efficiency correction, modified to provide an absolute efficiency. This systematic uncertainty is assigned equally to all three bins. We obtain the systematic uncertainties in Table 5.13. The overall uncertainties listed will be applied to each $d_{\mathrm{VV}}^{\mathrm{C}}$ bin in the final template.

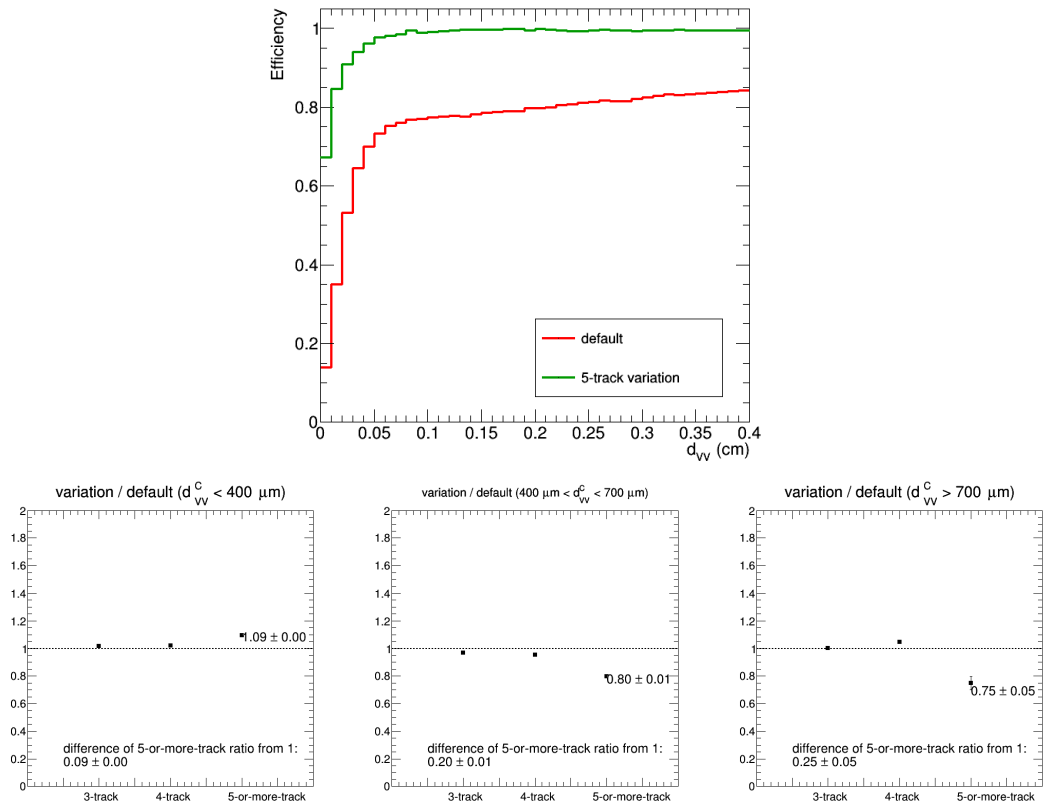

Figure 5.53: Distributions of the efficiency to keep pairs of vertices as a function of d_{VV}, from two methods with 2017 and 2018 data (top), ratios of simulated yields in the region $d^{\mathrm{C}}_{\mathrm{VV}} < 400\,\mu$m (bottom left), ratios of simulated yields in the region $400 < d^{\mathrm{C}}_{\mathrm{VV}} < 700\,\mu$m (bottom middle), ratios of simulated yields in the region $d^{\mathrm{C}}_{\mathrm{VV}} > 700\,\mu$m (bottom right).

Table 5.13: Fractional systematic shifts in the background template normalization factor for both 2017 and 2018. The overall systematic uncertainty is the sum in quadrature of the shifts assuming no correlations among the sources.

Systematic effect	2017	2018
Closure in 3-track control sample	0.02	0.01
Modeling of vertex survival efficiency	0.23	0.21
Variation of b-tag scale factors	0.04	0.03
Overall	0.24	0.21

Distribution of azimuthal angle between vertices

We must also test the assumption that the closure in 3-track events implies closure in ≥5-track events. When constructing the background template, the angular separation between vertices $\Delta\phi_{\mathrm{VV}}$ is modeled from the $\Delta\phi_{\mathrm{JJ}}$ distribution in 3-track vertices. The $\Delta\phi_{\mathrm{JJ}}$ distributions in ≥5-track one-vertex events and 3-track one-vertex events are consistent. This does not exclude differences in the angles between jets and vertices. To gauge this effect, we construct the template by sampling the $\Delta\phi_{\mathrm{VV}}$ value from a uniform distribution. The fractional change of the resulting template from the nominal template in each $d^{\mathrm{C}}_{\mathrm{VV}}$ bin is taken as the systematic uncertainty. Figure 5.54 shows the resulting $\Delta\phi_{\mathrm{VV}}$ distributions when drawing an angle from the two separate distributions along with the corresponding ratio in the predicted yields for each of the $d^{\mathrm{C}}_{\mathrm{VV}}$ bins.

Systematic uncertainties on the template construction using b-tagged and non-b-tagged events

The b tag efficiencies and fake rates are determined using simulated events in the phase space relevant to this analysis, and efficiency correction factors are applied to match the efficiencies and fake rates in data. We vary these based on measurements of the p_{T}-dependent b tagging efficiency [60] and take the fractional change of the resulting template as the systematic uncertainty. Similarly, we vary the b quark fraction in ≥5-track vertex events within the ranges observed in 3-track and 4-track vertex events, assigning the systematic uncertainty as the fractional change in the resulting template.

Varying the b-tagging data-to-simulation scale factors

The data-to-simulation scale factors ("SF") and their variations ("var") are provided

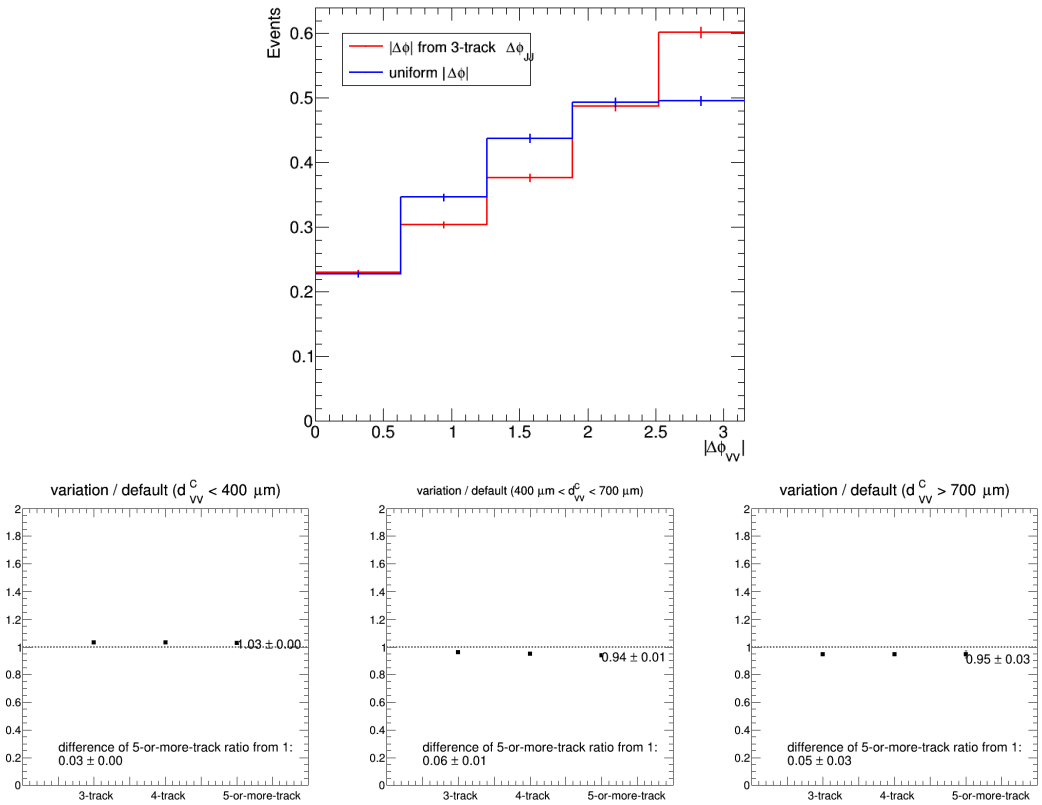

Figure 5.54: Constructed $\Delta\phi_{\mathrm{VV}}$ distribution using a variation in the $\Delta\phi_{\mathrm{VV}}$ input for 2017 and 2018 data (top), ratios of simulated yields in the region $d^{\mathrm{C}}_{\mathrm{VV}} < 400\,\mu$m (bottom left), ratios of simulated yields in the region $400 < d^{\mathrm{C}}_{\mathrm{VV}} < 700\,\mu$m (bottom middle), ratios of simulated yields in the region $d^{\mathrm{C}}_{\mathrm{VV}} > 700\,\mu$m (bottom right).

by the BTV POG in Refs. [**?**, **?**]. These scale factor variations are applied to jets as:

- SF ± var for b-jets and c-jets;
- SF × (1 ± var) for light jets.

As described in Refs. [**?**, **?**], the b-jets and c-jets are treated as fully correlated with one another (i.e. they are varied up or down simultaneously) while light jets are treated as uncorrelated with the b/c-jets (and thus the resulting light jet uncertainty is added in quadrature with that associated with the variations for b/c-jets).

A flat variation per jet flavor is used in this analysis, which is chosen to be 5% for b-jets, 18% for c-jets, and 28% for light jets. These values correspond to the maximum variations for b-jets and c-jets with 30 GeV $< p_{\mathrm{T}} <$ 600 GeV and the maximum variation for all light jets. While slightly larger variations do exist for b/c-jets outside of this p_{T} range, the majority of b/c-jets relevant for the analysis fall within this range, so this is still fairly conservative. The resulting systematic shifts from this approach are shown in Table 5.14 for 2017 plus 2018 simulation, which demonstrates that the b-tagging scale factor variations introduce uncertainties of a few percent, and only at high d_{VV}.

Table 5.14: Fractional systematic shifts in the 2017 plus 2018 MC background prediction in each $d^{\mathrm{C}}_{\mathrm{VV}}$ bin arising from varying the b-tagging data-to-simulation scale factors. The shift values and their statistical uncertainties are shown.

	d_{VV} range		
Systematic effect	0–400 μm	400–700 μm	700–40000 μm
Vary b/c-jet scale factors down	0.00 ± 0.00	0.00 ± 0.00	+0.01 ± 0.03
Vary b/c-jet scale factors up	0.00 ± 0.00	0.00 ± 0.00	−0.01 ± 0.03
Vary light jet scale factors down	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
Vary light jet scale factors up	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.01

Uncertainties due to variations on the ≥5-track two-vertex b-quark fraction

The two-vertex b-quark fraction ($F_{\text{2-vertex, b-quark}}$), is used to combine the b-tag sorted $d^{\mathrm{C}}_{\mathrm{VV}}$ templates from data, resulting in the final background template. However, $F_{\text{2-vertex, b-quark}}$ varies with track multiplicity—the $F_{\text{2-vertex, b-quark}}$ values measured using 2017 plus 2018 simulation are found to be 85% for 3-track events, 88% for 4-track events, and 95% for ≥5-track events. In other words, the fraction of two-vertex events containing a b-quark is larger for secondary vertices containing five tracks than for vertices containing only three tracks. As the ≥5-track two-vertex sample is used as the signal region in the analysis while the 3-track and 4-track two-vertex samples are used to assess e.g. closure tests, it is important to assign an uncertainty which can cover potential differences between

the closure region and the signal region used in the analysis.

One very conservative approach would be to vary $F_{\text{2-vertex, b-quark}}$ for ≥5-track events down to the 3-track value, and vary it up by a similar amount in the opposite direction. In other words:

$$\frac{F_{\text{2-vertex, b-quark}}(\geq\text{5-track})}{F_{\text{2-vertex, b-quark}}(\text{3-track})} = 1 + x$$
$$\implies \text{use } F_{\text{2-vertex, b-quark}}(\geq\text{5-track}) = \min\{F_{\text{2-vertex, b-quark}}(\text{3-track}) \times (1 + x \pm x), 1\} \tag{5.6}$$

However, this is an extremely conservative variation, as there is a clear trend for $F_{\text{2-vertex, b-quark}}$ to increase with increasing track multiplicity, and it is therefore unlikely that the $F_{\text{2-vertex, b-quark}}$ value from 3-track events is actually correct for the ≥5-track events. Moreover, there is data among both 3-track two-vertex and 4-track two-vertex events which can be used to assess closure. This provides confidence in using a slightly less extreme variation using the 4-track $F_{\text{2-vertex, b-quark}}$ value, though this is still quite conservative:

$$\frac{F_{\text{2-vertex, b-quark}}(\geq\text{5-track})}{F_{\text{2-vertex, b-quark}}(\text{4-track})} = 1 + x$$
$$\implies \text{use } F_{\text{2-vertex, b-quark}}(\geq\text{5-track}) = \min\{F_{\text{2-vertex, b-quark}}(\text{4-track}) \times (1 + x \pm x), 1\} \tag{5.7}$$

Since $F_{\text{2-vertex, b-quark}}(\text{4-track}) = 0.88$ and $F_{\text{2-vertex, b-quark}}(\geq\text{5-track}) = 0.95$ in 2017 plus 2018 simulation, $x = 0.08$, i.e. $F_{\text{2-vertex, b-quark}}(\geq\text{5-track})$ is varied down to 0.88 and up to $\min\{1.02, 1\} = 1$ (since the fraction of events containing a b-quark can never be more than 100%). Table 5.15 demonstrates the impact of this variation, where the largest uncertainty is found to be about 11% in the largest d_{VV} bin. Adding this in quadrature with the uncertainties in Table 5.14, the total systematic uncertainties on this method are found to be about 2% in the first d_{VV} bin, 1% in the second d_{VV} bin, and 11% in the last d_{VV} bin, which is a huge improvement compared to those seen in Ref. [61] for the old method.

Table 5.15: Fractional systematic shifts in the 2017 plus 2018 MC background prediction in each d_{VV}^{C} bin arising from varying the two-vertex b-quark fractions. The shift values and their statistical uncertainties are shown.

	d_{VV} range		
Systematic effect	0–400 μm	400–700 μm	700–40000 μm
Vary $F_{\text{2-vertex, b-quark}}$ down	$+0.02 \pm 0.01$	0.00 ± 0.00	-0.06 ± 0.09
Vary $F_{\text{2-vertex, b-quark}}$ up	-0.02 ± 0.01	0.00 ± 0.00	$+0.04 \pm 0.08$

Pileup

The dependence on pileup of the background template is studied by sorting the simulated one-vertex events into bins of the true number of pileup interactions. A separate d_{VV}^{C} template is constructed for each bin of number of pileup interactions, which would enhance any systematic dependence or correlation that could be introduced at varying levels of pileup. These templates are then combined into a single template. Figure 5.55 shows the d_{VV}^{C} distributions obtained following this procedure, compared with the default construction. The differences in predicted yields in each d_{VV}^{C} bin are less than 1%, so we take the effect to be negligible.

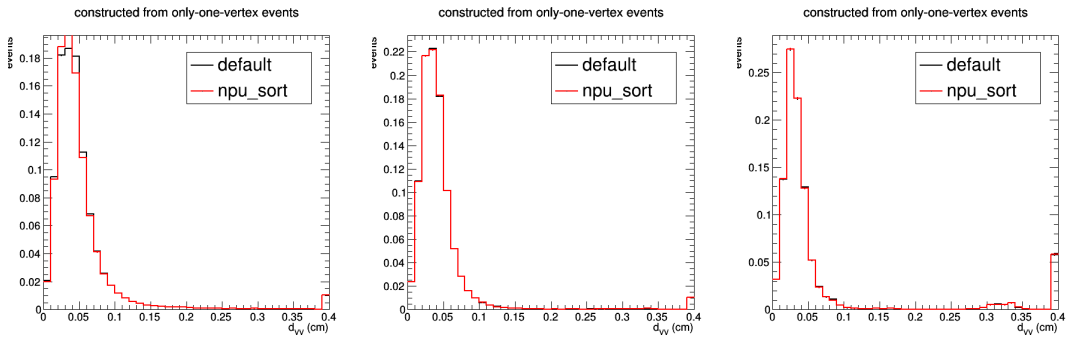

Figure 5.55: Background template, d_{VV}^{C}, constructed using the default method (black) and sorting by the true number of pileup interactions (red), for 3-track (left), 4-track (middle), and ≥5-track (right) vertices.

Table 5.16 and 5.17 summarizes the systematic uncertainty for each of these components among the three d_{VV} bins for 2017 and 2018, respectively. We assume no correla-

tions between these different effects and add all values and their uncertainties in quadrature to obtain the overall systematic uncertainty in each bin. The limits are computed assuming the first bin is anti-correlated with the second and third bins in the background systematic uncertainty and each bin is fully correlated across the different years while the statistical components of each bin are assumed to be uncorrelated.

Table 5.16: Fractional systematic shifts in the background prediction in each d_{VV}^{C} bin arising from varying the construction of the d_{VV}^{C} template for 2017. The overall systematic uncertainty is the sum in quadrature of the shifts assuming no correlations among the sources.

	d_{VV} range		
Systematic effect	0–400 μm	400–700 μm	700–40000 μm
Closure in 3-track control sample	0.03	0.01	0.18
Difference from 3-track vertices to 5-track vertices:			
Modeling of $\Delta\phi_{VV}$	0.02	0.06	0.05
Modeling of vertex survival efficiency	0.09	0.23	0.24
Variation of b-tag fraction	0.02	0.04	0.05
Variation of b-tag scale factors	0.00	0.01	0.01
5-track template normalization factor	0.24	0.24	0.24
Overall	0.26	0.34	0.39

Table 5.17: Fractional systematic shifts in the background prediction in each d_{VV}^{C} bin arising from varying the construction of the d_{VV}^{C} template for 2018. The overall systematic uncertainty is the sum in quadrature of the shifts assuming no correlations among the sources.

	d_{VV} range		
Systematic effect	0–400 μm	400–700 μm	700–40000 μm
Closure in 3-track control sample	0.01	0.15	0.68
Difference from 3-track vertices to 5-track vertices:			
Modeling of $\Delta\phi_{VV}$	0.03	0.06	0.03
Modeling of vertex survival efficiency	0.10	0.17	0.26
Variation of b-tag fraction	0.02	0.03	0.05
Variation of b-tag scale factors	0.00	0.00	0.01
5-track template normalization factor	0.21	0.21	0.21
Overall	0.23	0.32	0.76

5.8 Results and statistical interpretation

Table 5.18 summarizes the predicted 5-track two-vertex event yields in each of the three d_{VV} bins from the background and signal templates for three multijet signal lifetime points as well as the observation in data. No 5-track two-vertex events were observed in the 2017 and 2018 data.

Table 5.18: Predicted yields for the background-only normalized template, the predicted yields for multijet signals with mass of 800 GeV, production cross section of 0.3 fb, and $c\tau = 0.3, 1.0$, and 10 mm, and the observed yield in each d_{VV} bin. The uncertainty in the signal yields reflect the systematic uncertainties given in Table 5.12. In the background prediction, the first uncertainty corresponds to the statistical uncertainty given in Tables 5.8 and 5.9 while the second reflects the systematic uncertainty given in Tables 5.16 and 5.17.

		Predicted multijet signal yields			
d_{VV} range	Predicted background yield	0.3 mm	1.0 mm	10 mm	Observed
0–0.4 mm	$0.235 \pm 0.003 \pm 0.059$	0.7 ± 0.2	0.7 ± 0.1	0.20 ± 0.02	0
0.4–0.7 mm	$0.096 \pm 0.003 \pm 0.031$	0.8 ± 0.2	1.1 ± 0.2	0.10 ± 0.01	0
0.7–40 mm	$0.011 \pm 0.001 \pm 0.006$	0.8 ± 0.2	5.4 ± 0.9	12 ± 1	0

To extract the signal yield from the data, we perform a binned shape fit using extended maximum likelihood with three d_{VV} bins. Signal d_{VV} templates come directly from simulation with a template for each signal model, mass, and lifetime point. The background d_{VV}^{C} template is constructed from the one-vertex events in data. The overall normalizations of the signal and background templates are free parameters of the fit under the constraint that their total integrated yield matches the yield observed in the two-vertex event data. The results obtained from the fit depend on the relative yields in the three d_{VV} bins and their systematic uncertainties where the 2017 and 2018 datasets are treated independently and combined in the fit.

The upper limits on the signal cross section are determined by first assuming a uni-

form Bayesian prior for the cross section. For each signal mass and lifetime point, the signal efficiency is constrained by a log-normal prior with a corresponding width as determined from the overall systematic uncertainty in signal as summarized in Table 5.12. The shape uncertainty in the signal template arises from the statistical uncertainty of the simulation. For the background template, a log-normal prior is taken for each $d_{\mathrm{VV}}^{\mathrm{C}}$ bin for each dataset year with widths specified in Table 5.16 and 5.17 for 2017 and 2018, respectively.

The final fit combines this dataset together with limits set in 2015 and 2016 to achieve the full Run 2 result. The correlation between these datasets are treated the same as the correlation between the 2017 and 2018 datasets. Figure 5.56 shows the full Run 2 observed 95% confidence level upper limits on the product of the pair-production cross section with the square of the branching fraction ($\sigma\mathcal{B}^2$), as a function of mass and mean proper decay length. As a reminder from Sec. 2.4, the exclusion curves overlaid assume the gluino and top squark pair production cross sections at $\mathrm{NNLO}_{\mathrm{approx}}$+NNLL precision [25, 23, 24]. For the neutralino, the production cross sections are computed at NLO+NLL precision in a limit of mass-degenerate higgsino states $\tilde{\chi}_1^{\pm}$, $\tilde{\chi}_1^0$, and $\tilde{\chi}_2^0$ with all the other sparticles assumed to be heavy and decoupled [45, 46]. For all models, we assume a 100% branching fraction to the specified decay mode.

For the long-lived gluino, neutralino, and top squark in the RPV models described, pair production cross sections larger than 0.08 fb are excluded for masses between 800 and 3000 GeV and mean proper decay lengths between 1 and 25 mm. For mean proper decay lengths between 0.6 and 90 mm, the data exclude gluino masses up to 2500 GeV; for mean proper decay lengths between 0.6 and 70 mm, the data exclude neutralino masses up to 1100 GeV; and for mean proper decay lengths between 0.4 and 80 mm, the data exclude top squark masses up to 1600 GeV. These are the most stringent bounds on

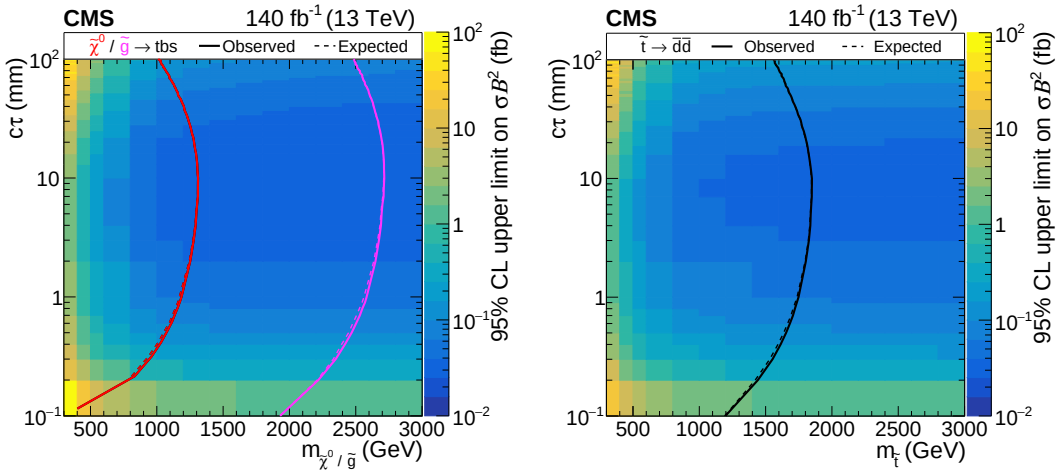

Figure 5.56: Observed 95% CL upper limits on cross section times branching fraction squared for the multijet (left) and dijet (right) signals as a function of mass and $c\tau$. The overlaid mass-lifetime exclusion curves assume pair-production cross sections for the neutralino (red) and gluino (pink) in multijet signals and top squark cross sections for the dijet signals with 100% branching fraction to each model's respective decay mode specified.

these models for mean proper decay lengths between 100 μm and 15 mm for all masses considered.

Figure 5.57 shows one-dimensional slices of the upper limit as a function of mass for several values of $c\tau$. Similarly, Fig. 5.58 shows the the upper limit as a function of $c\tau$ for a selection of masses.

At a specific signal point, a gluino with a mass of 800 GeV and mean proper decay length $c\tau$ of 1 mm in the 2017 and 2018 dataset alone, the computed 95% CL upper limit on $\sigma\mathcal{B}^2$ is 0.11 fb, compared with the limit from the 2015 and 2016 dataset of 0.3 fb. The improvements primarily arise from the increase in statistical precision due to the increased integrated luminosity of 101 fb^{-1} compared with 38.5 fb^{-1}. However, various other refinements in the analysis improvements in the background rejection of vertices in which tracks originate from separate primary vertices; improvements in the background systematic uncertainty by employing b tags to provide the b quark correla-

tion corrections in the background template, which removed the reliance on simulation in the background prediction; and improvements in the treatment of signal systematic uncertainties, namely in acquiring a better handle on the dominant uncertainty due to the vertex reconstruction efficiency. By combining these two datasets, the 95% CL upper limit for the same signal point moves further down to 0.08 fb.By combining these two datasets, the 95% CL upper limit for the same signal point moves further down to 0.08 fb.

5.9 Theory reinterpretation

While the search presented specifically addresses two models of RPV SUSY, the method described is not uniquely constrained to the models themselves. The results shown here broadly apply to other signal models in which the pair-produced long-lived particle decays into two or more jets in the final state. In this section, we present a set of generator-level selection requirements that, when applied, can approximate the reconstruction-level efficiency of this analysis on any given model and allow for a reinterpretation of the results presented.

Requirements are applied to the generated jet properties, along with the requirements on properties of the constituent long-lived particles and its daughter decay products. The generated jets are clustered from all final state particles, excluding neutrinos, using the anti-k_{T} algorithm with a distance parameter of 0.4. A jet is rejected if the fraction of energy shared by electrons is greater than 0.9, or similarly if the muon energy fraction is greater than 0.8. The daughter particles considered are the u, d, s, c, and b quarks and the electron, muon, and tau leptons from the long-lived particle decay. These daughter particles must have a transverse impact parameter with respect to the origin of at least

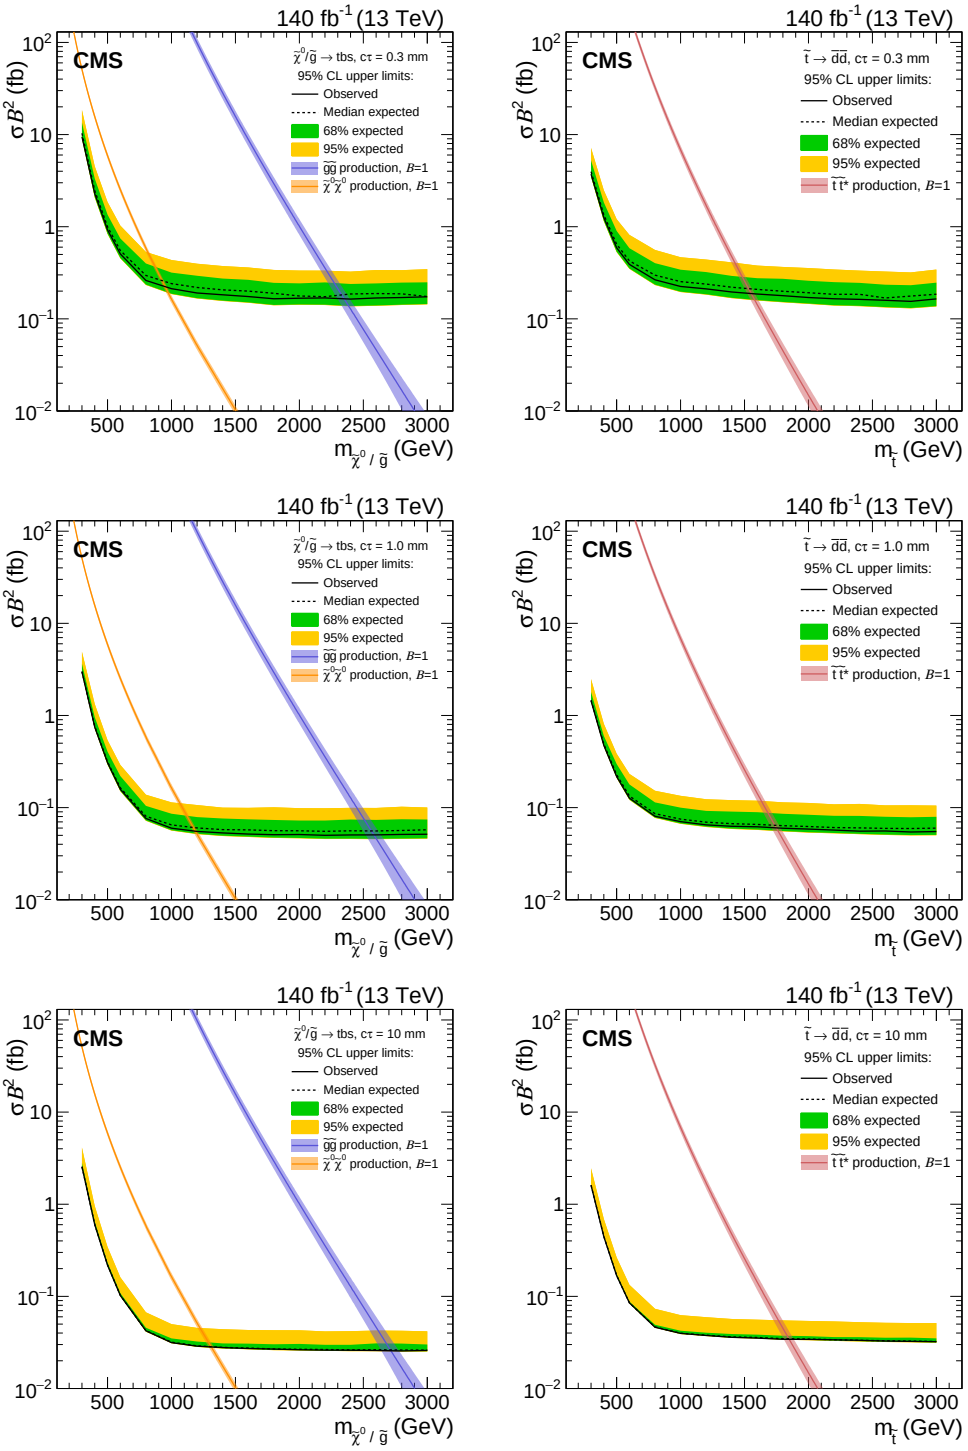

Figure 5.57: Observed and expected 95% CL upper limits on cross section times branching fraction squared as a function of mass for multijet signals (left) and dijet signals (right), for a fixed $c\tau$ of 0.3 mm (top), 1 mm (middle), and 10 mm (bottom) in the full Run 2 data set. The neutralino and gluino pair production cross section is overlaid for the multijet signals, and the top squark pair production cross section is overlaid for the dijet signals.

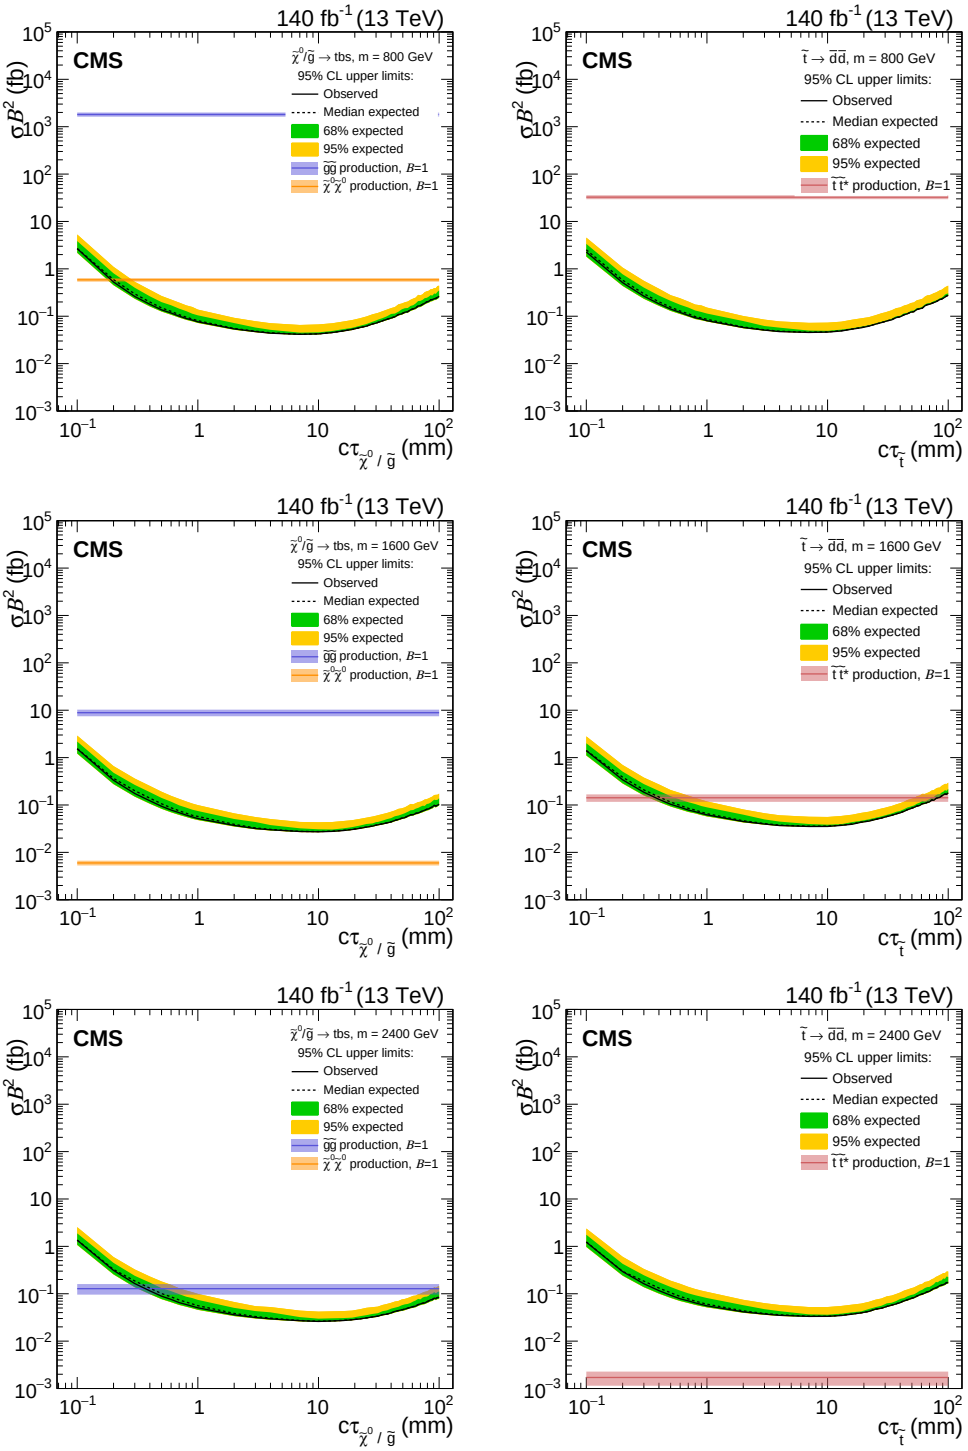

Figure 5.58: Observed and expected 95% CL upper limits on cross section times branching fraction squared as a function of $c\tau$ for multijet signals (left) and dijet signals (right), for a fixed mass of 800 GeV (top), 1600 GeV (middle), and 2400 GeV (bottom) in the full Run 2 data set. The neutralino and gluino pair production cross section is overlaid for the multijet signals, and the top squark pair production cross section is overlaid for the dijet signals.

0.1 mm. To be selected, generated jets and the daughter particles are required to satisfy $p_{\mathrm{T}} > 20\,\mathrm{GeV}$ and $|\eta| < 2.5$.

The following lists the generator-level selection requirements that approximate the reconstruction-level criteria:

- each event must contain at least four generated jets
- $H_{\mathrm{T}} > 1200\,\mathrm{GeV}$, where H_{T} is the scalar sum of the generated jet p_{T} that have $p_{\mathrm{T}} > 40\,\mathrm{GeV}$
- the distance of decay point from the origin in the *x*-*y* plane of each generated long-lived particle must be within 0.1 and 20 mm
- the Σp_{T} of the daughter particles of each long-lived particle must exceed 350 GeV to ensure sufficiently small uncertainty in d_{BV} and sufficiently large number of tracks per vertex. However, if the daughter particle is a b quark its Σp_{T} is scaled down by a factor of 0.65. This corrects for the reduced efficiency at reconstruction-level due to the lifetime of the b quark that can inhibit the association of the decay products to the reconstructed vertex.
- the distance between the two decay points of each long-lived particle must be greater than 0.4 mm

Following this prescription, the generator-level efficiency can approximate the reconstruction-level efficiency to within 20% accuracy for a wide variety of models that have a sufficiently high signal efficiency of 10% or more, and the results of this analysis can be reinterpreted in the context of any other signal model.

CHAPTER 6

BROADER CONTEXTS AND FUTURE DIRECTIONS

This chapter aims to evaluate the analysis described within the broader context of other analyses and searches for new physics beyond the standard model. It will also provide a brief discussion on the even broader question of where the high energy physics community should go in light of the cumulative results that have materialized over the past decade of running the LHC.

A renewed burst in searches for new physics beyond the standard model erupted immediately with the first collision data from the LHC in 2010 as unexplored regions of parameter space were suddenly unlocked with the increased energy and luminosity of the collider. With cautious optimism, these early searches pursued the striking signature of missing transverse energy, the mysterious hint and manifestation of an undetectable particle that was predicted by countless models of SUSY developed over the preceding decades. After anxiously and eagerly waiting many years in anticipation, these searches ultimately bore no fruit, tightly constraining the viability of SUSY. This forced the community to pivot and consider models where initial assumptions from these earlier models were relaxed in order to salvage this theoretical framework, which had provided elegant and satisfying resolutions to several unanswered questions of the standard model.

6.1 Other analyses

The search described in this thesis emerged in the immediate wake of this pivot to pursue the phenomenological signature of a decaying SUSY particle whose decay could be suppressed through various mechanisms, which in turn gave rise to an observable lifetime and displaced signature. One way of contextualizing the result is comparing

this search with prompt searches for RPV SUSY [62, 65], which search for resonances in the mass spectra of multiple reconstructed jets from the SM decay of the LSP. These searches have excluded pair-produced prompt gluinos decaying into trijet final states for masses up to 1500 GeV and prompt top squarks decaying into dijet final states for masses up to 520 GeV, which can be compared to the 2500 GeV and 1500 GeV limits of the same particles, respectively, for certain lifetime ranges of the same decay modes provided by this analysis. The potential long lifetime of the particle provides a powerful handle by moving the search region away from the SM background at the interaction region and allows this displaced vertex search to have better sensitivity to larger masses. Thus, for a more appropriate comparison, it is also important to contextualize the search within the backdrop of other analyses that employ techniques to probe similar regions of parameter space, namely other displaced searches.

As mentioned previously, there are various analyses in both ATLAS and CMS that exploit a particle's lifetime as a discriminator in the search for new physics [7, 8, 64]. The ATLAS searches share similar techniques such as vertex reconstruction, which are performed for events that satisfy other triggers, e.g. lepton or MET triggers, and thus provide complementarity to the H_{T} trigger used by this analysis. Additionally, the ATLAS searches require a single vertex, which improves the inclusiveness of the search and provides greater sensitivity to longer lifetimes, but at the same time, sacrifices the sensitivity to shorter lifetimes that a two vertex requirement provides. In a similar vein, the displaced jet search is another CMS displaced object search that has sensitivity to the same multijet and dijet signal models considered in this displaced vertex search. It uses a pair of dedicated triggers to find events with displaced pairs of jets and reconstructs at least one vertex in the event. A gradient boosted decision tree provides a multivariate discriminant to distinguish the displaced jet signal from background. The displaced jet analysis is ultimately sensitive to a wide range of lifetimes larger than 1 mm, but strug-

gles similarly like other single vertex analyses with sustaining high signal efficiency for shorter lifetimes. In this sense, the displaced vertex analysis described in this thesis plays an important role in the broad landscape of long-lived BSM searches as it provides coverage on less explored regions of parameter space for new physics and complements the discovery region of other searches. Combined with the theorist reinterpretation procedure, this search provides valuable feedback toward future theoretical models of BSM physics.

6.2 Future directions

On Friday, June 19, 2020, the CMS experiment submitted its 1000th publication, just over a full decade after the LHC began its designated research program with 7 TeV collisions on March 30, 2010. These publications are a mix of detector instrumentation-based articles, along with measurements of different properties of the top or Higgs physics and other SM particles, and a sizable proportion ($\sim 1/3$) are analyses in search for new BSM physics. (Although some measurements can, in some sense, also probe BSM physics if they differ from theoretical predictions, so this distinction is not necessarily a hard delineation.) The searches carried out so far across these hundreds of analyses have turned up empty with no significant excess beyond SM background predictions, placing tight constraints on the viability of large regions of BSM parameter space. An important question to step back and reflect upon is where the field of high energy physics should go following this 'crisis' in which no preferred direction has been illuminated by the LHC. The dearth of any discovery of new physics leaves the field with valuable information on where new physics has not yet been found but at the same time with no clear indication as to where it may lie that will be within the existing technological limitations of near future experimental endeavors.

With still many stones unturned, it may seem natural to stay the course by performing iterative refinements in existing technologies by constructing higher energy and higher luminosity colliders combined with more sensitive detector instrumentation for more accurate and precise event reconstruction. However, given the limited resources and without at least a singular well-defined objective as the Higgs search had provided for this generation of collider experiments at the energy frontier, it will be difficult to justify the billions of dollars it would take to pursue this direction of more general purpose collider and detector experiments. On the other end of the spectrum, smaller scale experiments with budgets two or three orders of magnitude smaller than the large collaborations provide a scalpel in identifying particular open questions to pursue and answer unambiguously. Ultimately, a nation's (or even the world's) high energy experimental portfolio should have a well-diversified mix of both these classes of experiments ranging from expensive general purpose but with large breadth and discovery potential as well as more affordable experiments with a more limited scope but more directed inquiry, with a healthy smattering in between these two extremes. The community will have to continue to pivot and continue the creative thinking with close collaboration between experiment and theory to determine where and how to look for new physics. It will be beneficial to re-evaluate the current landscape of existing experimental exclusions and begin ordering and prioritizing a set of questions to focus efforts, for example concentrating on finding and understanding "known unknowns" such as dark matter or better characterizing neutrino properties. This does seem to be the current direction of the field, which may be the best option given the present circumstances and information available, so to that, I say, "Keep at it! 👍"

CHAPTER 7

CONCLUSIONS AND SUMMARY

This dissertation has presented a search for pair-produced long-lived particles decaying into multijet and dijet final states from proton-proton collision events collected with the CMS detector at a center-of-mass energy of 13 TeV during the full Run 2 data collection period, corresponding to an integrated luminosity of 140 fb^{-1}.

No events were observed in the signal region in the 2017 and 2018 datasets, and no excess yield beyond the standard model prediction is observed in the full Run 2 dataset. At 95% confidence level, upper limits are set for an RPV SUSY model in which a long-lived neutralino or gluino decays into a multijet final state with top, bottom, and strange antiquarks. Signal pair-production cross sections larger than 0.08 fb are excluded for long-lived neutralino, gluino, and top squark masses between 800 and 3000 GeV and mean proper decay lengths between 1 mm and 25 mm. For the range of mean proper decay lengths between 0.6 and 90 mm, the data exclude gluino masses up to 2500 GeV. For a neutralino LSP, the data exclude neutralino masses up to 1100 GeV for mean proper decay lengths between 0.6 and 70 mm. Additionally, limits are placed for an RPV SUSY model in which a long-lived top squark decays into a dijet final state with two down antiquarks. The data exclude top squark masses up to 1600 GeV for mean proper decay lengths between 0.4 and 80 mm. These are the most stringent bounds on these models for $c\tau$ between 100 μm and 15 mm for all masses considered, complementing the results of the CMS displaced jet search [64]. While the search directly constrains these two RPV SUSY models, the techniques and methodology are generic and the results are applicable to other models of pair-produced long-lived particles that decay into jets. A method is provided to facilitate the reinterpretation of these results for alternative models.

APPENDIX A

CALCULATIONS FOR B-TAG WEIGHTS IN THE BACKGROUND ESTIMATION PROCEDURE

A.1 Normalization of events with and without b-tagged jets

A.1.1 Determining the number of events with and without b-quarks

To combine the d_{BV} distributions that are sorted into those with and those without at least one Tight b-tag, it is necessary to know the fraction of two-vertex events that contain at least one b-quark, as this fraction is used to normalize the distributions. In the "old" method, this was determined using the generator-level information in simulation; in this "new" method, the b-tagging efficiencies and fake rates can be unfolded based on the relationship between events with/without b-quarks and events with/without b-tagged jets. Generically, these events are related via

$$\begin{pmatrix} N_{\text{b-tag}} \\ N_{\text{b-tag veto}} \end{pmatrix} = \begin{pmatrix} \mathcal{E} & \mathcal{F} \\ \overline{\mathcal{E}} & \overline{\mathcal{F}} \end{pmatrix} \begin{pmatrix} N_{\text{b-quark}} \\ N_{\text{no b-quark}} \end{pmatrix}, \tag{A.1}$$

where

- $N_{\text{b-tag}}$ is the number of events containing at least one b-tagged jet,
- $N_{\text{b-tag veto}}$ is the number of events containing zero b-tagged jets,
- $\mathcal{E}$ is the event-level b-tagging efficiency,
- $\overline{\mathcal{E}} = (1 - \mathcal{E})$,
- $\mathcal{F}$ is the event-level b-tagging fake rate,

- $\overline{\mathcal{F}} = (1 - \mathcal{F})$,
- $N_{\text{b-quark}}$ is the number of events containing at least one b-quark,
- $N_{\text{no b-quark}}$ is the number of events containing zero b-quarks.

Equation A.1 can then be inverted to solve for $N_{\text{b-quark}}$ and $N_{\text{no b-quark}}$:

$$\begin{pmatrix} N_{\text{b-quark}} \\ N_{\text{no b-quark}} \end{pmatrix} = \frac{1}{\mathcal{E}\overline{\mathcal{F}} - \mathcal{F}\overline{\mathcal{E}}} \begin{pmatrix} \overline{\mathcal{F}} & -\mathcal{F} \\ -\overline{\mathcal{E}} & \mathcal{E} \end{pmatrix} \begin{pmatrix} N_{\text{b-tag}} \\ N_{\text{b-tag veto}} \end{pmatrix}. \tag{A.2}$$

This can then be simplified, resulting in

$$N_{\text{b-quark}} = \frac{1}{\mathcal{E} - \mathcal{F}} \left(-\mathcal{F}(N_{\text{b-tag}} + N_{\text{b-tag veto}}) + N_{\text{b-tag}} \right) \tag{A.3}$$

and

$$N_{\text{no b-quark}} = \frac{1}{\mathcal{E} - \mathcal{F}} \left(\mathcal{E}(N_{\text{b-tag}} + N_{\text{b-tag veto}}) - N_{\text{b-tag}} \right). \tag{A.4}$$

These event-level efficiencies and fake rates must be computed based on the per-jet efficiencies and fake rates. In a given event, the probability to find zero b-tagged jets is

$$(1 - \varepsilon)^{n_{\text{b-jets}}} (1 - f)^{n_{l\text{-jets}}}, \tag{A.5}$$

where ε is the per-jet b-tagging efficiency, f is the per-jet b-tagging fake rate, $n_{\text{b-jets}}$ is the number of b-jets in the event, and $n_{l\text{-jets}}$ is the number of light jets (u, d, s, c, g) in the event. Consequently, the probability to find at least one b-tagged jet in a given event is

$$1 - (1 - \varepsilon)^{n_{\text{b-jets}}} (1 - f)^{n_{l\text{-jets}}}. \tag{A.6}$$

This probability can then be used to compute the event-level efficiencies and fake rates via:

$$\mathcal{E} = \frac{\sum_{n_{\text{b-jets}}=1}^{\infty} \sum_{n_{l\text{-jets}}=0}^{\infty} N_{(n_{\text{b-jets}}, n_{l\text{-jets}})} \left[1 - (1 - \varepsilon)^{n_{\text{b-jets}}} (1 - f)^{n_{l\text{-jets}}} \right]}{\sum_{n_{\text{b-jets}}=1}^{\infty} \sum_{n_{l\text{-jets}}=0}^{\infty} N_{(n_{\text{b-jets}}, n_{l\text{-jets}})}} \tag{A.7}$$

and

$$\mathcal{F} = \frac{\sum_{n_{l\text{-jets}}=0}^{\infty} N_{(n_{\text{b-jets}}=0, n_{l\text{-jets}})} \left[1 - (1-f)^{n_{l\text{-jets}}}\right]}{\sum_{n_{l\text{-jets}}=0}^{\infty} N_{(n_{\text{b-jets}}=0, n_{l\text{-jets}})}}, \tag{A.8}$$

where $N_{(n_{\text{b-jets}}, n_{l\text{-jets}})}$ is the number of events with the specified values of $n_{\text{b-jets}}$ and $n_{l\text{-jets}}$. $\mathcal{E}$ and $\mathcal{F}$ are computed using simulated events which are corrected by the b-tagging efficiency data-to-simulation scale factors from [**?**] and [**?**].

A.1.2 Determining the fraction of two-vertex events with b-quarks

Using the number of events with and without b-quarks ($N_{\text{b-quark}}$ and $N_{\text{no b-quark}}$), the fraction of two-vertex events with b-quarks can be computed. Given the probability of reconstructing a vertex from a b-quark decay (denoted P_{vertex}), the probabilities for reconstructing zero, one, or ≥two vertices from b-quarks in an event containing two b-quarks are:

$$\begin{aligned} \text{zero vertices: } & (1 - P_{\text{vertex}})^2 \\ \text{one vertex: } & 2(1 - P_{\text{vertex}})P_{\text{vertex}} \\ \geq\text{two vertices: } & P_{\text{vertex}}^2. \end{aligned} \tag{A.9}$$

Misreconstructed tracks can also result in the reconstruction of fake vertices. The probabilities for reconstructing zero, one, or ≥two fake vertices in a given event are:

$$\begin{aligned} \text{zero vertices: } & 1 - f_{\text{vertex}} \\ \text{one vertex: } & (1 - f_{\text{vertex}})f_{\text{vertex}} \\ \geq\text{two vertices: } & f_{\text{vertex}}^2, \end{aligned} \tag{A.10}$$

where f_{vertex} is the probability to reconstruct a fake vertex from misreconstructed tracks. In events with (without) a b-quark, this probability to reconstruct a fake vertex is denoted

$f_{\text{vertex, b-quark}}$ ($f_{\text{vertex, no b-quark}}$). As a result, the total probabilities to reconstruct zero, one, or ≥two vertices in events with a b-quark are:

$$\begin{aligned}
P_{\text{0-vertex, b-quark}} &= (1 - P_{\text{vertex}})^2(1 - f_{\text{vertex, b-quark}}) \\
P_{\text{1-vertex, b-quark}} &= 2P_{\text{vertex}}(1 - f_{\text{vertex, b-quark}}) + (1 - P_{\text{vertex}})^2(1 - f_{\text{vertex, b-quark}})f_{\text{vertex, b-quark}} \\
P_{\text{2-vertex, b-quark}} &= P^2_{\text{vertex}}(1 - f_{\text{vertex, b-quark}}) + 2P_{\text{vertex}}(1 - f_{\text{vertex, b-quark}})f_{\text{vertex, b-quark}} \\
&\quad + (1 - P_{\text{vertex}})^2 f^2_{\text{vertex, b-quark}},
\end{aligned} \tag{A.11}$$

while in events without a b-quark, these probabilities are:

$$\begin{aligned}
P_{\text{0-vertex, no b-quark}} &= (1 - f_{\text{vertex, no b-quark}}) \\
P_{\text{1-vertex, no b-quark}} &= (1 - f_{\text{vertex, no b-quark}})f_{\text{vertex, no b-quark}} \\
P_{\text{2-vertex, no b-quark}} &= f^2_{\text{vertex, no b-quark}}.
\end{aligned} \tag{A.12}$$

With these probabilities in hand, it is possible to compute the b-quark fractions. The fraction of preselected and one-vertex events with b-quarks are denoted $F_{\text{presel, b-quark}}$ and $F_{\text{1-vertex, b-quark}}$, respectively. For preselected events, the b-quark fraction is:

$$F_{\text{presel, b-quark}} = \frac{N_{\text{b-quark}}}{N_{\text{b-quark}} + N_{\text{no b-quark}}} \tag{A.13}$$

while for one-vertex events, it is:

$$F_{\text{1-vertex, b-quark}} = \frac{P_{\text{1-vertex, b-quark}}N_{\text{b-quark}}}{P_{\text{1-vertex, b-quark}}N_{\text{b-quark}} + P_{\text{1-vertex, no b-quark}}N_{\text{no b-quark}}}. \tag{A.14}$$

Similarly, the fraction of two-vertex events with b-quarks can be written as

$$F_{\text{2-vertex, b-quark}} = \frac{P_{\text{2-vertex, b-quark}}C_{\text{b-quark}}N_{\text{b-quark}}}{P_{\text{2-vertex, b-quark}}C_{\text{b-quark}}N_{\text{b-quark}} + P_{\text{2-vertex, no b-quark}}C_{\text{no b-quark}}N_{\text{no b-quark}}}, \tag{A.15}$$

where $C_{\text{b-quark}}$ ($C_{\text{no b-quark}}$) is the efficiency correction factor for the d^{C}_{VV} template constructed from one-vertex events with (without) b-quarks. Equation A.15 can then be rewritten as

$$F_{\text{2-vertex, b-quark}} = \frac{a}{1 + a} \tag{A.16}$$

using the substitution

$$a = \frac{P_{\text{2-vertex, b-quark}} C_{\text{b-quark}} N_{\text{b-quark}}}{P_{\text{2-vertex, no b-quark}} C_{\text{no b-quark}} N_{\text{no b-quark}}}. \tag{A.17}$$

Since the goal is to compute $F_{\text{2-vertex, b-quark}}$, and the inputs $C_{\text{b-quark}}$, $C_{\text{no b-quark}}$, $N_{\text{b-quark}}$, and $N_{\text{no b-quark}}$ can be directly computed, all that remains to be determined is $\frac{P_{\text{2-vertex, b-quark}}}{P_{\text{2-vertex, no b-quark}}}$. As the probability to reconstruct fake vertices tends to be small, and only a small fraction of b-quarks result in a reconstructed vertex, the limits are considered where $f_{\text{vertex, b-quark}} \ll 1$, $f_{\text{vertex, no b-quark}} \ll 1$, and $P_{\text{vertex}} \ll 1$. In this limit, $\frac{P_{\text{1-vertex, b-quark}}}{P_{\text{1-vertex, no b-quark}}}$ and $\frac{P_{\text{2-vertex, b-quark}}}{P_{\text{2-vertex, no b-quark}}}$ can be computed using Equations A.11 and A.12 to be:

$$\frac{P_{\text{1-vertex, b-quark}}}{P_{\text{1-vertex, no b-quark}}} = \frac{2P_{\text{vertex}} + f_{\text{vertex, b-quark}}}{f_{\text{vertex, no b-quark}}} \tag{A.18}$$

and

$$\frac{P_{\text{2-vertex, b-quark}}}{P_{\text{2-vertex, no b-quark}}} = \frac{(P_{\text{vertex}} + f_{\text{vertex, b-quark}})^2}{f^2_{\text{vertex, no b-quark}}}. \tag{A.19}$$

An additional substitution can be made, namely:

$$s = \frac{2P_{\text{vertex}} + f_{\text{vertex, b-quark}}}{P_{\text{vertex}} + f_{\text{vertex, b-quark}}}. \tag{A.20}$$

This allows Equation A.19 to be rewritten in terms of $\frac{P_{\text{1-vertex, b-quark}}}{P_{\text{1-vertex, no b-quark}}}$ as:

$$\frac{P_{\text{2-vertex, b-quark}}}{P_{\text{2-vertex, no b-quark}}} = \frac{1}{s^2}\left(\frac{P_{\text{1-vertex, b-quark}}}{P_{\text{1-vertex, no b-quark}}}\right)^2, \tag{A.21}$$

and subsequently

$$a = \frac{1}{s^2}\left(\frac{P_{\text{1-vertex, b-quark}}}{P_{\text{1-vertex, no b-quark}}}\right)^2 \frac{C_{\text{b-quark}} N_{\text{b-quark}}}{C_{\text{no b-quark}} N_{\text{no b-quark}}}. \tag{A.22}$$

The only unknown that remains for determining $F_{\text{2-vertex, b-quark}}$ is s; a value of $s = 1$ assumes the probability of finding two vertices is equal to the square of the probability of finding one vertex (i.e. the vertices are uncorrelated). As we expect this to be the case, we use $s = 1$ as our nominal value, but assign an uncertainty (described in Section 5.7.2) based on variations to $F_{\text{2-vertex, b-quark}}$.

APPENDIX B

PREDICTING TWO-VERTEX YIELD FROM ONE-VERTEX EVENTS

The calculation for estimating the number of two-vertex background events uses many of the same mathematical arguments and assumptions from Appendix A with a focus on deriving an absolute yield of two-vertex events instead of relative proportions in different event classes. Many of the variables defined in the previous Appendix will be used throughout the calculation. We begin with the same assumption that the two-vertex background events can be broken into separate categories, namely events containing at least one b-quark and those that do not. We can start with an initial relationship between the absolute number of two-vertex events N_{2v} and the number of events with and without a b-quark and the respective probabilities of reconstructing a vertex in such an event:

$$N_{2v} = N_{\text{b-quark}} C_{\text{b-quark}} f^2_{\text{vertex, b-quark}} + N_{\text{no b-quark}} C_{\text{no b-quark}} f^2_{\text{vertex, no b-quark}} \tag{B.1}$$

Here, as defined before $N_{\text{no b-quark}}$ and $N_{\text{b-quark}}$ correspond to the number of events not containing a b-quark and ones that do, and $f_{\text{vertex, no b-quark}}$ and $f_{\text{vertex, b-quark}}$ are the probabilities of reconstructing a single vertex in the respective event category. The $C_{\text{b-quark}}$ and $C_{\text{no b-quark}}$ correspond to the efficiency correction factor to account for the probability that the vertex positions will overlap. It is helpful to remember that $f_{\text{vertex, b-quark}}$ is different for different track multiplicity vertices, but they will sum to unity, i.e. $f_{\text{0-vertex, b}} + f_{\text{3-track, b}} + f_{\text{4-track, b}} + f_{\geq\text{5-track, b}} = 1$. In general, $f_{n \text{ or } m} = f_{n\text{-track, b}} + f_{m\text{-track, b}}$, which, when used in equation B.1 and using a little algebra, yields:

$$
\begin{aligned}
N_{\text{2v, n or m}} = & \underbrace{(N_{\text{b-quark}} C_{\text{b-quark}} f^2_{n\text{-track, b}} + N_{\text{no b-quark}} C_{\text{no b-quark}} f^2_{n\text{-track, no b}})}_{N_{n\times n}} \\
& + \underbrace{(N_{\text{b-quark}} C_{\text{b-quark}} f^2_{m\text{-track, b}} + N_{\text{no b-quark}} C_{\text{no b-quark}} f^2_{m\text{-track, no b}})}_{N_{m\times m}} \\
& + \underbrace{2(N_{\text{b-quark}} C_{\text{b-quark}} f_{n\text{-track, b}} f_{m\text{-track, b}} + N_{\text{no b-quark}} C_{\text{no b-quark}} f_{n\text{-track, no b}} f_{m\text{-track, no b}})}_{N_{n\times m}}
\end{aligned}
\tag{B.2}
$$

We can derive a relation between $N_{1\text{v}}$, the number of one-vertex events, to the probabilities in a similar argument:

$$
N_{1\text{v}} = N_{\text{b-quark}} f_{\text{vertex, b-quark}} + N_{\text{no b-quark}} f_{\text{vertex, no b-quark}} \tag{B.3}
$$

The subscript denoting the track multiplicity of the vertex is omitted but implied. We can rewrite this equation in terms of $f_{\text{vertex, no b-quark}}$ as:

$$
f_{\text{vertex, no b-quark}} = \frac{N_{1\text{v}}}{N_{\text{no b-quark}} \left[\left(\frac{N_{\text{b-quark}}}{N_{\text{no b-quark}}} \right) \left(\frac{f_{\text{vertex, b-quark}}}{f_{\text{vertex, no b-quark}}} \right) + 1 \right]} \tag{B.4}
$$

$$
= \frac{N_{1\text{v}}}{N_{\text{no b-quark}} \left[\left(\frac{F_{\text{1-vertex, b-quark}}}{1 - F_{\text{1-vertex, b-quark}}} \right) + 1 \right]} \tag{B.5}
$$

The second line in equation B.5 arises from the relation $\frac{f_{\text{vertex, b-quark}}}{f_{\text{vertex, no b-quark}}} = \frac{F_{\text{1-vertex, b-quark}}}{1 - F_{\text{1-vertex, b-quark}}} \frac{1 - F_{\text{presel, b-quark}}}{F_{\text{presel, b-quark}}}$ and $\frac{N_{\text{b-quark}}}{N_{\text{no b-quark}}} = \frac{F_{\text{presel, b-quark}}}{1 - F_{\text{presel, b-quark}}}$. We can return to equation B.2 and expand one of the terms explicitly:

$$N_{2\text{v}} = N_{\text{no b-quark}} C_{\text{no b-quark}} f_{n\text{-track, no b}} f_{m\text{-track, no b}} \times \left[\frac{N_{\text{b-quark}}}{N_{\text{no b-quark}}} \frac{C_{\text{b-quark}}}{C_{\text{no b-quark}}} \frac{f_{n\text{-track, b}}}{f_{n\text{-track, no b}}} \frac{f_{m\text{-track, b}}}{f_{m\text{-track, no b}}} + 1 \right] \tag{B.6}$$

Substituting equation B.5 into this and simplifying and using $N_{\text{no b-quark}}/N_{\text{presel}} = 1 - F_{\text{presel, b-quark}}$, we finally have the following equation for the predicted number of two-vertex events:

$$N_{n\times m} = \left[\frac{N_{1\text{v}, n} N_{1\text{v}, m} C_{\text{no b-quark}}}{N_{\text{presel}}(1 - F_{\text{presel, b-quark}}) \left(\frac{F_{n\text{-track, b}}}{1-F_{n\text{-track, b}}} + 1 \right) \left(\frac{F_{m\text{-track, b}}}{1-F_{m\text{-track, b}}} + 1 \right)} \right] \times \left[\left(\frac{F_{n\text{-track, b}}}{1 - F_{n\text{-track, b}}} \right) \left(\frac{F_{m\text{-track, b}}}{1 - F_{m\text{-track, b}}} \right) \left(\frac{1 - F_{\text{presel, b-quark}}}{F_{\text{presel, b-quark}}} \right) \left(\frac{C_{\text{b-quark}}}{C_{\text{no b-quark}}} \right) + 1 \right] \tag{B.7}$$

Each of the independent variables are obtained from data. The efficiency correction terms ($C_{\text{b-quark}}$ and $C_{\text{no b-quark}}$) are derived from the integrated effect of applying the overlap efficiency curve for 5-track vertices in data used in the variation study for the systematic uncertainty indicated by the green curve in Figure 5.53. This particular curve was chosen as it provides the most conservative yield estimate that addresses this effect barring no correction at all. Other curves were studied and the final effect on the observed limit using the correction terms derived from other curves results in a difference of less than a percent, indicating a very weak sensitivity to this correction. With this, the values for each of these variables can be found in Table B.1 for the 2017 and 2018 data. These can be broken up into the constituent years, 2017 and 2018 in the two tables that follow, Tables B.2 and B.3. The last few tables show the resulting predicted values in the 2017 and 2018 data in the combined dataset for Table B.4 and its constituent years in Tables B.5 and B.6.

Table B.1: Input variables for predicting the number of two-vertex events for different vertex track multiplicities in 2017 and 2018 data.

Event Category	N_{presel}	$F_{\text{presel, b-quark}}$	$N_{1\text{v}}$		$F_{1\text{-vertex, b-quark}}$		$C_{\text{b-quark}}$	$C_{\text{no b-quark}}$
3-track	35802016	0.156	61818		0.496		0.909	0.888
4-track			14730		0.533		0.903	0.883
5-track			2211		0.643		0.895	0.874
3-track × 4-track			61818	14730	0.496	0.533	0.906	0.886
3-track × 5-track			61818	2211	0.496	0.643	0.903	0.883
4-track × 5-track			14730	2211	0.533	0.643	0.900	0.879

Table B.2: Input variables for predicting the number of two-vertex events for different vertex track multiplicities in 2017 data.

Event Category	N_{presel}	$F_{\text{presel, b-quark}}$	$N_{1\text{v}}$		$F_{1\text{-vertex, b-quark}}$		$C_{\text{b-quark}}$	$C_{\text{no b-quark}}$
3-track	15067777	0.166	32152		0.489		0.921	0.908
4-track			7838		0.509		0.917	0.903
5-track			1303		0.599		0.911	0.898
3-track × 4-track			32152	7838	0.489	0.509	0.919	0.905
3-track × 5-track			32152	1303	0.489	0.599	0.917	0.904
4-track × 5-track			7838	1303	0.509	0.599	0.915	0.901

Table B.3: Input variables for predicting the number of two-vertex events for different vertex track multiplicities in 2018 data.

Event Category	N_{presel}	$F_{\text{presel, b-quark}}$	$N_{1\text{v}}$		$F_{1\text{-vertex, b-quark}}$		$C_{\text{b-quark}}$	$C_{\text{no b-quark}}$
3-track	20734240	0.153	29666		0.504		0.888	0.860
4-track			6892		0.555		0.880	0.852
5-track			908		0.689		0.870	0.841
3-track × 4-track			29666	6892	0.504	0.555	0.884	0.856
3-track × 5-track			29666	908	0.504	0.689	0.880	0.852
4-track × 5-track			6892	908	0.555	0.689	0.876	0.847

Table B.4: Predicted and actual numbers of two-vertex events in 2017 and 2018 data for different vertex track multiplicity events.

Event Category	Predicted N_{2v}	Observed N_{2v}	Ratio (Obs./Pred.)
3-track	181.8	185	$1.02^{+0.08}_{-0.08}$
4-track	11.4	12	$1.05^{+0.40}_{-0.30}$
5-track	0.34	0	$0.00^{+5.38}_{-0.00}$
3-track × 4-track	90.9	101	$1.11^{+0.12}_{-0.11}$
3-track × 5-track	15.6	9	$0.58^{+0.26}_{-0.19}$
4-track × 5-track	3.92	1	$0.26^{+0.58}_{-0.22}$

Table B.5: Predicted and actual numbers of two-vertex events in 2017 data for different vertex track multiplicity events.

Event Category	Predicted N_{2v}	Observed N_{2v}	Ratio (Obs./Pred.)
3-track	101.5	113	$1.02^{+0.11}_{-0.09}$
4-track	6.91	9	$1.30^{+0.60}_{-0.42}$
5-track	0.24	0	$0.00^{+7.63}_{-0.00}$
3-track × 4-track	55.2	59	$1.07^{+0.16}_{-0.14}$
3-track × 5-track	10.2	5	$0.49^{+0.33}_{-0.21}$
4-track × 5-track	2.57	0	$0.00^{+0.72}_{-0.00}$

Table B.6: Predicted and actual numbers of two-vertex events in 2018 data for different vertex track multiplicity events.

Event Category	Predicted N_{2v}	Observed N_{2v}	Ratio (Obs./Pred.)
3-track	73.0	72	$0.99^{+0.13}_{-0.12}$
4-track	4.51	3	$0.67^{+0.64}_{-0.37}$
5-track	0.11	0	$0.00^{+16.6}_{-0.00}$
3-track × 4-track	36.2	42	$1.16^{+0.21}_{-0.18}$
3-track × 5-track	5.58	4	$0.72^{+0.56}_{-0.35}$
4-track × 5-track	1.40	1	$0.71^{+1.64}_{-0.59}$

BIBLIOGRAPHY

[1] About cms. Accessed: 2020-07-09.

[2] The CMS muon project: Technical Design Report. 1997.

[3] LHC Design Report Vol.1: The LHC Main Ring. 6 2004.

[4] CMS luminosity measurement for the 2017 data-taking period at $\sqrt{s} = 13\,\text{TeV}$. CMS Physics Analysis Summary CMS-PAS-LUM-17-001, CERN, 2018.

[5] CMS luminosity measurement for the 2018 data-taking period at $\sqrt{s} = 13\,\text{TeV}$. CMS Physics Analysis Summary CMS-PAS-LUM-18-002, CERN, 2019.

[6] Track impact parameter resolution in the 2017 dataset with the CMS Phase-1 Pixel detector. CMS Detector Performance Summary CMS-DP-2020-032, 2020.

[7] Morad Aaboud et al. Search for long-lived, massive particles in events with displaced vertices and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector. *Phys. Rev. D*, 97:052012, 2018.

[8] Georges Aad et al. Search for long-lived, massive particles in events with a displaced vertex and a muon with large impact parameter in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. 3 2020.

[9] W. Adam et al. The DAQ and control system for the CMS Phase-1 pixel detector upgrade. *JINST*, 14(10):P10017, 2019.

[10] S. Agostinelli et al. GEANT4—a simulation toolkit. *Nucl. Instrum. Meth. A*, 506:250, 2003.

[11] Juliette Alimena et al. Searching for Long-Lived Particles beyond the Standard Model at the Large Hadron Collider. 3 2019.

[12] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro. The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *JHEP*, 07:079, 2014.

[13] J. Alwall, S. Höche, F. Krauss, N. Lavesson, L. Lönnblad, F. Maltoni, M. L. Mangano, M. Moretti, C. G. Papadopoulos, F. Piccinini, S. Schumann, M. Treccani, J. Winter, and M. Worek. Comparative study of various algorithms for the

merging of parton showers and matrix elements in hadronic collisions. *Eur. Phys. J. C*, 53:473, 2008.

[14] N. Arkani-Hamed, S. Dimopoulos, G. F. Giudice, and A. Romanino. Aspects of split supersymmetry. *Nucl. Phys. B*, 709:3, 2005.

[15] Nima Arkani-Hamed and Savas Dimopoulos. Supersymmetric unification without low energy supersymmetry and signatures for fine-tuning at the LHC. *JHEP*, 06:073, 2005.

[16] Nima Arkani-Hamed, Arpit Gupta, David E. Kaplan, Neal Weiner, and Tom Zorawski. Simply unnatural supersymmetry. 2012.

[17] Asimina Arvanitaki, Nathaniel Craig, Savas Dimopoulos, and Giovanni Villadoro. Mini-split. *JHEP*, 02:126, 2013.

[18] Howard Baer and Xerxes Tata. *Weak Scale Supersymmetry*. Cambridge University Press, 2006.

[19] Richard D. Ball et al. Parton distributions for the LHC Run II. *JHEP*, 04:040, 2015.

[20] Richard D. Ball et al. Parton distributions from high-precision collider data. *Eur. Phys. J.*, C77(10):663, 2017.

[21] R. Barbier, C. Bérat, M. Besançon, M. Chemtob, A. Deandrea, E. Dudas, P. Fayet, S. Lavignac, G. Moreau, E. Perez, and Y. Sirois. *R*-parity violating supersymmetry. *Phys. Rept.*, 420:1, 2005.

[22] G. L. Bayatian et al. CMS Physics: Technical Design Report Volume 1: Detector Performance and Software. 2006.

[23] Wim Beenakker, Christoph Borschensky, Raphael Heger, Michael Krmer, Anna Kulesza, and Eric Laenen. NNLL resummation for stop pair-production at the LHC. *JHEP*, 05:153, 2016.

[24] Martin Beneke, Jan Piclum, Christian Schwinn, and Christopher Wever. NNLL soft and Coulomb resummation for squark and gluino production at the LHC. *JHEP*, 10:054, 2016.

[25] Christoph Borschensky, Michael Krämer, Anna Kulesza, Michelangelo Mangano, Sanjay Padhi, Tilman Plehn, and Xavier Portell. Squark and gluino production

cross sections in pp collisions at $\sqrt{s}$ = 13, 14, 33 and 100 TeV. *Eur. Phys. J. C*, 74:3174, 2014.

[26] Rym Bouchendira, Pierre Cladé, Saïda Guellati-Khélifa, Fran çois Nez, and Fran çois Biraben. New determination of the fine structure constant and test of the quantum electrodynamics. *Phys. Rev. Lett.*, 106:080801, Feb 2011.

[27] Jon Butterworth et al. PDF4LHC recommendations for LHC Run II. *J. Phys. G*, 43:023001, 2016.

[28] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. The anti-k_t jet clustering algorithm. *JHEP*, 04:063, 2008.

[29] Lorenzo Calibbi, Laura Lopez-Honorez, Steven Lowette, and Alberto Mariotti. Singlet-Doublet dark matter freeze-in: LHC displaced signatures versus cosmology. *JHEP*, 09:037, 2018.

[30] E. Chabanat and N. Estre. Deterministic annealing for vertex finding at CMS. In *14th International Conference on Computing in High-Energy and Nuclear Physics*, pages 287–290, 2005.

[31] S. Chatrchyan et al. The CMS Experiment at the CERN LHC. *JINST*, 3:S08004, 2008.

[32] S Chatrchyan et al. Performance and Operation of the CMS Electromagnetic Calorimeter. *JINST*, 5:T03010, 2010.

[33] S Chatrchyan et al. Precise Mapping of the Magnetic Field in the CMS Barrel Yoke using Cosmic Rays. *JINST*, 5:T03021, 2010.

[34] Serguei Chatrchyan et al. The Performance of the CMS Muon Detector in Proton-Proton Collisions at $\sqrt{s}$ = 7 TeV at the LHC. *JINST*, 8:P11002, 2013.

[35] Serguei Chatrchyan et al. Description and performance of track and primary-vertex reconstruction with the CMS tracker. *JINST*, 9(10):P10009, 2014.

[36] Raymond T. Co, Francesco D'Eramo, Lawrence J. Hall, and Duccio Pappadopulo. Freeze-In dark matter with displaced signatures at colliders. *JCAP*, 1512:024, 2015.

[37] Csaba Csáki, Yuval Grossman, and Ben Heidenreich. Minimal flavor violation

supersymmetry: a natural theory for *R*-parity violation. *Phys. Rev. D*, 85:095009, 2012.

[38] Csaba Csáki, Eric Kuflik, Salvator Lombardo, Oren Slone, and Tomer Volansky. Phenomenology of a long-lived LSP with *R*-parity violation. *JHEP*, 08:016, 2015.

[39] Yanou Cui, Lisa Randall, and Brian Shuve. A WIMPy baryogenesis miracle. *JHEP*, 04:075, 2012.

[40] Yanou Cui and Brian Shuve. Probing baryogenesis with displaced vertices at the LHC. *JHEP*, 02:049, 2015.

[41] Yanou Cui and Raman Sundrum. Baryogenesis for weakly interacting massive particles. *Phys. Rev. D*, 87:116013, 2013.

[42] G. D'Ambrosio, G. F. Giudice, G. Isidori, and A. Strumia. Minimal flavor violation: An Effective field theory approach. *Nucl. Phys.*, B645:155–187, 2002.

[43] JiJi Fan, Matthew Reece, and Joshua T. Ruderman. Stealth supersymmetry. *JHEP*, 11:012, 2011.

[44] JiJi Fan, Matthew Reece, and Joshua T. Ruderman. A stealth supersymmetry sampler. *JHEP*, 07:196, 2012.

[45] Benjamin Fuks, Michael Klasen, David R. Lamprea, and Marcel Rothering. Gaugino production in proton-proton collisions at a center-of-mass energy of 8 TeV. *JHEP*, 10:081, 2012.

[46] Benjamin Fuks, Michael Klasen, David R. Lamprea, and Marcel Rothering. Precision predictions for electroweak superpartner production at hadron colliders with Resummino. *Eur. Phys. J. C*, 73:2480, 2013.

[47] P. Gambino, G. F. Giudice, and P. Slavich. Gluino decays in split supersymmetry. *Nucl. Phys. B*, 726:35, 2005.

[48] G. F. Giudice and A. Romanino. Split supersymmetry. *Nucl. Phys. B*, 699:65, 2004.

[49] Lawrence J. Hall, Karsten Jedamzik, John March-Russell, and Stephen M. West. Freeze-in production of FIMP dark matter. *JHEP*, 03:080, 2010.

[50] Lawrence J. Hall and Mahiko Suzuki. Explicit *R*-parity breaking in supersymmetric models. *Nucl. Phys. B*, 231:419, 1984.

[51] Tao Han, Zongguo Si, Kathryn M. Zurek, and Matthew J. Strassler. Phenomenology of hidden valleys at hadron colliders. *JHEP*, 07:008, 2008.

[52] D. Hanneke, S.Fogwell Hoogerheide, and G. Gabrielse. Cavity Control of a Single-Electron Quantum Cyclotron: Measuring the Electron Magnetic Moment. *Phys. Rev. A*, 83:052122, 2011.

[53] JoAnne L. Hewett, Ben Lillie, Manuel Masip, and Thomas G. Rizzo. Signatures of long-lived gluinos in split supersymmetry. *JHEP*, 09:070, 2004.

[54] D. Hits and A. Starodumov. The CMS Pixel Readout Chip for the Phase 1 Upgrade. *JINST*, 10(05):C05029, 2015.

[55] David E. Kaplan, Markus A. Luty, and Kathryn M. Zurek. Asymmetric dark matter. *Phys. Rev. D*, 79:115016, 2009.

[56] Ian-Woo Kim and Kathryn M. Zurek. Flavor and collider signatures of asymmetric dark matter. *Phys. Rev. D*, 89:035008, 2014.

[57] Stephen P. Martin. A Supersymmetry primer. 1997. [Adv. Ser. Direct. High Energy Phys.18,1(1998)].

[58] Emanuel Nikolidakis and Christopher Smith. Minimal Flavor Violation, Seesaw, and R-parity. *Phys. Rev.*, D77:015021, 2008.

[59] A. M. Sirunyan et al. Particle-flow reconstruction and global event description with the CMS detector. *JINST*, 12:P10003, 2017.

[60] Albert M Sirunyan et al. Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV. Submitted to *JINST*, 2017.

[61] Albert M. Sirunyan et al. Search for long-lived particles with displaced vertices in multijet events in proton-proton collisions at $\sqrt{s} = 13\,\text{TeV}$. *Phys. Rev. D*, 98(arXiv:1808.03078. CMS-EXO-17-018-003):092011. 25 p, Aug 2018.

[62] Albert M Sirunyan et al. Search for pair-produced resonances decaying to quark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV. *Phys. Rev. D*, 98(11):112014, 2018.

[63] Albert M Sirunyan et al. Extraction and validation of a new set of CMS PYTHIA8 tunes from underlying-event measurements. 2019.

[64] Albert M Sirunyan et al. Search for long-lived particles decaying into displaced jets in proton-proton collisions at $\sqrt{s}$ = 13 TeV. *Phys. Rev. D*, 99(3):032011, 2019.

[65] Albert M Sirunyan et al. Search for pair-produced three-jet resonances in proton-proton collisions at $\sqrt{s}$ =13 TeV. *Phys. Rev. D*, 99(1):012010, 2019.

[66] A.M. Sirunyan et al. Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s}$ = 13 TeV. *JINST*, 13(06):P06015, 2018.

[67] Torbjörn Sjöstrand, Stefan Ask, Jesper R. Christiansen, Richard Corke, Nishita Desai, Philip Ilten, Stephen Mrenna, Stefan Prestel, Christine O. Rasmussen, and Peter Z. Skands. An Introduction to PYTHIA 8.2. *Comput. Phys. Commun.*, 191:159–177, 2015.

[68] Matthew J. Strassler and Kathryn M. Zurek. Echoes of a hidden valley at hadron colliders. *Phys. Lett. B*, 651:374, 2007.

[69] Matthew J. Strassler and Kathryn M. Zurek. Discovering the Higgs through highly-displaced vertices. *Phys. Lett. B*, 661:263, 2008.