A New Strategy to Search for Natural Squirks

Carson Tenney

A senior thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Bachelor of Science

Christopher B. Verhaaren, Advisor

Department of Physics and Astronomy

Brigham Young University

ABSTRACT

A New Strategy to Search for Natural Squirks

Carson Tenney
Department of Physics and Astronomy, BYU
Bachelor of Science

Squirks are theoretical particles predicted in many realizations of neutral naturalness, a framework which extends the Standard Model to address the electroweak hierarchy problem. For a two-squirk bound state with net electric charge, the annihilation branching ratios depend on the mass difference Δ between the squirks. In this work, I show that as Δ increases, the Δ-dependent branching ratios weaken experimental bounds on squirk masses. In particular, I find that some squirk masses less than 1 TeV have not yet been excluded by searches at the Large Hadron Collider (LHC). These masses are accessible at the LHC's current energy levels, so with a new strategy, we can search for squirks more effectively, increasing the discovery power of the LHC while maintaining its current collision energy. In this paper, I present one such new strategy, which is to search for events where the collision produces a pair of squirks with a glueball. The squirks annihilate into a pair of bosons, and the glueball creates a displaced vertex when it decays in the detector. I also present our simulations of searching for this combination of a bosonic resonance and a displaced vertex. I find that this strategy complements and extends current methods, probing a larger region of squirk mass parameter space than the region excluded by current search strategies, especially for large Δ.

Keywords: squirks, glueballs, neutral naturalness, $W\gamma$ resonance, WZ resonance, displaced vertex

ACKNOWLEDGMENTS

I'd like to thank my wife Emma for her loving support through the many hours I put into this project. Her commitment to me and my education provides me with the stability I need to be an effective student and researcher.

I'd also like to thank my advisor Chris Verhaaren for helping me contribute meaningfully to publishable research. Similarly, I'd like to thank Josh Forsyth for his patience in helping me understand the nuances of this project. Working with them taught me to think like a researcher, to communicate my research effectively, and to learn outside a structured course.

I'm grateful for the PHSCS 416 writing class taught by Scott Bergeson and Chanel Earl. This course gave me the guidance and resources I needed to make this thesis into something I can be proud of.

I'm grateful to the BYU College of Computational, Mathematical, and Physical Sciences and to the Physics and Astronomy Department for financial support in my RA position. I'm also grateful for the opportunities they provided for me to present my research at Student Research Conferences and at the APS Four Corners Annual Meeting. These conference experiences continue to inspire me to become a better scientific communicator.

The research I present in this thesis was funded in part by the National Science Foundation under Grant No. PHY-2210067.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The Large Hadron Collider (LHC) has been pivotal in confirming the Standard Model (SM) of particle physics, most notably with the discovery of the Higgs boson at the CMS and ATLAS detectors [1, 2]. While the SM successfully predicts many observable phenomena, it has known limitations. For example, the SM does not include gravity, does not explain neutrino masses, and does not describe dark matter. To address these limitations, particle physicists develop Beyond-the-Standard-Model (BSM) theories, and colliders like the LHC can test and constrain these models.

In running experiments at the LHC, scientists have found that the detectors generate more data than they could possibly store [3]. To work around this issue, physicists create search strategies by choosing specific signals (for example, a pair of photons with high enough energy) theorized to occur in the model being tested. Then, they program the detectors to look for those signals and to store only the data relevant to the search. After the experiment, they look at the data to see if the signal occurred more than the null hypothesis would predict. Observing more signal than the predicted background is evidence of a discovery! This method has already proved effective, evidenced by the discovery of the Higgs boson noted above.

In the near future, the LHC will receive upgrades to increase its luminosity, meaning it will collide more particles per second [4]. The upgraded LHC is known as the high-luminosity-LHC (HL-LHC). These upgrades will significantly increase LHC data output, so effective search strategies will be even more valuable for extracting useful data from experiments. Importantly, the HL-LHC will not increase the collision energy, so any improvements to search strategies cannot depend on increasing the energy of the collider. I recently helped develop a novel search strategy [5] that meets this criterion, and in this paper, I will discuss this strategy and its expected performance at the HL-LHC.

My collaborators and I developed our strategy to search more effectively for evidence of neutral naturalness [6], a framework of BSM theories. Neutral naturalness aims to address a limitation in the SM known as the electroweak hierarchy problem [7], which is that the SM does not explain why the Higgs mass is so much lower than the Planck mass, the characteristic scale of gravity. Even though one can explain the low Higgs mass using fine-tuned parameters, doing so is unsatisfying. Realizations of neutral naturalness posit additional symmetries beyond those of the SM, creating a more “natural” explanation for the observed Higgs mass. The symmetries lead to cancellations of high-energy corrections to the Higgs mass as calculated in quantum field theory, leaving the Higgs mass insensitive to high energy scales (like the Planck scale) and making the observed mass unsurprising. The “naturalness” in neutral naturalness refers to the model’s symmetry-based explanation of the Higgs mass.

In Ref. [5], we focused on models that combine supersymmetry (SUSY) [8] with a discrete symmetry [9] to explain the Higgs mass. In these models, the SM quarks have symmetry partners called squarks, where the “s” added to the front of “quarks” stands for “scalar” since the squarks have spin-0. In SUSY without the discrete symmetry, squarks are charged under the SM’s SU(3) gauge symmetry (associated with the strong nuclear force). However, including the discrete symmetry allows the squarks to be charged under a separate SU(3) and neutral under the SM SU(3). The squarks therefore do not feel the strong nuclear force of the SM, but they do interact with their own separate version of the strong force outside the SM. The "neutral" in neutral naturalness refers to the squarks' neutrality under the SM's strong nuclear force.

To maintain the symmetry arguments that explain the natural Higgs mass, the squarks' SU(3) must have approximately the same confinement scale as the SM SU(3) ($\gtrsim$ 1 GeV) [10]. Because collider data has excluded squark masses below about 100 GeV in neutral naturalness (see Sec. 2.1), all the squarks have larger masses than the confinement scale. These large masses define a "quirky" sector [11], so we call them squirks instead of squarks.

In this paper, the squirks I consider carry electroweak charge, and I focus on squirk pairs with net electric charge. When these pairs annihilate, they produce a W boson and either a photon γ or a Z boson. These bosonic pairs are also known as resonances, and previous searches by ATLAS and CMS [12–17] have already looked for $W\gamma$ and WZ resonances. They did not discover squirks in these searches, which places restrictions on the possible masses squirks could have. Interestingly, these bounds on squirk masses decrease as Δ, the mass difference between the squirks in the pair, increases. So, these searches have not adequately explored the large-Δ portion of squirk mass parameter space (see Sec. 2.1). Squirks with masses in this region of parameter space are light enough that they could be produced by the current collision energies of the LHC. So, with a better search strategy, we could explore more of squirk parameter space without upgrading the LHC's collision energy, thereby increasing the discovery potential of the LHC.

Knowing that discovering squirks would help explain the Higgs mass and knowing that squirks could be accessible at current LHC energies, we created our strategy [5] to search more effectively for squirks. Our search strategy combines searching for a $W\gamma$ or WZ resonance with searching for a displaced vertex (DV), which is when a particle decays away from the initial collision point, making the detected particles appear to have originated from a different point inside the detector. In our model, the DV is created by the decay of a glueball, which is part of the quirky sector. Because the quirky sector is defined by an SU(3) gauge symmetry, it contains gluons that behave very similarly to the gluons in the SM's SU(3) sector. Throughout this paper, when I refer to gluons I mean the gluons in the quirky sector. Where the distinction might be unclear, I explicitly include "SM" when talking about gluons in the SM, and I often describe the gluons in the quirky sector as "hidden" because they are not part of the SM. A glueball is simply a bound state of gluons, and glueballs can decay into SM particles by mixing with the Higgs (see Sec. 2.2). This decay has a long lifetime on the scales relevant to particle colliders, which is how the glueball gets far enough away from the collision point to form a DV when it decays.

To predict the performance of this resonance + DV search strategy, I ran simulations of proton-proton collisions at the LHC and applied the strategy to the data. Based on the simulations, I found that compared to previous searches, our strategy covers more squirk mass parameter space when glueballs decay within the detector. Conversely, when glueballs decay outside the detector, our strategy fails. Since our strategy depends heavily on the glueball decay, it cannot replace current methods, but it can extend their reach.

The rest of the paper proceeds as follows. In Chap. 2, I discuss background information on squirks and glueballs, and in Chap. 3, I discuss background information on the ATLAS detector. Both these chapters provide important context for the simulation, which I discuss in Chap. 4. I present the results in Chap. 5 and interpret them in Chap. 6. Finally, I give some concluding thoughts in Chap. 7, including a note on future work.

Chapter 2

Squirk Model

In this chapter, I describe the important parameters and quantities in the squirk model we studied. Section 2.1 describes the squirks themselves, and Sec. 2.2 describes glueballs, the bizarre particles key to the success of our search strategy. In these two sections, I cover much of the background information required to understand the methods used to simulate our search strategy as described in Chap. 4.

2.1 Squirks

The squirks I consider in this work form an electroweak doublet. One squirk is an up-type symmetry partner and therefore has an electric charge of $+2/3$. The other is a down-type symmetry partner with electric charge $-1/3$. I'll call these the up squirk and the down squirk respectively and use the symbols $\tilde{u}$ and $\tilde{d}$. With these two squirks, there are four possible squirk-antisquirk pairs: up-antiup, down-antidown, up-antidown, and down-antiup. In this work I focus on the latter two pairs because they have net electric charge. The up-antidown bound state has charge $+1$, so I call it Ψ_+. Similarly, the down-antiup bounds state has charge -1, so I call it Ψ_-. I refer to these two states collectively as $\Psi_\pm$.

Letting the up squirk mass be $m_{\tilde{u}}$ and the down squirk mass be $m_{\tilde{d}}$, I characterize the mass of a $\Psi_\pm$ state using two parameters: the total mass M and the mass difference Δ :

$$M = m_{\tilde{u}} + m_{\tilde{d}} \quad \text{and} \quad \Delta = m_{\tilde{u}} - m_{\tilde{d}} \; . \tag{2.1}$$

In using M and Δ, I assume that the binding energy of a $\Psi_\pm$ state is negligible. As we found in [5], these parameters have an important relationship to $\tan(\beta)$, a parameter in our meutral naturalness model. The relationship is

$$M\Delta = m_W^2 \frac{1 - \tan^2(\beta)}{1 + \tan^2(\beta)} \tag{2.2}$$

where m_W is the mass of the W boson. Because $\tan(\beta)$ must be a positive finite real number, Eq. 2.2 gives a constraint on Δ for a given M:

$$\frac{-m_W^2}{M} < \Delta < \frac{m_W^2}{M} \; . \tag{2.3}$$

I will refer to both Δ and $\tan(\beta)$ throughout this thesis, but because $\tan(\beta)$ is independent of M, I will use $\tan(\beta)$ rather than Δ when plotting.

The $\Psi_\pm$ state is not stable and will eventually annihilate to a pair of bosons. Conservation of electric charge limits the annihilation to the following channels:

$$\Psi_\pm \to W^\pm \gamma \quad \text{and} \quad \Psi_\pm \to W^\pm Z \; . \tag{2.4}$$

Since there are two annihilation channels, it is useful to know how frequently they occur. This frequency is given by a branching ratio (BR), sometimes called a branching fraction. By definition, the BR is a number between 0 and 1 giving the probability that a decay or an annihilation occurs through a particular channel. To calculate a BR, we need another quantity called a decay width. Even though we are working with an annihilation, by convention we still call it a “decay width” and not an “annihilation width” because the bound state (as a single object) is decaying. The decay width Γ of the bound state annihilation is defined relative to the half-life $t_{1/2}$ of that state's annihilation, or the time it takes for half of the states to annihilate. The relationship is

$$\Gamma = \frac{\hbar \ln(2)}{t_{1/2}} \ . \tag{2.5}$$

Each annihilation channel has its own decay width. Using Γ_i as the decay width of the ith annihilation channel, the BR for the ith channel is

$$\mathrm{BR}_i = \frac{\Gamma_i}{\sum_{j=1}^{n} \Gamma_j} \ , \tag{2.6}$$

where n is the total number of annihilation channels. For the squirk pair annihilation I consider in this work, the branching ratios are therefore

$$\mathrm{BR}(\Psi_\pm \to W^\pm\gamma) = \frac{\Gamma(\Psi_\pm \to W^\pm\gamma)}{\Gamma(\Psi_\pm \to W^\pm\gamma) + \Gamma(\Psi_\pm \to W^\pm Z)} \tag{2.7}$$

$$\mathrm{BR}(\Psi_\pm \to W^\pm Z) = \frac{\Gamma(\Psi_\pm \to W^\pm Z)}{\Gamma(\Psi_\pm \to W^\pm\gamma) + \Gamma(\Psi_\pm \to W^\pm Z)} \tag{2.8}$$

The decay widths to the channels given in Eq. 2.4 are calculated from the following formula adapted from Ref. [18]:

$$\Gamma(\Psi_\pm \to W^\pm V) = \frac{3|R(0)|^2}{32\pi^2 M^2} \sqrt{\lambda(M; m_W, m_V)} \sum |\mathcal{M}(\Psi_\pm \to W^\pm V)|^2 \ , \tag{2.9}$$

where V is a placeholder for either a photon γ or a Z boson, m_V is the mass of the V boson, R(0) is the radial part of the $\Psi_\pm$ wavefunction evaluated at the origin, and $\mathcal{M}$ is a matrix element calculated in quantum field theory. The λ function is

$$\lambda(M; m_1, m_2) = \left(1 - \left(\frac{m_1 + m_2}{M}\right)^2\right) \left(1 - \left(\frac{m_1 - m_2}{M}\right)^2\right) \ . \tag{2.10}$$

It shows up naturally when considering the momenta of the decay products when a particle of mass M decays into two particles, one of mass m_1 and one of mass m_2 (see Appendix A).

As we found in [5], the $W\gamma$ decay width is

$$\begin{aligned}\Gamma(\Psi_\pm \to W^\pm\gamma) &= \frac{3\alpha_W^2 s_W^2 |R(0)|^2}{2M^2} \left(1 - \frac{m_W^2}{M^2}\right) \left(\frac{1}{3} - \frac{\Delta}{M}\right)^2 \\ &= \frac{3\alpha_W^2 s_W^2 |R(0)|^2}{2M^2} \left(1 - \frac{m_W^2}{M^2}\right) \left(\frac{1}{3} - \frac{m_W^2}{M^2} \frac{1 - \tan^2(\beta)}{1 + \tan^2(\beta)}\right)^2 \ ,\end{aligned} \tag{2.11}$$

where $\alpha_W = g^2/(4\pi) \approx 0.0339$ is the SM weak sector coupling and $s_W \approx 0.481$ is the sine of the weak mixing angle. I obtained these numerical values from the Particle Data Group (PDG) [19]. The WZ decay width is

$$\begin{aligned}
\Gamma(\Psi_\pm \to W^\pm Z) = {} & \frac{3\alpha_W^2 |R(0)|^2 s_W^4}{2c_W^2 M^2}\sqrt{\lambda(M; m_W, m_Z)}\left\{\left(\frac{1}{3} - \frac{\Delta}{M}\right)^2\right. \\
& + \frac{M^4}{8m_Z^2 m_W^2}\left[\left(1 - \frac{\Delta^2}{M^2}\right)\lambda(M; m_W, m_Z) A_{WZ}\right. \\
& \left.\left. - \left(\frac{1}{3} - \frac{\Delta}{M}\right)\left(1 - \frac{m_Z^2 + m_W^2}{M^2}\right)\right]^2\right\} \\
= {} & \frac{3\alpha_W^2 |R(0)|^2 s_W^4}{2c_W^2 M^2}\sqrt{\lambda(M; m_W, m_Z)}\left\{\left(\frac{1}{3} - \frac{m_W^2}{M^2}\frac{1 - \tan^2(\beta)}{1 + \tan^2(\beta)}\right)^2\right. \\
& + \frac{M^4}{8m_Z^2 m_W^2}\left[\left(1 - \left(\frac{m_W^2}{M^2}\frac{1 - \tan^2(\beta)}{1 + \tan^2(\beta)}\right)^2\right)\lambda(M; m_W, m_Z) A_{WZ}\right. \\
& \left.\left. - \left(\frac{1}{3} - \frac{m_W^2}{M^2}\frac{1 - \tan^2(\beta)}{1 + \tan^2(\beta)}\right)\left(1 - \frac{m_Z^2 + m_W^2}{M^2}\right)\right]^2\right\} ,
\end{aligned} \tag{2.12}$$

where c_W is the cosine of the weak mixing angle and

$$\begin{aligned}
A_{WZ} = {} & \frac{\frac{1}{3}\left(1 - \frac{m_Z^2 + m_W^2}{M^2}\right) + \frac{c_W^2 \Delta}{s_W^2 M}\left(1 + \frac{m_Z^2 - m_W^2}{M^2}\right)}{\left(1 - \frac{m_Z^2 + m_W^2}{M^2}\right)^2 - \frac{\Delta^2}{M^2}\left(1 + \frac{m_Z^2 - m_W^2}{M^2}\right)^2} \\
= {} & \frac{\frac{1}{3}\left(1 - \frac{m_Z^2 + m_W^2}{M^2}\right) + \frac{c_W^2}{s_W^2}\frac{m_W^2}{M^2}\frac{1 - \tan^2(\beta)}{1 + \tan^2(\beta)}\left(1 + \frac{m_Z^2 - m_W^2}{M^2}\right)}{\left(1 - \frac{m_Z^2 + m_W^2}{M^2}\right)^2 - \left(\frac{m_W^2}{M^2}\frac{1 - \tan^2(\beta)}{1 + \tan^2(\beta)}\right)^2\left(1 + \frac{m_Z^2 - m_W^2}{M^2}\right)^2} .
\end{aligned} \tag{2.13}$$

Plugging the widths from Eqs. 2.11 and 2.12 into Eqs. 2.7 and 2.8 gives the BRs as a function of M and $\tan(\beta)$. Figure 2.1 plots these BRs as a function of $\tan(\beta)$ for a few representative values of M. Note that near $\tan(\beta) = 1$ (equivalently $\Delta = 0$), the $W\gamma$ channel dominates, but as $\tan(\beta)$ approaches 0 or infinity, the WZ channel becomes more prominent until it eventually surpasses the $W\gamma$ channel. I therefore expect that the $W\gamma$ portion of the search strategy will perform better for $\tan(\beta) \approx 1$ while the WZ portion of the search strategy will perform better for $\tan(\beta)$ far from 1. Also note that there is an approximate

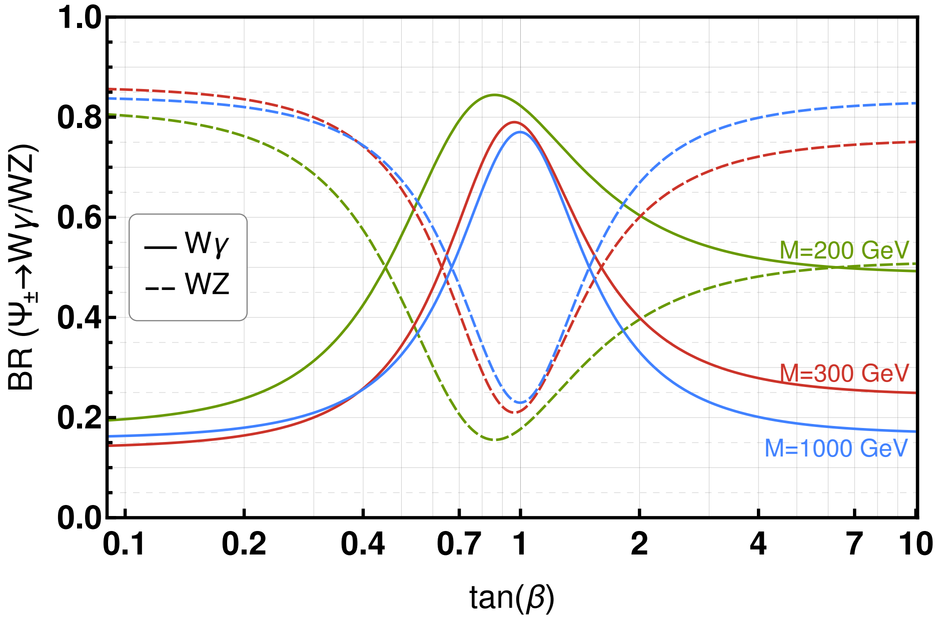

Figure 2.1 Branching ratios of a squirk bound state $\Psi_\pm$ to $W^\pm\gamma$ (solid) and $W^\pm Z$ (dashed) as a function of tan(β) for a few representative values of the bound state mass M. Green is for $M = 200$ GeV, red is for $M = 300$ GeV, and blue is for $M = 1000$ GeV. The $W\gamma$ signal dominates for tan(β) ≈ 1, and the WZ signal takes over as tan(β) approaches 0 or infinity. The turnover occurs closer to tan(β)=1 for larger values of M, and the log-scale asymmetry across tan(β)=1 is more pronounced for small M.

log-scale symmetry across tan(β) = 1 for larger M, but the plot becomes more asymmetric as M decreases. This symmetry will appear again in Fig. 2.2.

Detectors like ATLAS and CMS at the LHC have looked for $W\gamma$ and WZ pairs in previous searches. We can use these searches to see if any combinations of M and tan(β) have already been ruled out. By "ruled out", I mean that if squirks existed with that particular mass combination, the detectors would have already found them. Since they haven't found anything, we can rule out those mass combinations. Ruled-out mass combinations are often called "excluded", and I'll use that word frequently to describe ruled-out portions of parameter space.

To find the regions of parameter space excluded by experimental searches, I began by simulating the production of a $\Psi_{\pm}$ state using MadGraph [20]. MadGraph gives us a cross section σ for the production. The cross section is a number with dimensions of area that quantifies how likely a given process is to occur. Larger cross sections mean easier processes, or processes that are more likely to occur in a collider experiment. Cross sections are technically calculable by hand, but MadGraph is a simulation software designed to compute cross sections numerically and is therefore much easier to use. By running the MadGraph simulation of $\Psi_{\pm}$ production for several combinations of M and $\tan(\beta)$, I got the cross section as a function of these parameters.

Once I had the cross sections, I multiplied them by the branching ratio of the squirk pair into either $W\gamma$ or WZ. Depending on the search I wanted to compare my results to, I also multiplied by another branching ratio for the W or Z decays. For example, if a WZ search looked only for leptonic decays of the Z boson, then I'd multiply the cross section times the WZ branching ratio times the BR for Z decays into leptons like electrons and muons.

With these multiplications, I finally had an apples-to-apples comparison with the experimental data. If the cross section times the branching ratios is larger than the reported detector sensitivity, then the detector would have already seen the particles if they existed with that mass combination. If the value is lower than the reported sensitivity, then it is still possible that the particles could exist with that mass combination because the detectors aren't sensitive enough yet to see them. The exclusion condition can be summarized by the following inequality:

$$\sigma_{\mathrm{MG}}(M, \tan(\beta)) \cdot \mathrm{BR}_{WV}(M, \tan(\beta)) \cdot \prod_{\text{decays}} \mathrm{BR}(W \text{ or } Z) > \sigma_{\text{exp}} \tag{2.14}$$

where σ_{MG} is the cross section from MadGraph, BR_{WV} is the branching ratio to either $W\gamma$ or WZ as needed, the product over decays is for including W and Z decay branching ratios as I described above, and σ_{exp} is the experimental sensitivity reported by ATLAS or CMS. If the inequality in Eq. 2.14 is satisfied for a given combination of M and $\tan(\beta)$, then that combination is excluded by experiment.

I made comparisons to several searches [12–17] to determine the existing bounds on M and $\tan(\beta)$. The exclusion regions are plotted in Fig. 2.2. I found the red exclusion region by comparing simulations with Ref. [12], a $W\gamma$ search done by ATLAS at an 8 TeV collision energy. The green region comes from comparisons with Refs. [15, 16], $W\gamma$ searches done by CMS at a 13 TeV collision energy. The blue dashed region is a projection of these CMS exclusions into the HL-LHC. To make this projection, I first multiplied the CMS expected sensitivity by $\sqrt{138/3000}$, converting the expectation from the experimental luminosity of 138 fb^{-1} to the HL-LHC luminosity of 3000 fb^{-1}. Then I compared this new HL-LHC expectation with the simulation data. Notice that the figure does not contain any exclusions from the WZ searches [13, 14, 17]. These exclusion regions were all contained within the $W\gamma$ exclusions, so I did not include them in the plot. Also notice that the same approximate log-scale symmetry across $\tan(\beta) = 1$ that we saw in Fig. 2.1 appears here.

Figure 2.2 shows that for a given value of M, exclusions from current search methods weaken as $\tan(\beta)$ approaches 0 or infinity. These weakening bounds drove us to develop the search strategy we published in [5]. Our search strategy aims to explore the regions in Fig. 2.2 not covered by the blue region so that we can improve the expected reach of the HL-LHC into squirk mass parameter space.

2.2 Glueballs

To explore a larger region of squirk mass parameter space, our search strategy relies on a particle called a glueball, which creates a displaced vertex in the same event as a $W\gamma$ or WZ resonance. This section describes what glueballs are and how displaced vertices contribute to our search strategy.

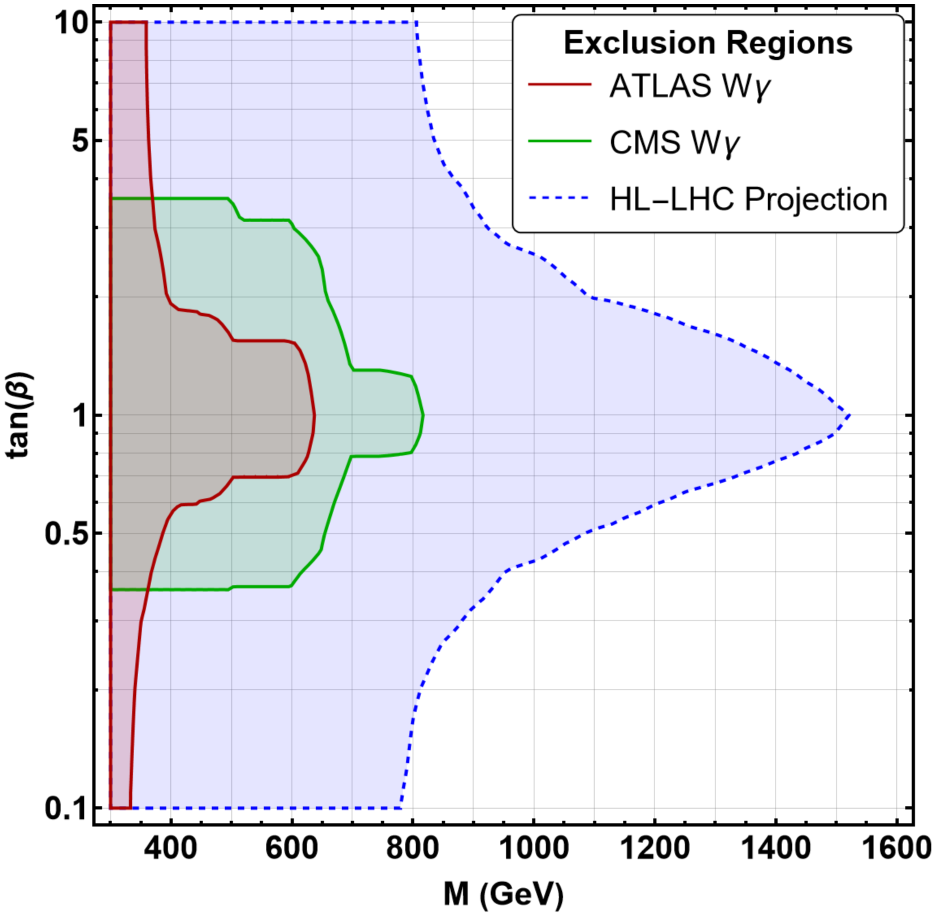

Figure 2.2 Experimental bounds on M and $\tan(\beta)$ from several LHC searches. The red region is excluded by the ATLAS $W\gamma$ search [12]. The green region is excluded by the CMS $W\gamma$ searches [15, 16]. The blue region is the projected exclusion for the HL-LHC based on the CMS exclusions in green. Any exclusion regions from other searches are contained within these regions and are therefore not included in the plot.

Within the model we consider, the squirks are charged under a new SU(3) gauge symmetry. Under this symmetry, the new gauge fields behave analogously to the gluon fields in the SM. For this reason, we call them hidden gluons. Since gluons are charged under their own gauge symmetry, they can interact with each other and even form bound states. A bound state made purely of gluons is called a glueball. The glueballs in the squirk sector are the lightest bound states of that sector because the squirk masses are so large compared to the confinement scale. The lightest of these glueballs has J^{PC} quantum numbers 0^{++}, meaning spin-0, even under parity, and even under charge conjugation respectively. Interestingly, these same quantum numbers describe the Higgs boson, which means that the 0^{++} hidden glueball can mix with the Higgs, allowing the 0^{++} to decay into SM particles through an off-shell Higgs.

The decay width for this process is

$$\Gamma(0^{++} \to X_{\mathrm{SM}} X_{\mathrm{SM}}) = |c_g|^2 \left(\frac{\alpha'}{6\pi} \frac{f_{0^{++}}}{v(m_h^2 - m_0^2)} \right)^2 \Gamma\left(h(m_0) \to X_{\mathrm{SM}} X_{\mathrm{SM}}\right) \ , \tag{2.15}$$

which is derived from [21] as explained in [5]. This equation has a lot to unpack. The placeholders $X_{\mathrm{SM}} X_{\mathrm{SM}}$ stand in for any two SM particles. The parameter c_g is related to the coupling of hidden gluons to the Higgs. As we found in the erratum of our paper [5], $1/16$ is a useful upper bound for c_g, so we consider it a free parameter within the range $0 < c_g \leq 1/16$. The parameter α' is the coupling of hidden gluons to squirks, v is the Higgs VEV, and m_h is the Higgs mass. The parameter m_0 is the mass of the 0^{++} glueball, which we take as a free parameter in our model. The decay width on the right is the decay width of an off-shell Higgs with mass m_0 to the two SM particles. We calculated these decay widths using HDECAY [22].

The matrix element $f_{0^{++}}$ was found in lattice calculations [23] to be $4\pi\alpha' f_{0^{++}}/m_0^3 = 2.69 \pm 0.27$. See Ref. [5] for the derivation of this result from the result presented in Ref. [23]. In our simulations, we replace $f_{0^{++}}$ with its equivalent value in terms of α' and m_0, conveniently removing alpha from the formula:

$$\Gamma(0^{++} \to X_{\mathrm{SM}} X_{\mathrm{SM}}) = |c_g|^2 \left(\frac{2.69}{24\pi^2} \frac{m_0^3}{v(m_h^2 - m_0^2)} \right)^2 \Gamma\left(h(m_0) \to X_{\mathrm{SM}} X_{\mathrm{SM}}\right) \ . \tag{2.16}$$

When working with squirk decay widths earlier, I used them to compute branching ratios. Beyond branching ratios, decay widths are also useful for computing the characteristic lifetime τ of an unstable state. τ is related to the half-life introduced earlier:

$$\tau = \frac{t_{1/2}}{\ln(2)} \ , \tag{2.17}$$

so

$$\tau = \frac{\hbar}{\Gamma} \ . \tag{2.18}$$

The lifetime of a 0^{++} hidden glueball depends on the free parameters m_0 and c_g. For some combinations of these parameters, the lifetime is large enough that the glueball can travel a noticeable distance within the detector before it decays. In these cases, we call the glueball a long-lived particle. Of course, the lifetime is still extremely short by human standards, but it is long enough to be noticeable on collider scales. In the detector, the long-lived decay of a hidden glueball creates a displaced vertex (DV): “displaced” because it is far from the initial collision and “vertex” because one particle becomes two. Since the hidden glueball is part of the squirk sector, it does not interact with the particles that make up the detector, so it is invisible until it creates the DV by decaying into SM particles.

The quantity $c\tau$ is known as the decay length and gives the characteristic distance an unstable state will travel before decaying. I have plotted the 0^{++} hidden glueball decay length as a function of m_0 and c_g in Fig. 2.3. The blue contours are in meters and show the wide range of values the decay length can have. The red region highlights distances from the original collision where the ATLAS detector would detect a DV. We expect that combinations of c_g and m_0 that lead to decay lengths in or near this region will lead to the highest rates of successful DV detection.

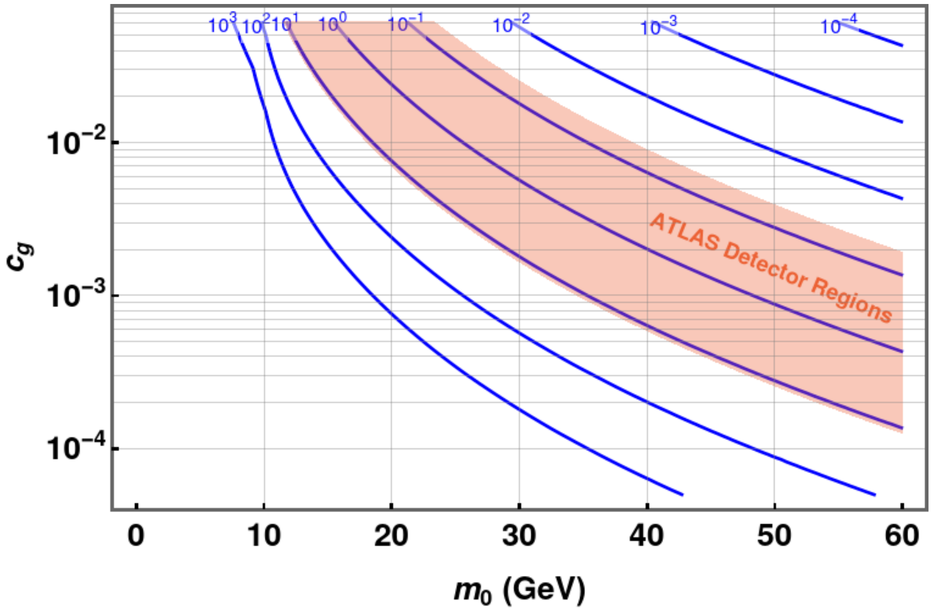

Figure 2.3 0^{++} glueball decay length in meters given by the blue contours as a function of the 0^{++} mass m_0 and the coupling parameter c_g. The red region highlights the distances from the collision point where the ATLAS detector is sensitive to displaced vertices.

Chapter 3

Detection

In this chapter, I outline a few important details necessary to understand the simulation of detection in Chap. 4. Section 3.1 gives an overview of the ATLAS detector subsystems as well as some key coordinates for labeling positions within the detector. In Sec. 3.2, I tabulate the cuts we used in the simulation to accurately represent the way ATLAS acquires data. On a similar note, Sec. 3.3 gives the DV reconstruction efficiencies for the different detector subsystems. These efficiencies allow the simulation to more accurately reflect the way ATLAS detects displaced vertices.

3.1 ATLAS Detector

The ATLAS detector is divided into four subsystems: the inner detector (ID) or tracker, the electromagnetic calorimeter, the hadronic calorimeter, and the muon system. I will refer to these as the Tracker, ECAL, HCAL, and MS respectively. These subsystems are each cylindrically symmetric about the beam axis to a good approximation. Each cylinder is nested inside the next, with the MS being the outermost. The detectors also have a mirror symmetry about the interaction point (IP), which is commonly used to define the origin of the

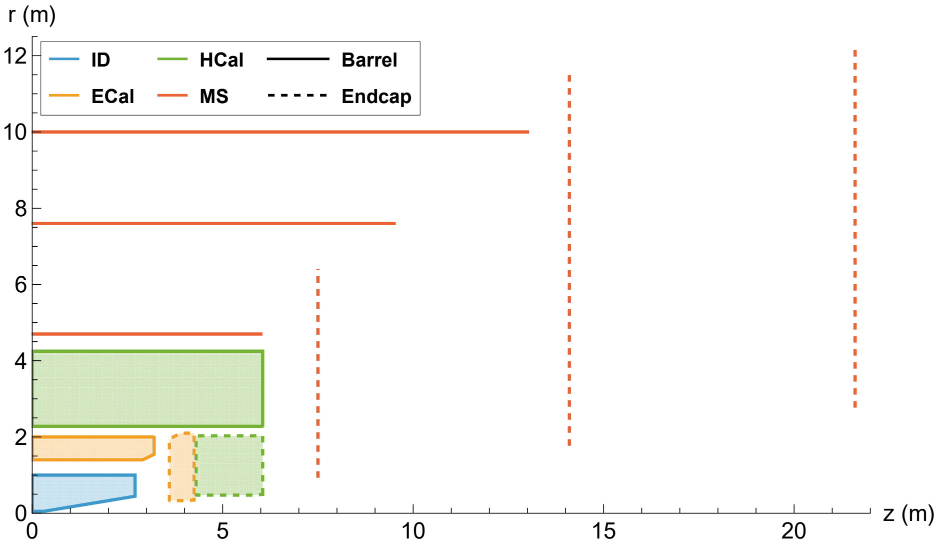

Figure 3.1 Simple characterization of the ATLAS detector subsystems.

collider coordinate system. The z axis is along the beam axis, so the detector subsystems are generally characterized by a cylindrical radius r and a distance along the beam axis $|z|$. The remaining cylindrical coordinate is the angular coordinate ϕ, which is useful in our simulation when we need to determine the angular separation of two particles, but it is not useful in characterizing the regions of the ATLAS detector. See Fig. 3.1 for my characterization of the ATLAS detector subsystems. Note that for every subsystem but the Tracker, I have distinguished between the barrel and endcap of the subsystem. This distinction will be important in Sec. 3.3.

Another important collider coordinate is called the pseudorapidity η. This is an angular coordinate related to the polar angle θ in spherical coordinates. The relationship is:

$$\eta = -\ln\left[\tan\left(\frac{\theta}{2}\right)\right] . \tag{3.1}$$

Since particles usually travel in straight lines from the IP, η is used to describe the trajectory's angle from the beam axis. This coordinate is often used in cuts for two reasons. First, particles with larger $|\eta|$ will go right back down the beam pipe instead of hitting the detectors, so there's no reason to keep those events. Second, particles with small $|\eta|$ have somehow ended up traveling perpendicular to the beam particles, which indicates interesting physics worth investigating.

3.2 Cuts

To make our simulation as realistic as possible, we looked at searches to determine what kind of cuts they made on the data to remove the SM background. We then applied those cuts in the simulation. For example, if the event's W boson didn't have enough energy to pass the cut, then that event was not counted as detected in the simulation. Making cuts is extremely important for making sure the data is statistically significant. Well-reasoned cuts can remove lots of SM background without removing a lot of the desired signal, which means even a little bit of the signal can be significant enough to indicate new physics. The cuts I used in the simulation are given in Tables 3.1-3.3. These tables were copied from our publication [5], and a description of the various terms and values can be found therein.

Quantity	**Leptonic Cuts**	**Hadronic Cuts**
p_T (GeV)	$p_T^e > 35$ $p_T^\mu > 30$	$p_T^J > 200$
p_T^γ (GeV)	$0.4\, m_T^{\ell\gamma} < p_T^\gamma < 0.55\, m_T^{\ell\gamma}$	$p_T^\gamma > 200$
p_T^{miss} (GeV)	$p_T^{\text{miss}} > 40$	-
$\eta^{\ell,J}$	$\eta^e < 2.1$ and $\eta^e \notin [1.37, 1.52]$ $\eta^\mu < 2.5$	$\eta^J < 2.0$
η^γ	$\eta^\gamma < 1.37$	

Table 3.1 Cuts for the $W\gamma$ resonance, including both leptonic and hadronic decays of the W.

Quantity	**Leptonic Cuts**	**Hadronic Cuts**
p_T (GeV)	$p_T^\ell > 25$ $p_T^\ell > 27$ (trigger-matched)	$p_T^{J_1} > 500$ (leading) $p_T^{J_2} > 200$ (subleading)
p_T^{miss} (GeV)	$p_T^{\text{miss}} > 25$	-
E_T^{miss} (GeV)	$E_T^{\text{miss}} > 25$	-
η^ℓ	$\eta^e < 2.47$ and $\eta^e \notin [1.37, 1.52]$ $\eta^\mu < 2.7$	$\eta^J < 2.0$
Other	$\mathcal{R}_{p_T} > 0.35$	-

Table 3.2 Cuts for the WZ resonance when both particles decay leptonically or both decay hadronically.

Quantity	**0-lepton Cuts**	**1-lepton Cuts**		**2-lepton Cuts**	
		merged	resolved	merged	resolved
R	-	$R < 1.0$	$R \geq 1.0$	$R < 1.0$	$R \geq 1.0$
p_T^ℓ (GeV)	-	$p_T^\ell > 30$			
p_T^{miss} (GeV)	$p_T^{\text{miss}} > 50$	$p_T^{\text{miss}} > 25$		-	
E_T^{miss} (GeV)	$E_T^{\text{miss}} > 250$	$E_T^{\text{miss}} > 100$	$E_T^{\text{miss}} > 60$	-	
p_T^J (GeV)	$p_T^J > 200$	$p_T^J > 200$	$p_T^{J_1} > 60$ $p_T^{J_2} > 45$	$p_T^J > 200$	$p_T^{J_1} > 60$ $p_T^{J_2} > 30$
η^J	$\eta^J < 2.0$	$\eta^J < 2.0$	$\eta^J < 4.5$	$\eta^J < 2.0$	$\eta^J < 2.5$
η^ℓ	-	$\eta^e < 2.47$ and $\eta^e \notin [1.37, 1.52]$ $\eta^\mu < 2.5$			
p_T^W (GeV)	-	$p_T^W > 200$	$p_T^W > 75$	-	
Other	$\mathcal{R}_{p_T} > 0.35$				

Table 3.3 Cuts for the WZ resonance when one particle decays leptonically and the other decays hadronically.

3.3 Vertex Reconstruction Efficiencies

Because our strategy relies on detecting a displaced vertex, we need to know how well the ATLAS detector can find a DV. Each detector subsystem has an efficiency for reconstructing DVs. A subsystem reconstructs a DV by tracing back the energy deposition until it pinpoints the location where the original particle decayed. ATLAS runs simulations to determine how often the reconstruction algorithms successfully find the location of the decay. Thus a subsystem's DV reconstruction efficiency is the fraction of simulated DVs that the subsystem successfully reconstructed. In my simulation, I use these efficiencies to make the detection portion of the simulation more realistic by allowing for some unsuccessful DV reconstruction (see Sec. 4.5). By taking these efficiencies into account, I make sure that my simulation does not overestimate the number of DVs detected.

The DV reconstruction efficiencies I used in my simulation are given in Table 3.4. Notice that the table does not include the ECAL. When detecting displaced vertices, the calorimeters ECAL and HCAL work together. Since the HCAL is the second of the two calorimeters that the particles will pass through, it will record more energy than the ECAL if the DV occurs within the calorimeters. Thus the HCAL stands in for both calorimeters when we consider DV detection.

Subsystem	Tracker	HCAL Barrel	HCAL Endcaps	MS Barrel	MS Endcaps
Efficiency	0.1	0.22	0.06	0.2	0.4

Table 3.4 DV reconstruction efficiencies for the different ATLAS subsystems used in my simulation.

Chapter 4

Simulation Methods

If we hope to recommend that the experiments at the LHC should implement our new search strategy, we need to show that we expect it to perform well and why. To determine how our strategy would perform at the LHC, I ran simulations of proton-proton collisions in the ATLAS detector and applied our search strategy to the simulated data. In this chapter, I describe these simulations in detail. Then, in Chap. 5, I discuss the results of the simulation, which reveal large portions of parameter space where our strategy complements current methods and even some portions where our strategy outperforms current methods.

Our search strategy looks for a bosonic resonance and a displaced vertex in the same event. For each event, I began by simulating the initial proton-proton collision at the center of the ATLAS detector. This collision produces a pair of squirks for the bosonic resonance and a hidden gluon for the displaced vertex. Section 4.1 discusses how I simulated the production of these particles.

The squirks form a bound state after production. This state must de-excite before it can annihilate into the $W\gamma$ or WZ resonance we're searching for. However, if the state beta decays before it fully de-excites, it will no longer annihilate into a $W\gamma$ or WZ pair, making it useless for our search strategy. I treat de-excitation and beta decay in Sec. 4.2.

After production, the hidden gluon hadronizes into a glueball. Only the 0^{++} glueball is useful in our search strategy, so my simulation of hadronization also included calculating the probability that the glueball was a 0^{++}. I discuss this portion of my simulation in Sec. 4.3. Once the glueball has formed, I can simulate the annihilation of the squirk pair, which I describe in Sec. 4.4.

After annihilation, the photon reaches the detector, but the W and Z bosons decay before they can get there. Detecting the bosonic resonance thus requires both direct detection of the photon and indirect detection of the W and Z via their decay products. Section 4.5 outlines my simulation of W and Z decays and my simulation of detecting the bosonic resonance using the cuts described in Sec. 3.2. Section 4.5 also describes my simulation of detecting the displaced vertex using the DV reconstruction efficiencies outlined in Sec. 3.3.

I conclude this chapter in Sec. 4.6 by describing how I convert the results of many simulated events into a prediction of the number of events that would be observed in an experiment. These predictions are then presented in Chap. 5.

4.1 Production

In this first step of the simulation, my goal was to simulate the production of the squirk pair and the hidden gluon from the initial proton-proton collision. I needed to simulate production for lots of different combinations of M and $\tan(\beta)$ to make sure I explored the parameter space. For each parameter combination, I needed the cross section and a list of simulated events. I needed the cross section because I was going to use it in a later step to convert simulated results to predictions of experimental results (see Sec. 4.6). I needed the simulated events because I was going to apply our search strategy on them to predict its performance at the LHC.

I used the software MadGraph [20] to perform these simulations and obtain the cross sections and events that I needed. My MadGraph script begins as follows:

```
import model MSSM_SLHA2

define su = ul ur

define su~ = ul~ ur~

define sd = dl dr

define sd~ = dl~ dr~

generate p p > w+ > su sd~ g / u u~ d d~ c c~ s s~

add process p p > w- > su~ sd g / u u~ d d~ c c~ s s~
```

The first line tells MadGraph to work with the MSSM_SLHA2 model, which contains the squirks. The next four lines define the squirks: up, anti-up, down, and anti-down. For our search strategy, squirks can be symmetry partners of either left- or right-handed fermions, which is why the `define` lines include both an `l` and an `r`. The final two lines tell MadGraph to simulate the processes of two protons to $\Psi_{\pm}$ and a gluon. I had to be quite explicit in these lines to make sure MadGraph treated the squirks correctly. The `/ u u~ d d~ c c~ s s~` at the end of each line is an exclusion ensuring that no SM quarks will show up as intermediate particles. This is extremely important because squirks do not interact directly with SM quarks. As an added measure, I required an intermediate W boson with the appropriate charge using `> w+/- >`. With these lines, I made sure that MadGraph simulated only the processes allowed in our squirk model.

Aside from a few lines of file i/o, the rest of my script contains instructions for runs. Each run is a simulation of many events all using the same parameters. Here's an example block of code with instructions for one run:

```
launch
```

```
analysis=OFF

set nevents 35000

set ebeam 6500

set msu1 100.0

set msd1 100.0

set run_tag M200_TB1

done
```

Let's walk through these lines and why I needed them. The lines `launch` and `done` are the bookends of the instructions for one run. Since I would be analyzing the output myself, I turned off MadGraph's analysis to help it run faster. I simulated 35,000 events for each run, which is enough to get statistically significant data but not so much that the simulation would take too long to run. I set each beam's energy to 6,500 GeV, creating a CoM energy of 13 TeV. I needed this CoM energy so that I could compare the simulation results to the projected exclusion regions for the HL-LHC at 13 TeV as shown earlier in Fig. 2.2. I calculated the up and down squirk masses `msu1` and `msd1` from M and $\tan(\beta)$ and defined the run tag to remind me later which values of M and $\tan(\beta)$ I used for this run.

As I mentioned before, I needed to explore the parameter space by considering lots of combinations of M and $\tan(\beta)$. So, I repeated this block of code several times by taking M from 200 GeV to 1600 GeV in steps of 50 and taking $\tan(\beta)$ from 0.1 to 1 in steps of 0.1 and from 2 to 10 in steps of 1 for each M-value. Since MadGraph was not simulating anything involving the hidden glueball, I did not include specific values for c_g or m_0 at this point in the simulation.

Each run produced the cross section and the list of events that I needed for a particular parameter combination. Since I discussed the significance of cross sections in Sec. 2.1, I'll focus on the events here. In MadGraph, each event is recorded as a list of incoming and outgoing particles. Each particle is recorded as a list of key numbers describing that particle. Table 4.1 contains an example of the outgoing particle lines from one of the simulated events. The first column is the Particle Data Group Identification number (PDGID), which tells us the type of particle described in the line. In Table 4.1, the first two particles are the squirks and the third particle is the gluon. The first set of ellipses in the table contains bookkeeping information that was not useful in my simulation, so I saved space in the table by not copying it. The next four columns are the components of the particle's four-momentum in the lab frame. The m column gives the mass of the particle. The second set of ellipses contains more information that was not useful in my simulation.

PDGID	...	p_x (GeV)	p_y (GeV)	p_z (GeV)	E (GeV)	m (GeV)	...
1000002	...	-144.215	+241.102	+502.163	584.034	100.0	...
-1000001	...	+154.608	-244.130	+198.610	364.622	100.0	...
21	...	-10.393	+3.028	+13.989	17.688	0.0	...

Table 4.1 Example of outgoing particle lines from a simulated event. The first column lists the ID of the particle described by that row. The columns p_i are the components of the particle's lab-frame three-momentum. E is the particle's lab-frame energy, and m is the particle's mass. Ellipses indicate information in the output file not used in my simulation.

Once I obtained cross sections and events from MadGraph by running the script I described in this section, I began applying the search strategy to the events as described in the next few sections.

4.2 De-excitation and Beta Decay

After the squirk pair is produced, the squirks are drawn back together by the hidden force. They will oscillate around each other until they have radiated enough energy (remember accelerating charged particles radiate photons) to fall into the ground state of their bound state, analogous to the s-orbital in atomic physics. This process is called de-excitation, and it is only after the squirks have totally de-excited that they can annihilate to produce the bosonic resonance our strategy searches for [11].

In our paper [5], we found that the maximum time it takes for a squirk pair to de-excite to the ground state is approximately

$$t_{\text{de-excite}} \approx \frac{3M^3}{64\pi\alpha\sigma^2} \frac{\left(\frac{E_T}{M} - 1\right)\left(1 - \frac{\Delta^2}{M^2}\right)^2}{1 + \frac{\Delta^2}{M^2}} \ , \tag{4.1}$$

where E_T is the total energy of the produced squirks in their CoM frame, α is the QED coupling, and σ is the string tension related to the confinement scale Λ' by $\sigma \approx 3.6(\Lambda')^2$.

To simulate de-excitation, we boost into the squirks' CoM frame, replace the squirk pair with one particle of mass M and zero momentum, and then boost back into the lab frame to get the lab-frame four-momentum of the de-excited squirk pair. To boost in and out of different reference frames, I used the formulas derived in Appendix B.

One of the issues that arises with a nonzero Δ is that the heavier squirk can beta decay into the lighter squirk, changing a $\Psi_\pm$ into a Ψ_0. This prevents the state from producing a $W\gamma$ or WZ pair, so the state is no longer useful for our search. We account for beta decay by comparing the timescales for beta decay and de-excitation to calculate a probability that the squirk pair de-excites before it beta decays. In our paper, we found the beta decay timescale to be roughly:

$$\tau_\beta \sim \left(10^{-12}\ \text{s}\right)\left(\frac{10\ \text{GeV}}{\Delta}\right)^5 \ . \tag{4.2}$$

The probability that the squirk pair completely de-excites before it beta decays is

$$P_{\text{de-excite}} = e^{-\frac{t_{\text{de-excite}}}{\tau_\beta}} \ . \tag{4.3}$$

In the simulation, I generated a random number between 0 and 1. If the number was greater than $P_{\text{de-excite}}$, then I considered the state beta-decayed and thus no longer useful to the search strategy. Unfortunately, this removed some events that could have otherwise contributed to our signal, but it was more physically accurate to consider the possibility of beta decay.

4.3 Formation of the Glueball

Each event from MadGraph contains two squirks and a hidden gluon as final state particles. Our search strategy looks for events where that gluon hadronizes into a glueball. To simulate the process of hadronization, we begin with conservation of energy and momentum. In the three particles' CoM frame, the total energy is

$$E = \sqrt{M^2 + p_g^2} + p_g \ , \tag{4.4}$$

where we have assumed that the squirk bound state has de-excited so we can treat it as one particle of mass M. p_g is the magnitude of the gluon's three-momentum. When the gluon hadronizes into a glueball, it picks up mass m_0, so the total energy becomes

$$E = \sqrt{M^2 + p_G^2} + \sqrt{m_0^2 + p_G^2} \ , \tag{4.5}$$

where p_G is the magnitude of the glueball's three momentum. We can determine the momentum of the glueball by setting these equations equal to each other:

$$\sqrt{M^2 + p_g^2} + p_g = \sqrt{M^2 + p_G^2} + \sqrt{m_0^2 + p_G^2} \ . \tag{4.6}$$

This equation gives us a constraint on m_0. Notice that $p_G^2 \geq 0$ by definition and that the case $p_G^2 = 0$ gives us the maximum allowed m_0 for a given M and p_g:

$$\sqrt{M^2 + p_g^2} + p_g = M + m_{0,\max} \ . \tag{4.7}$$

Dividing both sides by p_g, we have:

$$\frac{m_{0,\max}}{p_g} = 1 + \sqrt{1 + \frac{M^2}{p_g^2}} - \frac{M}{p_g} \ . \tag{4.8}$$

Since $m_0 \leq m_{0,\max}$, we get:

$$\frac{m_0}{p_g} \leq 1 + \sqrt{1 + \frac{M^2}{p_g^2}} - \frac{M}{p_g} \ . \tag{4.9}$$

For convenience, we define

$$f_m = \frac{m_0}{p_g} \quad \text{and} \quad f_M = \frac{M}{p_g} \ , \tag{4.10}$$

so the inequality becomes

$$f_m \leq 1 + \sqrt{1 + f_M^2} - f_M \ . \tag{4.11}$$

The easiest way to get physical intuition from this inequality is by first considering the large-M limit (where $M \gg p_g$). In this limit, the inequality becomes

$$f_m \leq 1 \ , \qquad\qquad p_g \geq m_0 \ . \tag{4.12}$$

So glueball hadronization requires that the gluon momentum be bigger than the glueball mass. This requirement makes sense because the gluon's momentum gets converted to mass during hadronization, so there needs to be enough momentum p_g to convert into a mass of m_0. If we remove the large-M assumption, we get

$$m_0 \leq p_g + \sqrt{M^2 + p_g^2} - M \ . \tag{4.13}$$

We see from this inequality that m_0 is allowed to be slightly larger than p_g. When we remove the large-M assumption, we allow both the gluon momentum and the squirk momentum to be converted to glueball mass. In this case the gluon no longer needs to carry all the momentum needed for conversion to the glueball mass m_0 because some of the mass comes from the squirk momentum.

Equation 4.6 also gives us a constraint on p_G. Notice that $m_0^2 \geq 0$ by definition and that the case $m_0^2 = 0$ gives us the maximum allowed p_G for a given M and p_g:

$$\sqrt{M^2 + p_g^2} + p_g = \sqrt{M^2 + p_{G,\max}^2} + p_{G,\max} \ , \tag{4.14}$$

$$p_{G,\max} = p_g \ . \tag{4.15}$$

Since $p_G \leq p_{G,\text{max}}$, we have:

$$p_G \leq p_g \ . \tag{4.16}$$

So glueball hadronization requires that the glueball momentum be smaller than the gluon momentum. This requirement makes sense because the glueball picks up mass, and the energy stored in that mass can only come from the momentum of the gluon and squirks. Picking up mass thus requires a decrease in momentum.

Now that we've discussed the physically meaningful constraints from Eq. 4.6, let's solve it! Solving for p_G is analogous to solving for p in Appendix A. Making the appropriate substitutions into the results of Appendix A and rearranging some terms gives the following result for p_G:

$$p_G = p_g \sqrt{\left(1 - \left(\frac{f_m}{1+\sqrt{1+f_M^2} - f_M}\right)^2\right)\left(1 - \left(\frac{f_m}{1+\sqrt{1+f_M^2} + f_M}\right)^2\right)} \ . \tag{4.17}$$

We see that the constraint in Eq. 4.11 keeps p_G real-valued and that the structure inside the square root of Eq. 4.17 keeps $p_G \leq p_g$ as we found in Eq. 4.16. In the simulation, I use Eq. 4.17 to replace the gluon in the event with a glueball having four-momentum

$$p_G^\mu = \begin{pmatrix} \sqrt{m_0^2 + p_G^2} \\ p_G \ \hat{d} \end{pmatrix} \ , \tag{4.18}$$

where $\hat{d}$ is the direction vector of the gluon's three-momentum before hadronization. I also modify the squirk pair's four-momentum to maintain conservation of energy and momentum:

$$p_{\text{squirks}}^\mu = \begin{pmatrix} \sqrt{M^2 + p_G^2} \\ -p_G \ \hat{d} \end{pmatrix} \ . \tag{4.19}$$

Another important note on glueball formation is that the 0^{++} glueball is not the only glueball that the gluon could hadronize into. Using formulas from Ref. [24], we calculate the probability that the gluon hadronizes into a 0^{++} glueball for each event. Using a random number to compare against the probability, we determine which events have a 0^{++} glueball (and thus are useful) and which events do not.

4.4 Annihilation

We simulate annihilation by working in the squirks' CoM frame. There, we assume that their relative velocity is small enough that the energy of the system is simply their combined mass M. We can then treat the annihilation as if it were a decay of one particle of mass M into two particles. The resulting energy and momentum of the decay products in this case is a well-known result of conservation of relativistic momentum. The results are derived in Appendix A. For a particle of mass M decaying into a particle of mass m_1 and a particle of mass m_2, the four-momenta of the decay products are

$$p_1^\mu = \begin{pmatrix} \sqrt{p^2 + m_1^2} \\ p\hat{d} \end{pmatrix} \quad \text{and} \quad p_2^\mu = \begin{pmatrix} \sqrt{p^2 + m_2^2} \\ -p\hat{d} \end{pmatrix} \tag{4.20}$$

where

$$p = \frac{M}{2}\sqrt{\left(1 - \left(\frac{m_1 + m_2}{M}\right)^2\right)\left(1 - \left(\frac{m_1 - m_2}{M}\right)^2\right)} \tag{4.21}$$

and $\hat{d}$ is a direction vector chosen randomly for each annihilation in our simulation. Since we know the mass M of the squirk pair and the masses of the annihilation products $W^\pm$, Z, and γ, simulating annihilation with the equations above is quite straightforward. Once we chose a random direction vector $\hat{d}$, we calculate the four-momentum of annihilation products and replace the squirks in the event with the $W\gamma$ or WZ pair.

4.5 Detection

Since our search strategy looks for both a bosonic resonance and a displaced vertex in the same event, simulation of detection takes place in two parts. For the bosonic resonance, the W and Z bosons decay into other particles before they reach the detector. So, we look for the decay products of a W or Z to know that one was there. Each of these particles can decay into leptons or into quarks. In our simulation, we compare a random number to the branching ratios of these decay modes (which we obtain from the PDG [19]) to decide which type of decay will happen. From there, we use the same formulas given in Sec. 4.4 to generate four-momenta for the decay products. We then check the four-momenta of the decay products against the appropriate cuts as described in Sec. 3.2.

If the bosonic resonance passes the cuts, we move on to the displaced vertex. Instead of simulating the decay position of each glueball to see which detector subsystem it decays in, we use the glueball decay width and its trajectory through the detector to calculate a probability that it decays within each detector subsystem. The probability that the glueball decays between positions L_1 and L_2 in the detector is

$$P(L_1, L_2) = e^{-\left(\frac{L_1}{c\tau}\right)\frac{1}{\gamma\beta}} - e^{-\left(\frac{L_2}{c\tau}\right)\frac{1}{\gamma\beta}} \ , \tag{4.22}$$

where β is the glueball's speed in the lab frame (using the speed of light $c = 1$), γ is the relativistic factor $\gamma = (1 - \beta^2)^{-1/2}$, and $c\tau$ is the decay length of the glueball, which is related to the decay width Γ given in Eq. 2.16 by the equation

$$c\tau = \frac{c\hbar}{\Gamma} \ . \tag{4.23}$$

I determine the positions L_1 and L_2 using the glueball's trajectory through the detector and the detector characterizations given in Sec. 3.1. L_1 is the distance from the initial collision to the point where the glueball enters the subsystem, and L_2 is the distance from the initial collision to the point where the glueball exits the subsystem.

To convert the probability to a number of observed events, we multiply the probability by the subsystem's DV reconstruction efficiency as given in Sec. 3.3. The result is a fraction of an event, and the accumulation of these fractional events over the course of the simulation eventually leads us to a total number of observed events in each detector subsystem.

4.6 Prediction of Observed Events

We are interested in determining the fraction of simulated events that our search strategy could detect. To see if an event is detected, we first check if the squirk pair de-excites before it beta decays. If it does, we next check if the gluon hadronizes into a 0^{++} glueball. If it does, our next step is to randomly select a decay mode for the W boson based on its branching ratios. For the WZ simulations, we also select a decay mode for the Z. Based on the randomly selected decay mode, we check the decay products against the appropriate cuts from the resonance searches. If the decay products pass the cuts, we have passed the resonance requirements, so we turn our attention to the glueball. Using the glueball's trajectory and decay length, we calculate the probability that it would decay in the Tracker, HCal, and MS. For each subsystem, we multiply this probability by the appropriate displaced vertex reconstruction efficiency. We add these three numbers to a running total for each detector subsystem, meaning that we have added a fraction of an event to each subsystem's counter of detected events.

Once we have worked through this process for all the events, we can divide each subsystem's counter by the total number of events to determine the fraction of simulated events that our search strategy was able to detect in a given detector subsystem. To convert this fraction into a prediction of the number of events observed in an experiment, we multiply the luminosity $\mathcal{L}$ of the experiment (3000 fb^{-1} for the HL-LHC) times the cross section σ of the process (given by MadGraph as explained in Sec. 4.1) times the branching ratio $\text{BR}(\Psi_{\pm} \to WV)$ of either $V = \gamma$ or $V = Z$ (given in Eqs. 2.7 and 2.8). This gives us an expected number of events in the experiment, which we can then multiply by the fraction f from our simulation to determine how many of those experimental events our search strategy would detect. Thus for each combination of the parameters M, $\tan(\beta)$, m_0, and c_g, the predicted number N of resonance + DV events observed in one subsystem at the HL-LHC is

$$N_{WV} = \mathcal{L} \cdot \sigma(pp \to \Psi_{\pm} g) \cdot \mathrm{BR}(\Psi_{\pm} \to WV) \cdot f(M, \tan(\beta), m_0, c_g) \ . \tag{4.24}$$

Chapter 5

Results

For each combination of the parameters M, $\tan(\beta)$, m_0, and c_g, my simulation (described in Chap. 4) predicts how many experimental events at the HL-LHC our search strategy would detect, one prediction for each detector subsystem. Each subsystem gets its own prediction because differences in size and DV reconstruction efficiency give each subsystem a unique sensitivity to the resonance + DV signal. The predictions are organized into grids of plots in Figs. 5.1-5.3, which I copied from the erratum of our publication [5].

Each plot grid is organized by the glueball portion of the parameter space: the 0^{++} glueball mass m_0 determines a plot's horizontal position in the grid, and the glueball coupling parameter c_g determines a plot's vertical position in the grid. These values are labeled on the upper and right axes respectively. This grid structure allows for easy comparisons with the glueball decay length plot in Fig. 2.3, comparisons I discuss in Chap. 6.

Each plot in a grid displays the squirk portion of the parameter space: the x-axis is the total mass M of the squirk pair, and the y-axis is the supersymmetry parameter $\tan(\beta)$ used to determine the mass difference between the squirks in the pair. Notice that each plot takes $1 \leq \tan(\beta) \leq 10$, even though I previously considered $0.1 \leq \tan(\beta) \leq 10$. Due to the approximate log-scale symmetry across $\tan(\beta) = 1$ as seen in Fig. 2.2, I expect similar results for $\tan(\beta) > 1$ and $\tan(\beta) < 1$, so I save space by plotting results for only $1 \leq \tan(\beta) \leq 10$, knowing that the results for $0.1 \leq \tan(\beta) \leq 1$ are qualitatively similar.

To display the simulation data, each plot contains two sensitivity regions. The sensitivity regions for the $W\gamma$ resonance + DV signal are shown in solid orange, and the WZ + DV sensitivity regions are in solid blue. I define a sensitivity region to be a region of parameter space where the simulation predicts three or more detected events. My collaborators and I chose such a small number of events as the cutoff for sensitivity because there is essentially no SM background for the resonance + DV signal, and with no background, detecting even a few events is evidence of a discovery. We chose the number three specifically because the authors of Ref. [25], who had a similarly background-free signal, made the same choice. So, the sensitivity regions mark where our strategy would detect enough events to claim a discovery.

In each plot, I have included the exclusion regions from Fig. 2.2 for comparison with my simulation data. The black regions are exclusions from previous $W\gamma$ searches, and the dashed orange regions are the projected exclusions for $W\gamma$ searches at the HL-LHC. Comparing these regions to my simulated data shows that our search strategy is expected to perform quite well: it complements and improves current methods, as I discuss in more detail in Chap. 6.

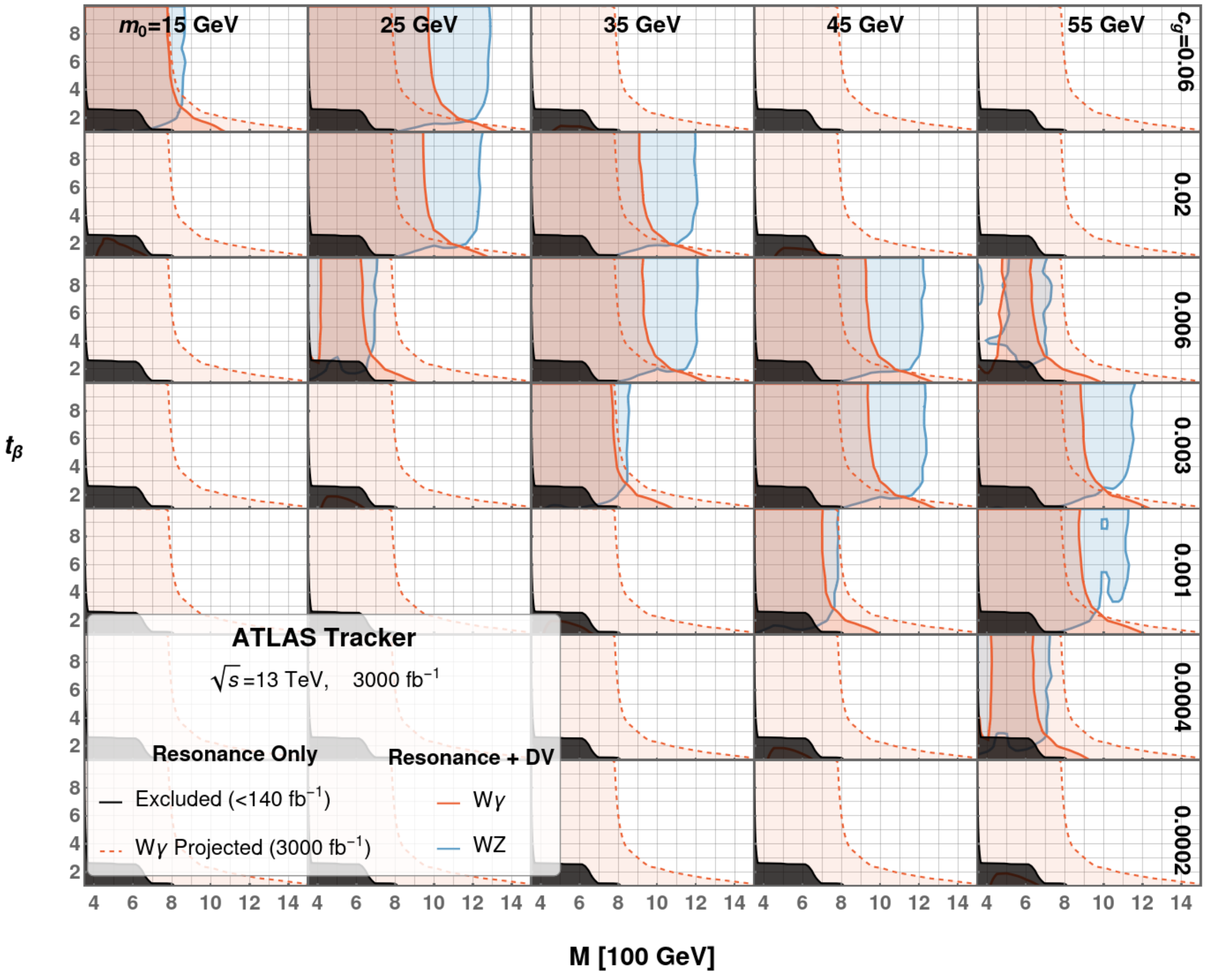

Figure 5.1 Sensitivity regions for our resonance + displaced vertex search strategy in the ATLAS Tracker. The grid is organized by the glueball parameters m_0 (mass) and c_g (coupling) as labeled on the upper and right axes respectively. For each combination of glueball parameters, the plot displays data in squirk mass parameter space with total mass M on the x-axis and supersymmetry parameter $\tan(\beta)$ on the y-axis. The solid orange (blue) regions are the sensitivity regions for detecting a $W\gamma$ (WZ) resonance + displaced vertex. Sensitivity is defined as predicted detection of three or more events. Exclusion regions from Fig. 2.2 are included in each plot for comparison. Black regions are exclusions from previous $W\gamma$ resonance searches, and dashed orange regions are projected exclusions for $W\gamma$ searches at the HL-LHC. For comparison with these projected exclusions, sensitivity regions were simulated at 13 TeV CoM energy and 3000 fb^{-1} integrated luminosity.

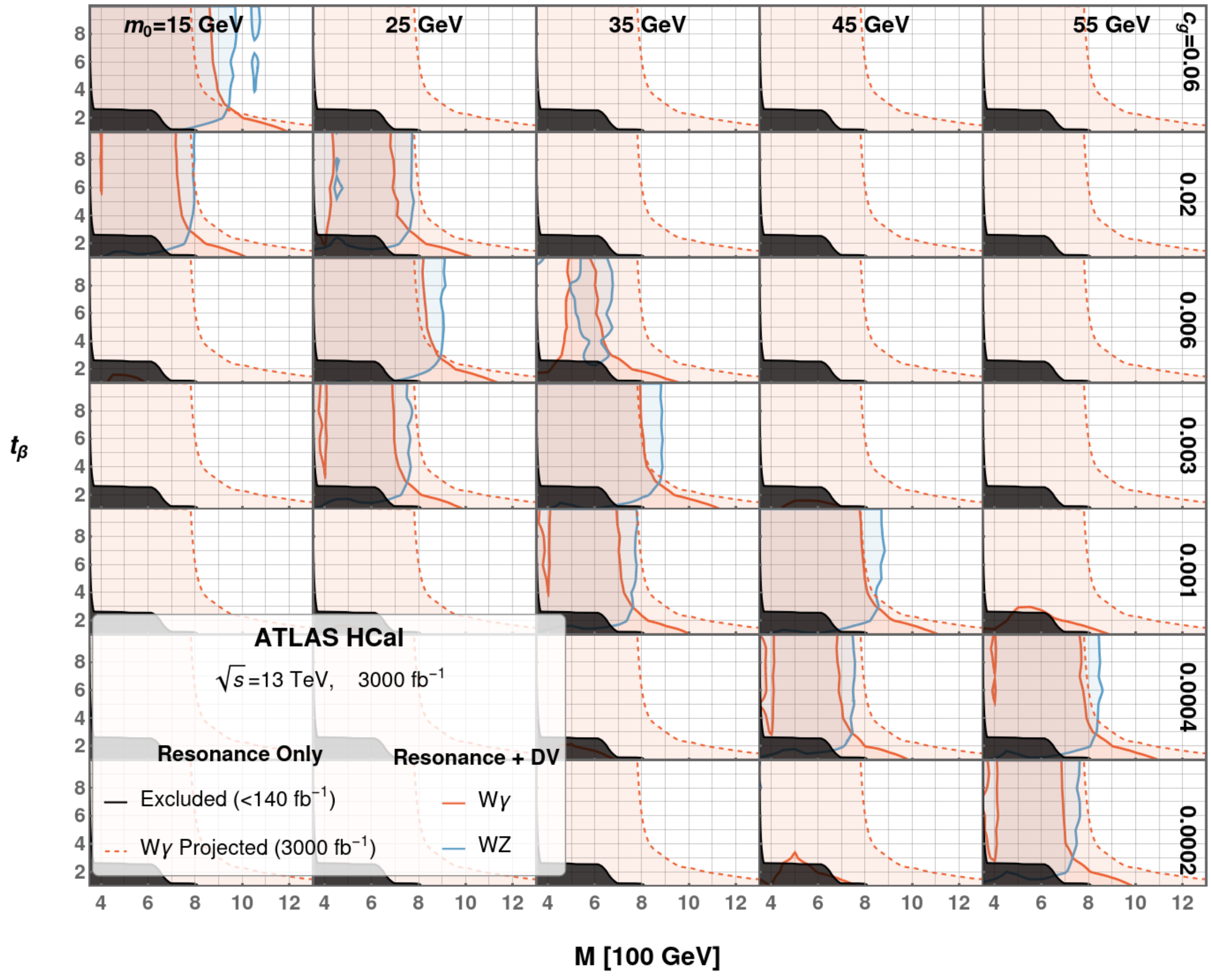

Figure 5.2 Sensitivity regions for our resonance + displaced vertex search strategy in the ATLAS Hadronic Calorimeter. The grid is organized by the glueball parameters m_0 (mass) and c_g (coupling) as labeled on the upper and right axes respectively. For each combination of glueball parameters, the plot displays data in squirk mass parameter space with total mass M on the x-axis and supersymmetry parameter $\tan(\beta)$ on the y-axis. The solid orange (blue) regions are the sensitivity regions for detecting a $W\gamma$ (WZ) resonance + displaced vertex. Sensitivity is defined as predicted detection of three or more events. Exclusion regions from Fig. 2.2 are included in each plot for comparison. Black regions are exclusions from previous $W\gamma$ resonance searches, and dashed orange regions are projected exclusions for $W\gamma$ searches at the HL-LHC. For comparison with these projected exclusions, sensitivity regions were simulated at 13 TeV CoM energy and 3000 fb^{-1} integrated luminosity.

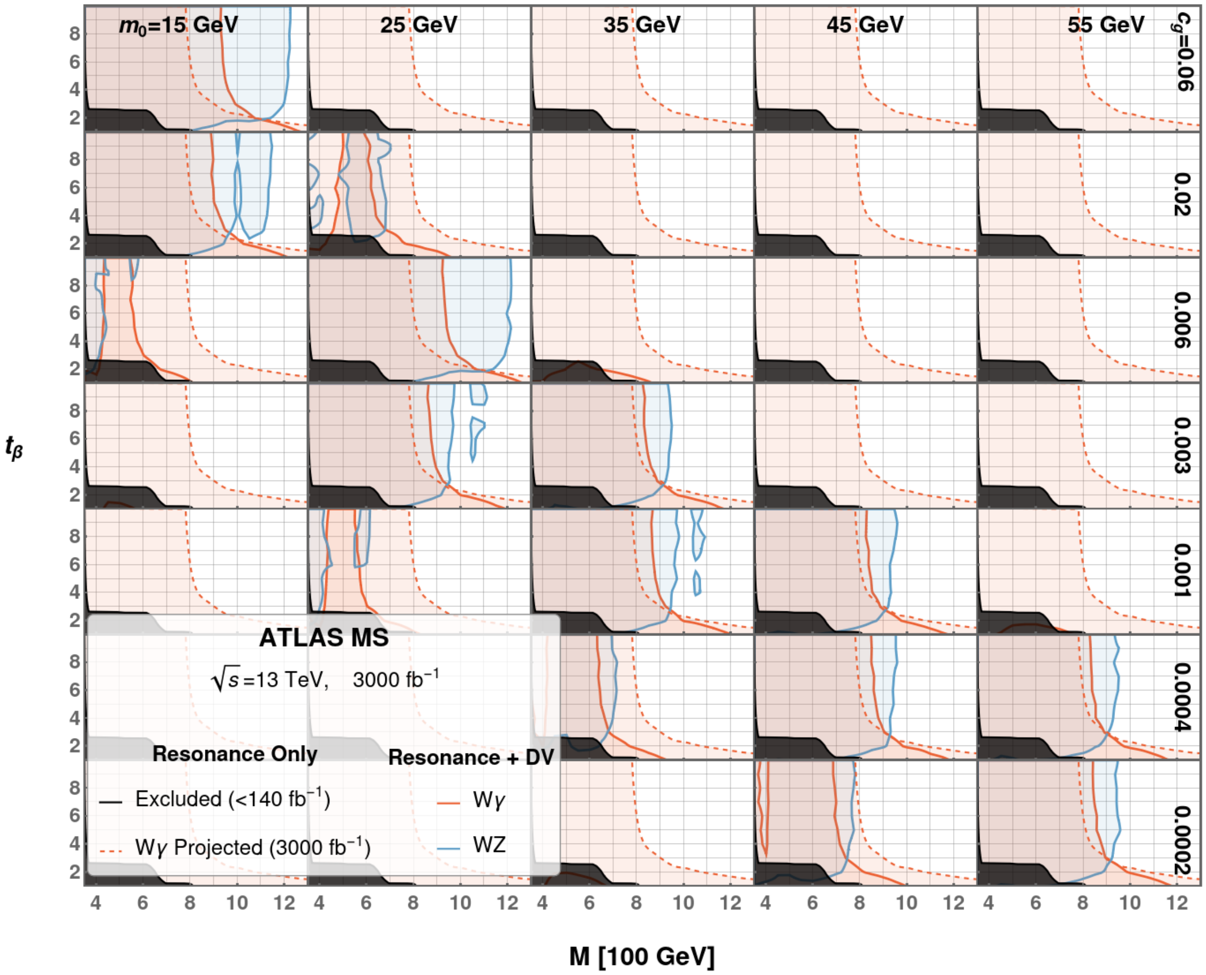

Figure 5.3 Sensitivity regions for our resonance + displaced vertex search strategy in the ATLAS Muon Spectrometer. The grid is organized by the glueball parameters m_0 (mass) and c_g (coupling) as labeled on the upper and right axes respectively. For each combination of glueball parameters, the plot displays data in squirk mass parameter space with total mass M on the x-axis and supersymmetry parameter $\tan(\beta)$ on the y-axis. The solid orange (blue) regions are the sensitivity regions for detecting a $W\gamma$ (WZ) resonance + displaced vertex. Sensitivity is defined as predicted detection of three or more events. Exclusion regions from Fig. 2.2 are included in each plot for comparison. Black regions are exclusions from previous $W\gamma$ resonance searches, and dashed orange regions are projected exclusions for $W\gamma$ searches at the HL-LHC. For comparison with these projected exclusions, sensitivity regions were simulated at 13 TeV CoM energy and 3000 fb^{-1} integrated luminosity.

Chapter 6

Discussion

The search strategy I outlined in this thesis looks for the annihilation of a squirk bound state and the long-lived decay of a hidden glueball. The squirk bound state has two crucial parameters: the total mass M of the two squirks and the parameter $\tan(\beta)$ used to find the mass difference Δ between the squirks. The hidden glueball also has two crucial parameters: the mass m_0 of the 0^{++} glueball and the coupling parameter c_g. To predict the performance of our search strategy at the HL-LHC, I ran simulations for almost 20,000 combinations of these parameters. The results of the simulations are presented in Figs. 5.1–5.3, which display grids of plots, one grid each for the ATLAS Tracker, HCal, and MS. Each plot has M on the x-axis and $\tan(\beta)$ on the y-axis and is placed in the grid according to its m_0 and c_g values as labeled on the upper and right axes of the grid. The dark region in each plot shows which $\tan(\beta)$-values have been excluded by experimental searches for a given M-value. The dashed orange region in each plot shows the projected exclusions for the HL-LHC, and the solid orange and blue regions show where our new search strategy would be sensitive to squirk signals based on the results of the simulation.

The plot grids are organized by m_0 and c_g for easier comparison with Fig. 2.3, which plots the decay length of the glueball as a function of m_0 and c_g. This decay length determines the probability that the glueball decays within the detector (see Eq. 4.22). When this probability is low, our strategy loses efficacy. You can see this effect in the upper-right and lower-left corners of each plot grid. There are no sensitivity regions in these portions of parameter space because the glueball decays too soon (upper-right) or too late (lower-left) to be detected.

On the other hand, our strategy is most effective when there's a high probability that the glueball decays within the detector. The red highlighted portion of Fig. 2.3 shows the approximate range of distances from the collision point for which the ATLAS detector is sensitive to displaced vertices. Notice that this region lines up very nicely with the band of plots in each grid containing sizable sensitivity regions. For these combinations of glueball parameters, the likelihood of the glueball decaying in the detector is much higher, creating larger sensitivity regions.

In comparing the three plot grids, you'll notice that the band of plots with sizable sensitivity regions is not in the same place in each grid. For example, the Tracker's sensitivity regions are closer to the upper-right corner of the grid than the MS's sensitivity regions. The band of sensitivity seems to shift when we look at a different ATLAS subsystem. We can explain this effect by comparing the layout of the ATLAS detector (see Fig. 3.1) to the red highlighted portion of Fig. 2.3. Notice that moving from upper-right to lower-left in Fig. 2.3 increases the decay length. As the decay length increases, the glueball will be more likely to decay in detector subsystems farther from the collision point. So, since the Tracker is closest to the collision and the MS is farthest, we expect the Tracker to detect displaced vertices for decay lengths in the upper-right portion of the red highlighted region, and we expect the MS to detect displaced vertices for decay lengths in the lower-left portion of that region. Thus the relative positions of the detector subsystems explain the shifting sensitivity band we observed in the plot grids.

In addition to the different positions of the detector subsystems, each subsystem's DV reconstruction efficiency (see Sec. 3.3) also plays a role in the plot grids. You'll notice that the sensitivity regions in the HCAL are smaller than the regions in the Tracker and the MS. The smaller HCAL sensitivity regions come from the behavior of the HCAL: it only records data about energy deposition instead of completely reconstructing the displaced vertex. In order to find a displaced vertex in the HCAL, the vertex must satisfy more stringent cuts than for the Tracker or MS, which do reconstruct displaced vertices. These more stringent cuts lead to a smaller effective efficiency for detecting displaced vertices, and thus smaller discovery regions in the HCAL.

In each plot, I have included the projected exclusion region for $W\gamma$ searches at the HL-LHC (see Fig. 2.2). Since the sensitivity regions come from simulations of the HL-LHC, it is useful to compare these regions to the projected exlusion to see how our search strategy compares to current methods. We see that outside the sensitivity band of each grid, our search strategy performs poorly while the current methods are unaffected. Because our strategy is so sensitive to the glueball parameters, it cannot compete with current search methods across all of glueball parameter space. However, within the sensitivity band, the sensitivity regions cover more squirk parameter space than the projected exclusion, especially for $\tan(\beta) \neq 1$. We see then that our strategy cannot replace current methods, but it can complement those methods by exploring more squirk parameter space for certain combinations of glueball parameters.

Additionally, our strategy complements current methods when the sensitivity regions overlap the projected exclusion region. Since squirk masses in the projected exclusion region have not yet been officially ruled out, it is still possible that the HL-LHC could discover squirks with masses in this region. Because the sensitivity regions can overlap with the projected exclusion region, our strategy provides an interesting follow-up on any discoveries in this region. If the standard $W\gamma$ searches discover something in this region, detectors could use our strategy to confirm that the signal comes from squirks and not some other BSM physics. They could also use our strategy to constrain the glueball parameters in the squirk system they've discovered.

In many plots in the sensitivity band of each grid, a search for the WZ + DV signal performs better than a search for the $W\gamma$ + DV signal when $\tan(\beta) \neq 1$. The superiority of the WZ signal results from the branching ratios of squirk pair annihilation (see Fig. 2.1). The WZ branching ratio is gets larger as $\tan(\beta)$ approaches 0 or infinity, so the WZ + DV signal is more common for $\tan(\beta) \neq 1$. As the signal gets more common, our strategy detects it more frequently, leading to large WZ sensitivity regions for $\tan(\beta) \neq 1$.

Chapter 7

Conclusion

In this thesis, I described a new strategy to search for squirks by looking for a bosonic resonance and a displaced vertex in the same event. To predict how this strategy would perform at the HL-LHC, I ran simulations of producing the signal and detecting it in the ATLAS detector using our strategy. The results of the simulation show that our new strategy complements current methods and explores more squirk parameter space, increasing the potential for discovering squirks at the LHC.

In studying the bound state of a squirk and an anti-squirk, we noticed that when the state has net electric charge, the only available decay channels are $W\gamma$ and WZ. As the difference between the masses of the two squirks increases, the $W\gamma$ branching ratio shrinks, eventually dropping below the WZ branching ratio (see Fig. 2.1). We found that this shrinking $W\gamma$ branching ratio weakens the bounds on squirk masses set by standard searches at the LHC [12–17, 26]. As shown in Fig. 2.2, these weaker bounds leave unexplored portions of squirk mass parameter space that are accessible at energies already achievable at the LHC. We therefore developed our search strategy to explore these regions of parameter space without any upgrades to collider technology beyond those already planned for the HL-LHC.

In running the simulations to predict the strategy's performance at the HL-LHC, I made several choices to keep the simulation as realistic as possible. I made sure to account for the possibility that the squirk pair might beta decay before it annihilates (see Sec. 4.2. A beta-decayed state cannot contribute to our signal because it cannot produce a $W\gamma$ or a WZ resonance when it annihilates. It was therefore crucial that I consider beta decay to make sure I wasn't overestimating the number of events our strategy would detect. Additionally, I used a simplified mock-up of the ATLAS detector to apply the DV reconstruction efficiencies in the correct regions of the detector subsystems. I also included cuts from previous $W\gamma$ and WZ searches to make sure my simulation of detection was comparable to the actual behavior of the ATLAS detector (see Sec. 4.5).

To adequately cover the parameter space, I ran the simulation for each combination of the following parameters: $\tan(\beta)$, the total mass M of the squirks, the lightest glueball mass m_0, and the hidden glueball coupling parameter c_g. I took M from 200 to 1600 GeV in steps of 50, $\tan(\beta)$ from 0.1 to 1 in steps of 0.1 and 2 to 10 in steps of 1, m_0 from 10 to 50 GeV in steps of 10, and $c_g \in [0.0002, 0.0004, 0.001, 0.003, 0.006, 0.02, 0.06]$. I find that these choices of parameter values give a broad view of the expected performance of the search strategy, including its weaknesses and strengths as discussed in Chap. 6.

The results of my simulation demonstrate that for favorable glueball decay lengths, our search strategy explores a larger region of squirk mass parameter space than previous searches at the LHC. Because the strategy depends on the glueball parameters, it cannot completely replace current search methods. However, it can complement current methods by confirming that an observed signal comes from squirks and not some other BSM physics.

Because this work focused on the electrically-charged squirk bound state, future work should consider the neutral bound state. In contrast to the charged case, the neutral bound state's branching ratio to two hidden gluons is overwhelmingly larger than any of its other branching ratios. Thus any future work with neutral bound states ought to consider the possibility of multiple hidden glueballs in the same event, leading to several displaced vertices.

Appendix A

Momenta of Two-Body Decay Products

Two-body decay is when an unstable particle ("mother" particle) decays into two particles ("daughter" particles). The daughters' four-momenta in the mother's rest frame are well-known results that we use in our simulations of two-body decay. To derive these four-momenta, we consider a particle of mass M and four-momentum P^μ that decays into two particles; the first has mass m_1 and four-momentum p_1^μ, and the second has mass m_2 and four-momentum p_2^μ. Conservation of four-momentum in the mother's rest frame gives us the following equation:

$$\begin{aligned} P^\mu &= p_1^\mu + p_2^\mu \\ \begin{pmatrix} M \\ \vec{0} \end{pmatrix} &= \begin{pmatrix} \sqrt{|\vec{p}_1|^2 + m_1^2} \\ \vec{p}_1 \end{pmatrix} + \begin{pmatrix} \sqrt{|\vec{p}_2|^2 + m_2^2} \\ \vec{p}_2 \end{pmatrix} \end{aligned} \tag{A.1}$$

where $\vec{p}_1$ and $\vec{p}_2$ are the three-momenta of the first and second daughter particles, respectively. The spatial components of this equation reveal that

$$\vec{p_1} = -\vec{p_2}, \tag{A.2}$$

so we define a new variable:

$$p = |\vec{p_1}| = |\vec{p_2}|. \tag{A.3}$$

We now turn our attention to the temporal components of Eq. A.1:

$$M = \sqrt{p^2 + m_1^2} + \sqrt{p^2 + m_2^2} \tag{A.4}$$

We want to solve for p so we can express the daughters' four-momenta entirely in terms of the three masses. To begin, we define two new variables,

$$m_{\text{tot}} = m_1 + m_2 \quad \text{and} \quad \Delta m = m_1 - m_2, \tag{A.5}$$

and rewrite the right-hand side of Eq. A.4:

$$M = \sqrt{p^2 + \left(\frac{m_{\text{tot}} + \Delta m}{2}\right)^2} + \sqrt{p^2 + \left(\frac{m_{\text{tot}} - \Delta m}{2}\right)^2}. \tag{A.6}$$

I should note that rewriting this expression in terms of m_{tot} and Δm is not strictly necessary to solve the equation. However, my preferred expression of the solution uses m_{tot} and Δm, so I will use them in this derivation.

We begin solving for p by squaring both sides:

$$\begin{aligned} M^2 = p^2 + \left(\frac{m_{\text{tot}} + \Delta m}{2}\right)^2 + p^2 + \left(\frac{m_{\text{tot}} - \Delta m}{2}\right)^2 \\ + 2\sqrt{\left[p^2 + \left(\frac{m_{\text{tot}} + \Delta m}{2}\right)^2\right]\left[p^2 + \left(\frac{m_{\text{tot}} - \Delta m}{2}\right)^2\right]} \end{aligned} \tag{A.7}$$

We then move everything but the square root to the left-hand side and simplify some things:

$$M^2 - 2p^2 - \frac{1}{2}m_{\text{tot}}^2 - \frac{1}{2}\Delta m^2 = 2\sqrt{\left[p^2 + \left(\frac{m_{\text{tot}} + \Delta m}{2}\right)^2\right]\left[p^2 + \left(\frac{m_{\text{tot}} - \Delta m}{2}\right)^2\right]}. \tag{A.8}$$

Now we square both sides again:

$$\begin{aligned} M^4 + 4p^4 + \frac{1}{4}m_{\text{tot}}^4 + \frac{1}{4}\Delta m^4 - 4M^2p^2 - M^2m_{\text{tot}}^2 - M^2\Delta m^2 + 2p^2m_{\text{tot}}^2 + 2p^2\Delta m^2 + \frac{1}{2}m_{\text{tot}}^2\Delta m^2 = \\ 4\left[p^4 + p^2\left(\frac{m_{\text{tot}} + \Delta m}{2}\right)^2 + p^2\left(\frac{m_{\text{tot}} - \Delta m}{2}\right)^2 + \left(\frac{m_{\text{tot}} + \Delta m}{2}\right)^2\left(\frac{m_{\text{tot}} - \Delta m}{2}\right)^2\right] \end{aligned} \tag{A.9}$$

and simplify the right-hand side:

$$M^4+4p^4+\frac{1}{4}m_{\text{tot}}^4+\frac{1}{4}\Delta m^4-4M^2p^2-M^2m_{\text{tot}}^2-M^2\Delta m^2+2p^2m_{\text{tot}}^2+2p^2\Delta m^2+\frac{1}{2}m_{\text{tot}}^2\Delta m^2 = 4p^4+2p^2m_{\text{tot}}^2+2p^2\Delta m^2+\frac{1}{4}(m_{\text{tot}}^2+2m_{\text{tot}}\Delta m+\Delta m^2)(m_{\text{tot}}^2-2m_{\text{tot}}\Delta m+\Delta m^2)$$

$$M^4+4p^4+\frac{1}{4}m_{\text{tot}}^4+\frac{1}{4}\Delta m^4-4M^2p^2-M^2m_{\text{tot}}^2-M^2\Delta m^2+2p^2m_{\text{tot}}^2+2p^2\Delta m^2+\frac{1}{2}m_{\text{tot}}^2\Delta m^2 = 4p^4+2p^2m_{\text{tot}}^2+2p^2\Delta m^2+\frac{1}{4}m_{\text{tot}}^4+\frac{1}{4}\Delta m^4-\frac{1}{2}m_{\text{tot}}^2\Delta m^2. \tag{A.10}$$

Many terms cancel, and after isolating the p^2 term on the left-hand side, we have:

$$-4M^2p^2 = -M^4+M^2m_{\text{tot}}^2+M^2\Delta m^2-m_{\text{tot}}^2\Delta m^2 \tag{A.11}$$

Multiplying by -1 and factoring the right-hand side gives us

$$4M^2p^2 = (M^2-m_{\text{tot}}^2)(M^2-\Delta m^2). \tag{A.12}$$

Dividing both sides by M^4 and isolating p^2 gives us

$$p^2 = \frac{M^2}{4}\left(1-\left(\frac{m_{\text{tot}}}{M}\right)^2\right)\left(1-\left(\frac{\Delta m}{M}\right)^2\right), \tag{A.13}$$

so the magnitude of both daughters' three-momentum in the mother's rest frame is

$$p = \frac{M}{2}\sqrt{\left(1-\left(\frac{m_{\text{tot}}}{M}\right)^2\right)\left(1-\left(\frac{\Delta m}{M}\right)^2\right)}. \tag{A.14}$$

If the first daughter particle travels in the direction $\hat{d}$ in the mother's rest frame, then the daughters' four-momenta are

$$p_1^\mu = \begin{pmatrix}\sqrt{p^2+m_1^2}\\ p\hat{d}\end{pmatrix} \quad \text{and} \quad p_2^\mu = \begin{pmatrix}\sqrt{p^2+m_2^2}\\ -p\hat{d}\end{pmatrix} \tag{A.15}$$

where p is given by Eq. A.14. We us these expressions in our simulations, and the direction $\hat{d}$ is chosen randomly for each simulated decay.

Appendix B

Lorentz Boost in an Arbitrary Direction

Textbooks that introduce the idea of a Lorentz boost often simplify the concept (for good reason) by defining the direction of the boost to be the x-direction. The transformation matrix in the $\{x, y, z\}$ basis is thus

$$\Lambda^{\mu}{}_{\nu} = \begin{pmatrix} \gamma & \gamma\beta & 0 \\ \gamma\beta & \gamma & 0 \\ 0 & 0 & \mathbb{I} \end{pmatrix}, \tag{B.1}$$

where β is the relative speed of the reference frames (we use the convention that the speed of light $c = 1$), and $\gamma = (1 - \beta^2)^{-1/2}$ as usual. The $\mathbb{I}$ is an identity matrix for the y- and z-components of the object that Λ acts on. This identity matrix represents a crucial idea, namely that a Lorentz boost does not affect the components of a four-vector perpendicular to the direction of the boost.

Using that idea, we can generalize the boost to any direction using a simple change of basis. Let's say that the unprimed frame is moving in the direction $\hat{d}$ at speed β relative to the primed frame. Note that $\hat{d}$ is a unit vector. For some four-vector X^{μ} in the unprimed frame, we would write

$$X^{\mu} = (X^0, \vec{X}). \tag{B.2}$$

In the regular $\{x, y, z\}$ basis, this looks like

$$X^\mu = (X^0, \vec{X} \cdot \hat{x}, \vec{X} \cdot \hat{y}, \vec{X} \cdot \hat{z}) = (X^0, X^1, X^2, X^3) \tag{B.3}$$

where $\hat{x}$ is the unit vector in the x-direction, and similarly for $\hat{y}$ and $\hat{z}$. Since the boost will only affect the component of $\vec{X}$ in the $\hat{d}$-direction, it would be more helpful to use $\hat{d}$ as one of the basis vectors. In the $\hat{d}$ basis, we have

$$X^\mu = (X^0, \vec{X} \cdot \hat{d}, \vec{X} - (\vec{X} \cdot \hat{d})\hat{d}) = (X^0, X_\parallel, \vec{X}_\perp) \tag{B.4}$$

where the first spatial component $X_\parallel$ is the component of $\vec{X}$ in the direction of $\hat{d}$, and $\vec{X}_\perp$ represents the two remaining components of $\vec{X}$ perpendicular to $\hat{d}$. The convenience of choosing this basis is that $\vec{X}_\perp$ will be completely unaffected by a boost in the $\hat{d}$-direction.

With this new choice of basis, we can now write the transformation to the primed frame as

$$\begin{pmatrix} X'^0 \\ X'_\parallel \\ \vec{X}'_\perp \end{pmatrix} = \begin{pmatrix} \gamma & \gamma\beta & 0 \\ \gamma\beta & \gamma & 0 \\ 0 & 0 & \mathbb{I} \end{pmatrix} \begin{pmatrix} X^0 \\ X_\parallel \\ \vec{X}_\perp \end{pmatrix}. \tag{B.5}$$

Expanding this gives us

$$\begin{pmatrix} X'^0 \\ X'_\parallel \\ \vec{X}'_\perp \end{pmatrix} = \begin{pmatrix} \gamma X^0 + \gamma\beta X_\parallel \\ \gamma\beta X^0 + \gamma X_\parallel \\ \vec{X}_\perp \end{pmatrix}. \tag{B.6}$$

The spatial components of these vectors can be rewritten in a suggestive way:

$$\begin{pmatrix} X'^0 \\ \vec{X}' \end{pmatrix} = \begin{pmatrix} \gamma X^0 + \gamma\beta X_\parallel \\ \vec{X} + \begin{pmatrix} \gamma\beta X^0 + (\gamma - 1) X_\parallel \\ \vec{0} \end{pmatrix} \end{pmatrix}. \tag{B.7}$$

On the left-hand side, we have combined $X'_\parallel$ and $\vec{X}'_\perp$ into $\vec{X}'$. On the right-hand side, we pulled one factor of $X_\parallel$ out of the first spatial component and combined it with the $\vec{X}_\perp$ to make $\vec{X}$. This leaves $(\gamma - 1)$ factors of $X_\parallel$ in the first spatial component. To simplify this expression further, note that in the basis we've chosen,

$$\hat{d} = \begin{pmatrix} 1 \\ \vec{0} \end{pmatrix}, \tag{B.8}$$

which gives us

$$\begin{pmatrix} X'^0 \\ \vec{X}' \end{pmatrix} = \begin{pmatrix} \gamma X^0 + \gamma\beta X_\parallel \\ \vec{X} + [\gamma\beta X^0 + (\gamma - 1)X_\parallel]\hat{d} \end{pmatrix}, \tag{B.9}$$

or more explicitly,

$$\begin{pmatrix} X'^0 \\ \vec{X}' \end{pmatrix} = \begin{pmatrix} \gamma X^0 + \gamma\beta(\vec{X} \cdot \hat{d}) \\ \vec{X} + [\gamma\beta X^0 + (\gamma - 1)(\vec{X} \cdot \hat{d})]\hat{d} \end{pmatrix}. \tag{B.10}$$

We use this equation in our simulations when boosting between different reference frames, especially between the lab frame and particle frames.

Bibliography

[1] S. Chatrchyan *et al.*, “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC,” Phys. Lett. B **716,** 30–61 (2012).

[2] G. Aad *et al.*, “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC,” Phys. Lett. B **716,** 1–29 (2012).

[3] M. Aaboud *et al.*, “Performance of the ATLAS trigger system in 2015,” European Physical Journal. C, Particles and Fields 77 (2017).

[4] O. Aberle *et al.*, “High-Luminosity Large Hadron Collider (HL-LHC): Technical Design Report,” Technical Report No. CERN-2020-010, CERN (2020) .

[5] J. Forsyth, M. Low, C. Tenney, and C. B. Verhaaren, “Visible collider signals of natural quirks,” Phys. Rev. D **111,** 115021 (2025).

[6] B. Batell, M. Low, E. T. Neil, and C. B. Verhaaren, “Review of neutral naturalness,” Physics Reports **1165,** 1–49 (2026).

[7] N. Craig, “Naturalness: A Snowmass White Paper,”, 2022.

[8] S. P. Martin, “A Supersymmetry primer,” Adv. Ser. Direct. High Energy Phys. **18,** 1–98 (1998).

[9] G. Burdman, Z. Chacko, H.-S. Goh, and R. Harnik, “Folded supersymmetry and the LEP paradox,” JHEP **02,** 009 (2007).

[10] D. Curtin and C. B. Verhaaren, “Discovering Uncolored Naturalness in Exotic Higgs Decays,” JHEP **12,** 072 (2015).

[11] J. Kang and M. A. Luty, “Macroscopic strings and “quirks” at colliders,” Journal of High Energy Physics **2009,** 065–065 (2009).

[12] G. Aad *et al.*, “Search for new resonances in Wγ and Zγ final states in pp collisions at sqrt(s) = 8 TeV with the ATLAS detector,” Physics Letters B **738,** 428–447 (2014).

[13] G. Aad *et al.*, “Search for heavy diboson resonances in semileptonic final states in pp collisions at sqrt(s) = 13 TeV with the ATLAS detector,” The European Physical Journal C 80 (2020).

[14] G. Aad *et al.*, “Search for resonant WZ production in the fully leptonic final state in proton–proton collisions at sqrt(s) = 13 TeV with the ATLAS detector,” The European Physical Journal C 83 (2023).

[15] A. Tumasyan *et al.*, “Search for Wγ resonances in proton-proton collisions at $\sqrt{s} = 13$ TeV using hadronic decays of Lorentz-boosted W bosons,” Phys. Lett. B **826,** 136888 (2022).

[16] A. Hayrapetyan *et al.*, “Search for a resonance decaying to a W boson and a photon in proton-proton collisions at sqrt(s) = 13 TeV using leptonic W boson decays,” Journal of High Energy Physics 2024 (2024).

[17] A. Tumasyan *et al.*, “Search for heavy resonances decaying to WW, WZ, or WH boson pairs in a final state consisting of a lepton and a large-radius jet in proton-proton collisions at sqrt(s) = 13 TeV,” Physical Review D 105 (2022).

[18] S. P. Martin, "Diphoton decays of stoponium at the CERN Large Hadron Collider," Physical Review D 77 (2008).

[19] S. Navas *et al.*, "Review of particle physics," Phys. Rev. D **110,** 030001 (2024).

[20] 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).

[21] J. E. Juknevich, "Pure-glue hidden valleys through the Higgs portal," JHEP **08,** 121 (2010).

[22] A. Djouadi, J. Kalinowski, and M. Spira, "HDECAY: A Program for Higgs boson decays in the standard model and its supersymmetric extension," Comput. Phys. Commun. **108,** 56–74 (1998).

[23] H. B. Meyer, "Glueball matrix elements: A Lattice calculation and applications," JHEP **01,** 071 (2009).

[24] A. Batz, T. Cohen, D. Curtin, C. Gemmell, and G. D. Kribs, "Dark sector glueballs at the LHC," Journal of High Energy Physics 2024 (2024).

[25] G. Ripellino, M. Vande Voorde, A. Gallén, and R. Gonzalez Suarez, "Searching for long-lived dark scalars at the FCC-ee," JHEP **06,** 143 (2025).

[26] G. Aad *et al.*, "Search for resonances decaying into photon pairs in 139 fb-1 of pp collisions at sqrt(s) = 13 TeV with the ATLAS detector," Physics Letters B **822,** 136651 (2021).

Index