Search for New Light Higgs Bosons in Boosted Tau Final States with the CMS Experiment

By

Francesca Shun-Ning Annarosa Ricci-Tam
B.S. (University of California, Davis) 2010
M.S. (University of California, Davis) 2012

Dissertation

Submitted in partial satisfaction of the requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Office of Graduate Studies

of the

University of California

Davis

Approved:

Chair Maxwell Chertok

Robin Erbacher

Michael Mulhearn

Committee in Charge

2016

Contents

Acknowledgments

I would like to thank my graduate advisor, Maxwell Chertok, for his patient help, advice, and friendship during my development as a high-energy physics researcher; and also Robin Erbacher and Michael Mulhearn for serving on my thesis committee. My appreciation extends to the rest of the UC Davis CMS group as well, for their welcoming and supportive working atmosphere.

During my years at CERN, I have learned a lot from working with my colleagues. There was Rachel Yohay, whose mentorship taught me much of what I know about data analysis techniques and CMS software as we worked together on what was to be my thesis research. Mauro Dinardo, Martina Malberti, and Gino Bolla all patiently gave me their expert advice and answered my questions during my hardware work with the CMS forward pixel detector. I would like to thank Dan Duggan and Gaelle Boudoul for teaching me about pixel offline software, Petra Merkel and Yana Osborne for getting me started in tracker geometry simulation work, and Viktor Veszprémi for his guidance during my tracker material budget group convenership.

I am grateful to the friends from around the world that I have met along the path to my PhD, who are too many to name here. Finally, I dedicate this thesis to my family. – To Papà, who has always taught me to have confidence in my own intellect and logic, and who has shown me by example that the mind of a physicist can tackle any challenge. To Mammina, who has encouraged my interest in reading and the study of nature. To my sisters Daniela and Chiara, who have always inspired my admiration with their boldness, imagination, and curiosity.

Preface

Elementary particle physics is the study of the most fundamental components of matter: determining what these components are, and how they behave. The word "atom" is derived from the ancient Greek word for "indivisible"; the discovery that atoms are not in fact indivisible but are themselves made up of smaller components – electrons, protons, and neutrons – set the stage for the ongoing search for ever-smaller and more fundamental building blocks of matter.

Decades of research have culminated in the present day with the Standard Model of fundamental physics. The search for the Higgs boson, the last major puzzle piece of the Standard model, carried on for many years, The most recent validation of the Standard Model was the discovery of the Higgs boson in 2012, which resolved the longstanding paradox of electroweak gauge boson masses. I feel fortunate to have started my work at CERN just at that exciting time when the Higgs boson discovery was announced, and when much work still lay ahead to uncover what lies beyond the Standard Model.

Various theories of physics beyond the Standard Model (BSM) have been developed, and they involve many unknown and interdependent parameters. Just as the Higgs boson mass was once unknown and had to be narrowed down by years of experimental research and data analysis, the parameters of BSM physics theories are still in the process of being increasingly constrained by experimental data, but there is a wide possible range of values open to them. This means that many different models of BSM physics, based on different combinations of assumptions for these parameters, still remain to be experimentally tested.

The collaboration of thousands of researchers at the European Organization for Nuclear Research (CERN) are currently engaged in analyzing data from the Large Hadron Collider for signals of processes predicted by the many models of BSM physics that there are. The research done for this dissertation has been part of that effort – my little contribution to our collective chipping-away at the unknown, inching towards a clearer picture of fundamental physics.

Francesca Shun-Ning Annarosa Ricci-Tam
March 2016
Physics

Search for New Light Higgs Bosons in Boosted
Tau Final States with the CMS Experiment

Abstract

In this dissertation, I present a search for non-standard decays of a Standard Model-like Higgs boson to pairs of light bosons, as predicted in models with extended Higgs sectors. In two Higgs doublet models, including the next-to-minimal supersymmetric standard model, the Higgs boson can decay into a pair of light pseudoscalars a.

In this search, the gluon fusion, W and Z associated Higgs, and vector boson fusion production channels for the Higgs are all considered, and the decay $H \to aa$ with $a \to \tau\tau$ is reconstructed from the tau decay products. The final state is characterized by one isolated high p_T muon plus at least one highly boosted pair of taus, of which one of the taus is required to decay to a muon.

Using 19.7 fb^{-1} of 8 TeV center of mass pp collision data recorded by the Compact Muon Solenoid experiment at the Large Hadron Collider, a counting experiment is performed in a region of high di-tau invariant mass. We have found no excess of events above the Standard Model backgrounds, and the observed data is used to set upper limits on the branching ratio $\mathrm{Br}(H \to aa)\mathrm{Br}^2(a \to \tau\tau)$. These results are equally applicable to decays of the SM-like Higgs boson to a pair of light scalars h.

Chapter 1

Theoretical introduction

1.1 The Standard Model

This chapter presents an overview of some of the main theoretical developments leading up to the Standard Model of fundamental physics, followed by an overview of the theory of supersymmetry and how it addresses some of the Standard Model's deficiencies.

1.1.1 A little background

Classical physics differentiates clearly between matter particles, which behave as localized objects, and radiation, which behaves as a wave. However, when physicists explored phenomena at the subatomic scale, the classical picture was found to be incorrect. The discovery of phenomena such as the photoelectric effect and the Compton effect, which point to the quantization of radiation, challenged classical assumptions about the wave behavior of light. Similarly, observations of diffraction behavior in electron beams revealed that beams of electrons can in fact behave not only like pointlike bodies but also like light waves [1]. The quest to understand this wave-particle duality led to the development of quantum mechanics to explain physical phenomena at the microscopic scale.

In quantum mechanics, the physical state of a particle or system of particles is characterized by a wavefunction. Measurable properties of particles, such as position and momentum, can all be derived from the wavefunction. The wavelike behavior of particles at the subatomic scale is reflected in the plane-wave wavefunction for free particles, and the bound-state wavefunctions that are related to standing waves, with discrete (quan-

tized) energy levels available to the particles in the bound system.

Quantum field theory developed from the need to describe the dynamics of relativistic elementary particles, for which nonrelativistic quantum mechanics is insufficient. For instance, consider the Klein-Gordon equation

$$(\partial^\mu \partial_\mu + m^2)\phi = 0 \tag{1.1}$$

which is the relativistic wave equation of motion for relativistic spin-0 particles. The plane-wave solutions to the Klein-Gordon equation have positive and negative energies, which correspond to particles with positive and negative probability densities. The latter concept is clearly nonphysical, but in the formalism of QFT, the negative-energy solutions acquire a physical interpretation[2, 3]. For each particle, there exists an antiparticle with identical mass and spin but opposite charge. The solution ϕ to the Klein-Gordon equation is not a single-particle wavefunction, but rather a scalar field, whose excitations correspond to the creation and annihilation of particles and antiparticles. The antiparticle corresponds to the negative-energy portion of the solution to the Klein-Gordon equation.

The same logic, applied to Dirac equation of motion for spin-$\frac{1}{2}$ particles,

$$(i\gamma^\mu \partial_\mu - m)\psi = 0 \tag{1.2}$$

led to the prediction of the existence of the positron, the antiparticle of the electron. Experimental support for this theory first came with the discovery of the positron in cloud-chamber studies of cosmic rays [4], confirming the existence of antiparticles and validating the description of particle physics with quantum field theory.

Particles interact via four fundamental forces: the strong force, electromagnetism, the weak force, and gravity. So far, the first three types of interactions have been successfully described by quantum field theories. The dynamics and interactions of fields are derived from the Lagrangian density $\mathcal{L}$, which is the quantum field theory analogue of the classical Lagrangian L. While L, the difference between the total kinetic and potential energy of a system of particles, is a function of the generalized coordinates and momenta of those particles, $\mathcal{L}$ is a function of fields and their spacetime derivatives.

By Noether's theorem, the invariance of $\mathcal{L}$ under continuous transformations of the fields implies the conservation of a current. Such transformations that leave $\mathcal{L}$ invariant can be expressed in terms of the generators of a symmetry group. Each fundamental interaction is governed by the invariance of $\mathcal{L}$ under local (i.e., spacetime-dependent) phase transformations known as gauge transformations, and transitions between states are constrained by the quantum numbers and conserved current associated with that interaction.

The invariance requirement for $\mathcal{L}$ necessitates the transformation of spacetime derivatives ∂_μ in the Lagrangian via the generators of the symmetry group. These generators correspond to gauge fields whose excitations are the gauge bosons that mediate the fundamental interaction. As a simple illustration, in the QED Lagrangian which obeys the symmetry of the unitary group U(1):

$$\begin{aligned} \mathcal{L}_{QED} &= \bar{\psi}(i\gamma^\mu D_\mu - m)\psi - \frac{1}{4}F^{\mu\nu}F_{\mu\nu} \\ &= \bar{\psi}(i\gamma^\mu D_\mu - m)\psi - \frac{1}{4}(\partial^\mu A^\nu - \partial^\nu A^\mu)(\partial_\mu A_\nu - \partial_\nu A_\mu) \end{aligned} \tag{1.3}$$

the gauge covariant derivative D_μ is given by:

$$D_\mu = \partial_\mu + ieA_\mu \tag{1.4}$$

where A_μ is the field corresponding to the photon, the gauge boson of the QED theory. Once A_μ is introduced and the gauge covariant derivative is thus defined, $\mathcal{L}$ is invariant under all U(1) transformations $\psi \to \psi' = e^{i\chi(x)}\psi$.

Developing field theory descriptions for the fundamental forces has led to predictions of the existence of many new particles. Over the decades, high-energy physics experiments have confirmed the existence of these particles. The Standard Model is the quantum field theory that provides the most successful description to date of all experimentally observed fundamental particles and their interactions. The following section presents a summary of the current state of the Standard Model and its categorization of all known fundamental particles.

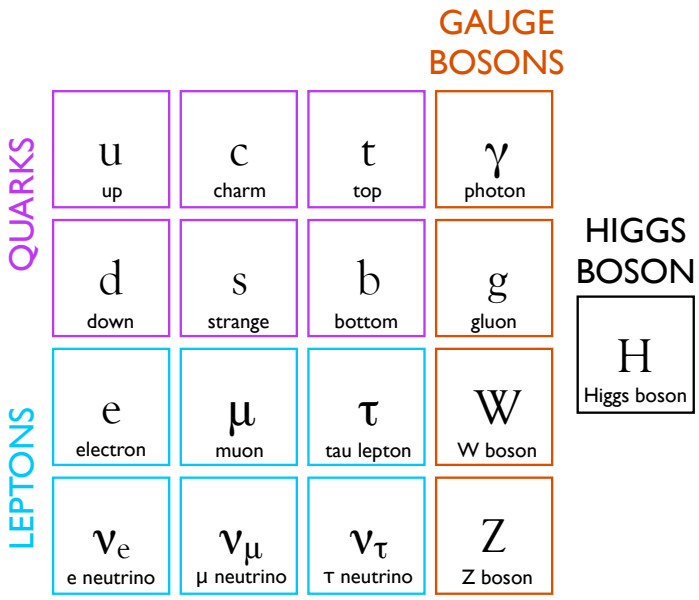

Figure 1.1. Standard Model particles.

1.1.2 Particles of the Standard Model

A tabular display of all experimentally observed fundamental particles is shown in Figure 1.1.

Leptons and quarks are fermions – particles with half-integer spin whose dynamics obey the Dirac equation (see Equation 1.2). Leptons fall into three generations called "flavors," and have electric charge equal to integer multiples of the elementary charge $e = 1.6{\cdot}10^{-19}$ C. The negatively charged leptons are the electron, the muon, and the tau (in increasing order of mass), and each is associated with an extremely light, neutral particle called a neutrino. Quarks have fractional multiples of the elementary electric charge and also possess another quantum property known as "color" charge, whose implications will be explained shortly. There are three generations of quarks and a total of six different quark flavors (two per generation). For each quark and lepton, there exists an associated antiparticle.

The strong, electromagnetic, and weak interactions occur via the exchange of gauge bosons, which obey Bose-Einstein statistics and have spin 1. Each type of interaction involves the coupling of particles to the gauge field associated with that interaction. The

theory describing the strong force is called quantum chromodynamics, or QCD. Quantum electrodynamics (QED) describes the electromagnetic force. At energies above $\mathcal{O}(100)$ GeV, the electromagnetic interaction force unifies with the weak force, and the unified force is described by the electroweak theory.

The Standard Model belongs to the symmetry group SU(3)×SU(2)×U(1). The QCD Lagrangian $\mathcal{L}_{QCD}$ obeys the symmetry of the special unitary group SU(3), while the electroweak portion $\mathcal{L}_{EWK}$ obeys SU(2)×U(1) symmetry.

The strong force is mediated by gluons, which are colorless, electrically neutral, and massless; only quarks and gluons, which possess nonzero color charge, can participate in strong interactions. In QCD, the eight generators of the SU(3) group give rise to eight gauge fields G^a_μ, whose linear combinations correspond to gluons. The conserved quantity in QCD interactions is "color" charge. An unusual feature of the strong force is that as the momentum transfer of the interaction increases, the strength of the interaction decreases. Thus, for high-energy interactions, the QCD coupling is small enough that perturbation theory can be applied to Feynman diagrams and quarks can be treated like free particles – a property known as asymptotic freedom. The behaviour of the strong force at low energies, where QCD becomes non-perturbative, is still not well understood; one consequence of the low-energy scale behaviour of the strong force is color confinement, which means that quarks and antiquarks cannot be found free but can only exist in bound states called hadrons. The SU(3) invariance of hadronic wavefunctions restricts the only possible hadronic states to be SU(3) singlets – i.e., states with a net zero color charge.

The quarks that make up a hadron, the gluons that bind them, and the fleeting quark-antiquark pairs that these gluons produce are collectively known as partons. The probability density for each parton to be found with a fraction x of the total hadronic 4-momentum is described by its parton distribution function (PDF), which is determined by the strong interactions among the various partons in the hadron. In practice, PDFs cannot be calculated from theory alone because of the non-perturbative nature of QCD interactions at low energies; thus, they can only be measured experimentally [4].

Any particle with electric charge can participate in electromagnetic interactions, which

are mediated by electrically neutral, massless photons. $W^{\pm}$ or Z^0 bosons are the carriers of the weak force, which is responsible for such processes as nuclear decays (they will hence be referred to as W and Z, dropping the charge superscript unless it is necessary to mention their charges explicitly). W^+ and W^- are one another's antiparticles, while Z and the photon are their own antiparticles. Unlike the gluon and photon, the W and Z bosons are massive.

Low-energy weak processes such as beta decays were first described by Fermi via a simple four-point interaction [5]. This, however, does not explain the experimental observation of parity violation in weak decays, and the predicted cross-section for weak decays blows up for high-energy processes ($q^2 > \mathcal{O}(100\text{ GeV})^2$) in the four-point interaction model. Electroweak theory developed in response to these issues; it predicted that weak interactions occur via parity-violating axial vector currents mediated by massive vector bosons [6]. In electroweak theory, the gauge fields B_μ (from the U(1) group) and $\mathbf{W}^1{}_\mu$, $\mathbf{W}^2{}_\mu$, and $\mathbf{W}^3{}_\mu$ (from the SU(2) group) give rise to the electroweak gauge bosons [7].

However, a problem arises from the fact that the Lagrangian for the electroweak gauge bosons can only be invariant under SU(2)×U(1) transformations if the masses of its gauge bosons are zero. Although the photon is known to be massless, the W and Z bosons are clearly not. A prediction for the mass of the W was first derived from measurements of the lifetime of the muon [3], and the W and the Z were later both discovered in e^+e^- collisions in the LEP experiment at CERN [8, 9].

The massive gauge boson paradox is resolved by the concepts of spontaneous symmetry breaking and the Higgs mechanism [3]. The latter involves the introduction of a complex scalar field whose vacuum expectation value is not zero but instead one of multiple nonzero minima of the scalar field potential. The choice of one of these vacuum expectation values breaks the symmetry of the scalar potential. Figure 1.2 shows an illustration for the simplified case of a real scalar potential with two local minima. When the Lagrangian of such a scalar field is expressed in terms of a perturbation of the field about its vacuum expectation value, this results in a mass term for the perturbation that corresponds to a scalar boson called a Goldstone boson.

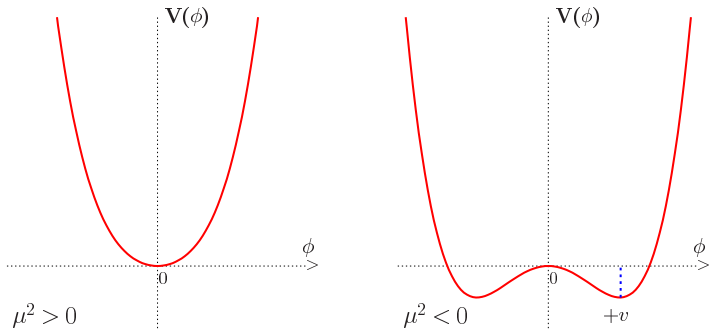

Figure 1.2. Illustration of spontaneous symmetry breaking in the case of a real scalar field ϕ, with potential $\mathrm{V}(\phi) = \frac{1}{2}\mu^2\phi^2 + \frac{1}{4}\lambda\phi^4$. For $\mu^2 > 0$ (Left), the potential has a minimum at zero and thus the vacuum expectation value of the field is zero. For $\mu^2 < 0$ (Right), the potential has two nonzero minima at $v = \pm\sqrt{-\frac{\mu^2}{\lambda}}$; in nature, the symmetry is broken by the choice of one of these two possible values for the vacuum expectation value of the field. Image copied from [10].

To complete the picture, the spontaneous symmetry of the scalar field is embedded in the SU(2)×U(1) symmetry. The simplest model of electroweak symmetry-breaking by the Higgs mechanism requires two complex scalar fields. When the combined Lagrangian of the scalar fields and electroweak gauge fields is expressed as an expansion about the chosen vacuum expectation values of the scalar fields, the Lagrangian ends up with terms quadratic in the B_μ, $\mathbf{W}^1{}_\mu$, $\mathbf{W}^2{}_\mu$, and $\mathbf{W}^3{}_\mu$ fields, which are the mass terms for the gauge bosons. Via an appropriate gauge transformation, the Goldstone bosons resulting from the symmetry breaking disappear by being absorbed into the longitudinal degree of freedom of the $\mathbf{W}^i{}_\mu$ fields, and the mixing of the four gauge fields results in the $W^\pm$ bosons, which are linear combinations of $\mathbf{W}^1{}_\mu$ and $\mathbf{W}^2{}_\mu$, and a Z boson and a photon, both of which are linear combinations of $\mathbf{W}^3{}_\mu$ and B_μ. When this mixing is accounted for in the Lagrangian, the only mass terms that remain are the ones for the W and Z bosons, while the photon is massless. The existence of the Higgs field also generates the masses of the fermions via Yukawa interactions between fermions and the Higgs field.

Thus, the electroweak theory predicts the existence of a Higgs boson. This is the final particle is represented in Figure 1.1, the only known scalar particle in the Standard Model. The experimental search for this Higgs boson has carried on for decades after its existence was first predicted, and has finally culminated in its discovery at the LHC

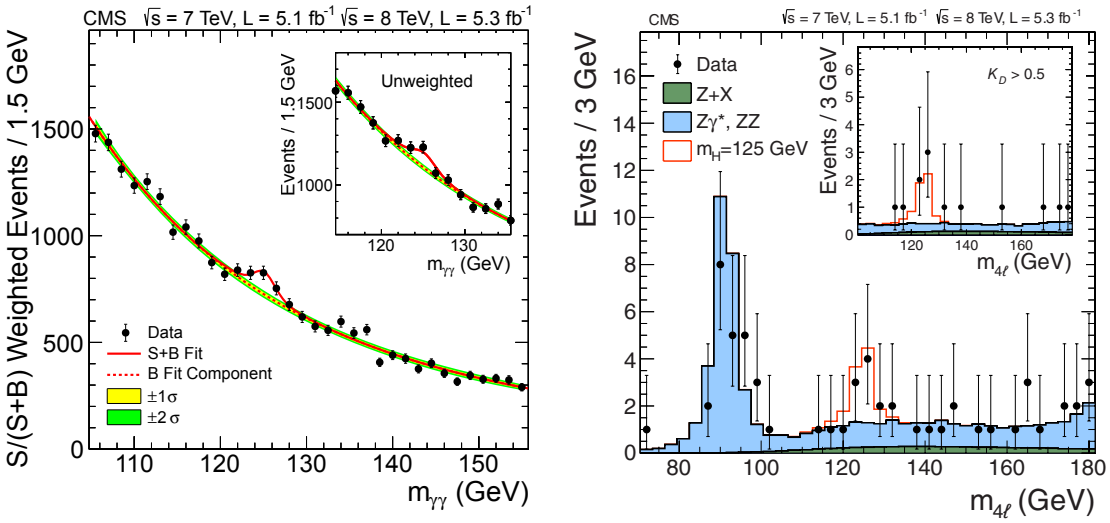

Figure 1.3. (Left) Observation of the Higgs resonance in the $H \to \gamma\gamma$ channel at the CMS experiment. The reconstructed diphoton events were grouped into different categories based on different kinematic variables before being combined into the final result; the inset shows the same diphoton mass spectrum, in which the events have not received a weight based on the signal-to-background ratio of their category. (Right) Observation of the Higgs resonance in the $H \to ZZ \to 4l$ channel. The inset shows the same spectrum after a further selection step using the probability ratio K_D of the signal and background hypotheses. The combined $H \to \gamma\gamma$ and $H \to ZZ \to 4l$ channels yield a best-fit Higgs mass of 125 $\pm$0.4 (stat.) $\pm$0.5 (syst.) GeV. [11]

collider at CERN [12, 11] (see Figure 1.3), providing solid validation for this last major prediction of the Standard Model.

1.2 Deficiencies of the Standard Model

Although the Standard Model has been successful in describing a wide range of experimental results, there is clear evidence that it is not complete. To name a few of its shortcomings [4]:

- **Gravity:** The four fundamental forces are the strong force, electromagnetism, the weak force, and gravity. The Standard Model accounts for the first three, but the way in which gravity, which is 10^{32} times weaker than the weak force, factors into the Standard Model is still unknown.

- **Neutrino oscillations:** The Standard Model treats neutrinos as massless particles.

However, there is strong experimental evidence to the contrary. The observation of neutrino flavor oscillations, which cannot occur if neutrinos were massless, suggests that the observed electron, muon, and tau neutrino flavors are not in fact mass eigenstates, but rather that the observed neutrino flavor states are superpositions of mass eigenstates. What the mass eigenvalues are and why they are so small, are mysteries that remain to be resolved.

- **Dark matter and energy:** Matter constitutes only about 30% of the total mass-energy density of the universe; the remainder is comprised of "dark matter" and "dark energy," but their exact nature is unknown, and their presence is only deduced indirectly from their gravitational effects, since dark matter does not emit or absorb electromagnetic radiation. Thus, it is suspected that dark matter must be made up of some different type of particle not accounted for by the Standard Model.

- **Grand unification:** The running coupling constants of the strong, electromagnetic, and weak interactions generally approach one another at increasingly higher energy scales. This suggests that these three fundamental interactions (and perhaps also gravity) might be manifestations of one unified field theory and thus may unify at some high energy scale; some theories predict this scale to be on the order of 10^{16} GeV.

- **The hierarchy problem:** The Standard Model predicts that the mass of the Higgs boson receives loop corrections from leptons, quarks, and whatever other particles may couple to the Higgs; Figure 1.4 shows two example Feynman diagrams for loop corrections from fermions and scalar particles.

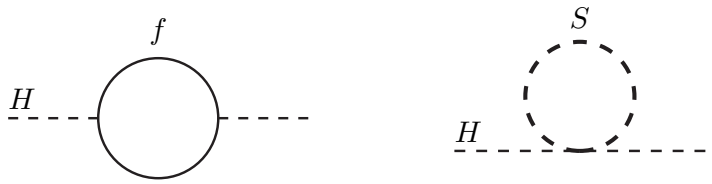

Figure 1.4. One-loop corrections to the Higgs mass from a fermion f (Left) and from a scalar S (Right) [13].

To first order, the corrections to the squared Higgs mass from the diagrams in Figure 1.4 are:

$$\Delta m_H^2 = -\frac{|\lambda_f|^2}{8\pi^2}\Lambda_{cutoff}^2 + \frac{\lambda_S}{16\pi^2}[\Lambda_{cutoff}^2 - 2m_S^2 \ln(\frac{\Lambda_{cutoff}}{m_S})] + \cdots \tag{1.5}$$

where λ_f and λ_S are the Higgs couplings to f and S respectively, and Λ_{cutoff} is the cutoff energy scale at which the Standard Model breaks down and new physics takes over. Examples of Λ_{cutoff} include the grand unification scale ($\mathcal{O}(10^{14}$ GeV)) at which the running couplings of the strong, electromagnetic, and weak forces unify, or the Planck scale ($\mathcal{O}(10^{19}$ GeV)) at which the quantum effects of gravity become significant. Thus, the Higgs mass is expected to receive corrections that are many orders of magnitude larger than its experimentally measured value. The fact that the measured Higgs mass is so much smaller than the predicted corrections is surprising, as it implies the existence of some mechanism, unexplained by the Standard Model, by which these corrections are reduced or cancelled. This constitutes what is known as the hierarchy problem of the Standard Model.

These deficiencies indicate that there is physics beyond the Standard Model, and that more elements are needed to provide an accurate description of nature.

1.3 Supersymmetry

The theory of supersymmetry provides a means to address the hierarchy problem. This theory posits that there exists a certain symmetry – called supersymmetry – that relates fermions to bosons [13]. For each fermion, there exists a corresponding bosonic parter particle, and likewise for each boson there exists a fermionic partner; such partner particles are referred to as superpartners. Under a supersymmetric transformation, a fermion turns into its superpartner and vice versa. For each loop correction to the Higgs mass from a Standard Model particle, there would be another loop correction from its supersymmetric partner. Since fermionic loops and bosonic loops have opposite signs, the loop corrections from Standard Model particles are neatly cancelled by loop corrections from their supersymmetric partners.

As supersymmetry predicts the existence of a superpartner for every Standard Model particle, it is of great interest to look for experimental evidence of these superpartners, none of which have yet been observed in nature. The question is where to look, since there is nothing that specifies the masses of supersymmetric particles or the other numerous other free parameters in the theory of supersymmetry. These free parameters must be constrained by experimental measurements.

The masses of supersymmetric particles must be different from the masses of their Standard Model counterparts, otherwise the supersymmetric particles would have been observed already. In order to have this inequality in masses, supersymmetry must be a broken symmetry. There are many models for the mechanism of supersymmetry breaking; but in general, in order to achieve the stabilization of the Higgs mass, the symmetry breaking scale is expected to be on the order of 1-10 TeV [14]. Since the mass splittings between Standard Model particles and their superpartners are determined by the supersymmetry breaking parameters, the masses of the lightest supersymmetric particles are expected to be around the same scale as well. Thus, if supersymmetric particles exist, it could be possible to discover them in high-energy collider experiments such as the LHC at CERN.

The simplest model incorporating supersymmetry into the Standard Model is the Minimal Supersymmetric Standard Model (MSSM). The details will not be covered here, but a comprehensive review can be found in [13]. The Higgs sector of the MSSM consists of two chiral Higgs supermultiplets H_u and H_d, of which the former couples to up-type quarks and the latter couples to down-type quarks and charged leptons. The superpotential W_{MSSM} of the MSSM, which defines the most general non-gauge interactions for the superfields, is:

$$W_{MSSM} = \bar{u} y_u Q H_u - \bar{d} y_d Q H_d - \bar{e} y_e L H_d + \bar{\mu} H_u H_d \tag{1.6}$$

However, the MSSM comes with deficiencies of its own. One problem, called the μ-problem, stems from the fact that the μ parameter in Equation 1.6, which must have the dimension of mass, must be of the order of magnitude of the electroweak scale in

order to provide the Higgs doublets with vacuum expectation values on the order of the electroweak scale. However, its magnitude is expected to be more naturally close to new physics cutoff scales such as the Planck scale, which are significantly larger. Fine-tuning its magnitude to that of the electroweak scale thus seems arbitrary and without theoretical motivation.

Another problem in the MSSM is that the mass of the stop (the supersymmetric partner of the top quark) must be quite large – on the order of 1-10 TeV – in order for the lightest CP-even Higgs boson to have a mass greater than 115 GeV without large stop mixing [15]. However, such large stop masses would produce large loop corrections to the Higgs mass and quartic coupling, thus reviving the need for fine-tuning. This is referred to as the "little hierarchy problem" [16].

The μ-problem is avoided in the Next-to-Minimal Supersymmetric Standard Model (NMSSM) [17], which extends the Higgs sector of the MSSM by a scalar Higgs singlet field S and ultimately predicts a total of seven Higgs bosons – a pair of charged ones, three neutral scalars, and two neutral pseudoscalars. The NMSSM superpotential is:

$$W_{NMSSM} = W_{MSSM} + \lambda S H_u H_d + \frac{1}{3}\kappa S^3 + \frac{1}{2}\mu_s S^2 \tag{1.7}$$

The coupling of the singlet field S to the doublet fields naturally generates an effective μ term via the expectation value of S. Thus, $\mu = \lambda < S >$, with the desired magnitude near the electroweak scale. The singlet field also provides a mechanism to raise the mass of the lightest CP-even Higgs without requiring large stop masses, thus significantly reducing the little hierarchy problem [15, 18, 19] The fact that the NMSSM circumvents these problems, combined with the fact that searches for evidence of the MSSM at the LHC have so far been fruitless, makes the NMSSM of great interest to study. The new physics search presented in this dissertation is a probe into the Higgs sector of the NMSSM.

Chapter 2

Experimental apparatus: the LHC and the CMS detector

The Large Hadron Collider (LHC) is currently the world's most powerful particle collider. Constructed by the European Organization for Nuclear Research (CERN) between 1998 and 2008, its goal is to elucidate the electroweak symmetry breaking mechanism and probe physics beyond the Standard Model with the highest collision energies yet attained.

The research described in this dissertation was conducted on data collected from 8-TeV proton-proton collisions in the LHC by the Compact Muon Solenoid (CMS) detector. This chapter presents a brief overview of the LHC, followed by a description of the CMS detector, its component subsystems, and its data acquisition system.

2.1 Large Hadron Collider

The Large Hadron Collider (see [20] for a detailed description) is a circular accelerator and collider of high-energy particles, straddling the French-Swiss border near Geneva, Switzerland. This synchrotron consists of 2 rings for counter-circulating proton or ion beams, 27 km in circumference and located at depths as low as 175 m underground in the tunnel previously occupied by the LEP collider ring.

A system of 1,232 superconducting dipole magnets, cooled to temperatures below 2 K with superfluid helium and generating magnetic fields above 8 T, steer the two particle beams (either protons or lead ions) along their trajectory through the accelerator ring;

each dipole magnet has a twin-bore design with two apertures that allow both beams to pass through. Quadrupole magnets are used at various points along the ring to keep the beams focused and reduce dispersion, and also at four interaction points (IPs) to cause the beams to intersect and collide. These interaction points are where the ATLAS, CMS, LHCb, and ALICE detectors are located.

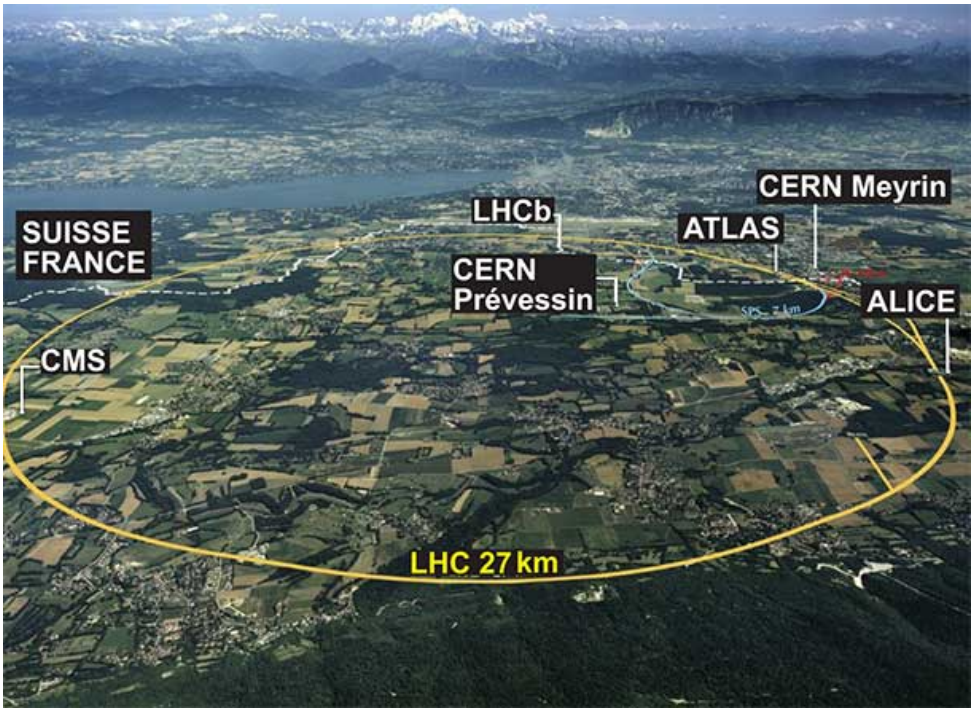

Figure 2.1. Aerial view of the Swiss-French border near Geneva, with the path of the LHC ring superimposed [21].

Protons, obtained by ionizing hydrogen gas, are accelerated first via a linear accelerator to a kinetic energy of 50 MeV and injected into the Proton Synchrotron Booster, which accelerates them to 1.4 GeV. From there, they are injected into the Proton Synchrotron, which accelerates them further to 25 GeV, and then into the Super Proton Synchrotron, which brings them to an energy of 450 GeV before finally injecting them into the LHC ring, in which they are accelerated to the desired center-of-mass energy for collisions.

Each beam circulating in the LHC contains 2808 bunches of protons, with a bunch spacing of 25 ns. Each bunch has a maximum intensity of $1.15{\cdot}10^{11}$ protons, limited by the mechanical aperture of the LHC and by the tuning performed to control the linear

effects of beam-beam interactions at the collision point.

For a given physics process, the event rate N from a beam collision is given by:

$$N = \mathcal{L} \cdot \sigma_{process} \tag{2.1}$$

where $\mathcal{L}$ is the luminosity of the colliding beams (the number of collisions per unit time per unit cross-sectional area) and $\sigma_{process}$ is the cross-section for the process of interest. The design luminosity of the LHC is driven by the needs of the experiments that take place around its ring.

The LHC is home to several experiments: the general-purpose detectors ATLAS (A Toroidal LHC Apparatus) and CMS (Compact Muon Solenoid), the detector LHCb that focuses on B physics, the detector TOTEM that focuses on elastic proton scattering, and the heavy-ion detector ALICE. The highest luminosity requirements come from ATLAS and CMS, which aim for a peak luminosity of 10^{34} $\text{cm}^{-2}\text{s}^{-1}$; by the logic of Equation 2.1, a high luminosity is needed for obtaining an event rate high enough to be sensitive to Higgs boson production and to rare physics processes beyond the Standard Model.

During a physics run, the collider luminosity progressively decreases, primarily due to beam losses from collisions. Factoring in the effects of all contributions to beam losses, the luminosity has an exponential decay constant of roughly 14.9 h. It is not practical to keep the beams colliding until they have exhausted all their energy, so after a certain point (the optimum run time is around 12 hours), the degraded beams are dumped, so that fresh beams can be generated at peak luminosity to refill the LHC ring.

At peak operation, the total LHC beam current is 0.584A, resulting in a stored energy of 362 MJ. Combined with the total electromagnetic energy of 600 MJ stored in the magnet system, the LHC during operation has a total stored energy of over 1 GJ. All of this energy has to be absorbed safely in a beam dump; this is achieved by a dedicated beam dumping system of magnets to divert the beams out of the ring and thirty-five 24-ton blocks of carbon for absorbing the beams.

The LHC has been designed to collide protons at a maximum center-of-mass energy of 14 TeV. During its first run, it operated at center-of-mass energy 7 TeV from 2010-

2012, and then at 8 TeV from 2012-2013. After a long shutdown for planned repairs and maintenance, the LHC was restarted in 2015, operating at 13 TeV, with the intention of eventually increasing to 14 TeV.

2.2 Compact Muon Solenoid

The Compact Muon Solenoid (CMS) detector is a nearly hermetic general-purpose detector at the LHC, gathering data from the collision of proton-proton and heavy ion beams to study a wide range of physics processes. This experiment is characterized by a powerful superconducting solenoid magnet that produces a 4 T magnetic field; the paths of charged particles are deflected by the magnetic field, allowing their momenta and charges to be accurately reconstructed from their trajectories.

When the proton bunches cross at the CMS interaction point, hard scattering interactions between the partons of colliding protons can occur. All the particles produced in a bunch crossing pass through and are recorded by a layered system of roughly concentric cylindrical subdetectors. These particles include not only the products of the hard scattering process, but also a lot of extraneous activity as well, such as particles produced via radiation, soft showers coming from partons that do not participate in the hard scattering, and the products of pileup events, which occur when a single bunch crossing results in more than one scattering event [22]. The CMS detector records the passage of all of these particles; this data must be reconstructed with sophisticated software into a picture of the actual event for physics analysis.

The centre of the official CMS coordinate system is at the interaction point, with the x-axis pointing radially inwards towards the centre of the LHC ring and the y-axis pointing vertically upward out of the ground. The z-axis lies along the beam line; as the CMS coordinate system is right-handed, the positive z-axis points in the anticlockwise beam direction. The radial coordinate in the xy plane is denoted by r, and the azimuthal angle ϕ is measured from the x-axis to the y-axis. As a measure of the angle relative to the beam axis (the z-axis), the CMS collaboration uses a quantity called pseudorapidity (η) instead of θ because, as the colliding particles travel at speeds comparable to the speed

of light, the pseudorapidity converges with a Lorentz-invariant quantity called rapidity y:

$$y = \frac{1}{2}\ln(\frac{E + p_z}{E - p_z}) \tag{2.2}$$

The pseudorapidity is defined as:

$$\eta = -\ln[\tan(\frac{\theta}{2})] \tag{2.3}$$

where θ is the polar angle measured from the z-axis to the xy plane.

The entire detector is a cylindrical structure 21.6 m in length (along the z axis) and 14.6 m in diameter. At the center of the ensemble is a silicon tracker subdetector for reconstructing accurately the momenta of charged particles. Encircling it is the electromagnetic calorimeter (ECAL), a homogenous calorimeter that uses scintillating lead tungstate crystals to reconstruct electrons and photons with high energy resolution. Outside the electromagnetic calorimeter is the hadron calorimeter (HCAL), a sampling calorimeter with alternating layers of brass and scintillator that measures the energy deposited in hadronic showers. The solenoid magnet is located outside the hadronic calorimeter, together with an iron yoke system that provides a return flux for the magnetic field. The outermost layer of the CMS detector is the muon tracking system for identifying and reconstructing the trajectories of muons. Figure 2.2 shows the overall layout of the CMS detector's subsystems.

The solenoid magnet and the various CMS subdetectors will be described in more detail in the rest of this chapter. A more complete description of the CMS detector can be found in [24]; unless otherwise specified, all information in this chapter is derived from this source.

2.2.1 Tracker

The tracker detector is the innermost subsystem of the CMS detector. Its purpose is to reconstruct with high resolution the trajectories and momenta of charged particles with transverse momentum 1 GeV and upwards, and to provide high impact parameter resolution for reconstructing the positions of secondary vertices, which are displaced from

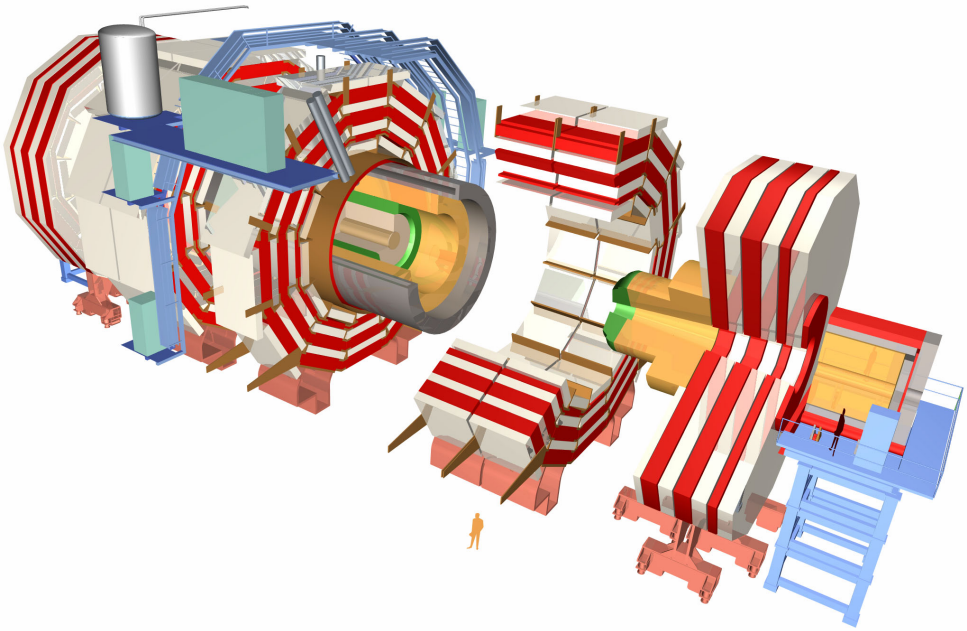

Figure 2.2. Illustration of the various layers of the CMS detector, with a human on the ground to show relative size [23].

the point of collision and generally are a characteristic signature of long-lived particles such as those from heavy-flavor processes. Since the tracker detector receives higher irradiation than any other part of the detector due to its proximity to the beam line and the interaction point (IP), it has been built to survive and operate in a high-flux environment, with thousands of particles passing through its volume every 25 ns when the LHC is running at its design luminosity. Thus the tracker has been designed with these challenges in mind in order to yield good position and time resolution for track reconstruction.

The tracker detector is composed of two subdetectors. The one closest to the beam line is the pixel detector; its sensors are 100 μm $\times$ 150 μm pixels, which receive an occupancy on the order of 10^{-4} per pixel per bunch crossing. The second and larger subdetector is the silicon strip tracker, whose sensors are silicon strips; since it is located at larger radii than the pixel detector, the fluence of particles that reach it is lower and thus the granularity of the silicon strips can be considerably lower than that of the pixel detector, with strip areas ranging from 10 cm $\times$ 80 μm at intermediate radii for the inner strip

tracker (20 < r < 55 cm) and 25 cm × 180 μm for the outer strip tracker (55 < r < 110 cm), resulting in an occupancy of about 2-3% for the inner tracker and 1% for the outer tracker. These subsystems are illustrated in Figure 2.3. The overall envelope of the tracker volume measures 5.8 m in length along the z axis and 2.5 m in diameter.

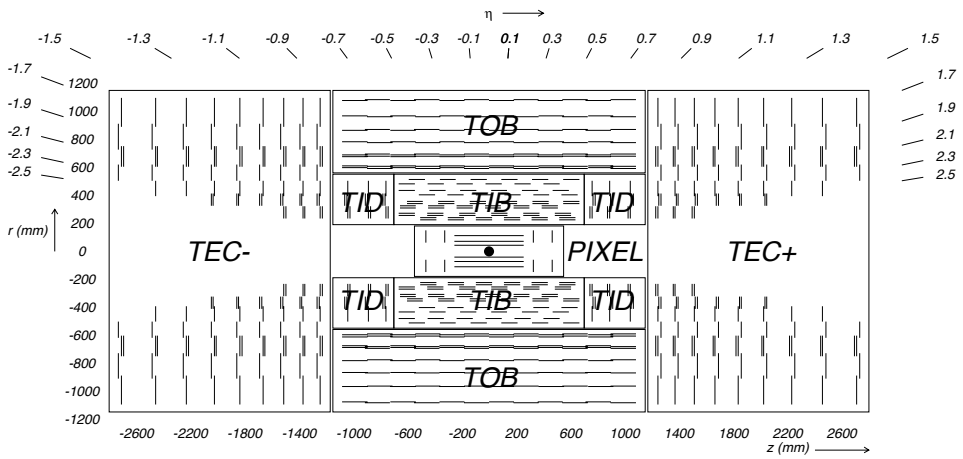

Figure 2.3. Layout of the CMS tracker detector, showing the pixel detector and the subsystems of the silicon strip tracker; each line represents a detector module [25].

In total, the CMS tracker detector is comprised of 200 m^2 of active silicon sensors, providing a coverage of up to $|\eta| < 2.5$. Because of the high particle fluence passing through the tracker volume during operation, radiation-hard sensors and electronics are required to withstand the radiation dosage; also, an efficient system system of cooling tubes carrying chilled liquid C_6F_{14} pervades the tracker volume, keeping it at or below -10° C during operation, thus minimizing radiation damage to the sensors caused either by direct irradiation or by the annealing of radiation-induced defects in the silicon crystal structure through thermal agitation.

Energy loss in the tracker material via multiple scattering, nuclear interactions, gamma conversions, and bremsstrahlung can all interfere with tracking efficiency. Also, bremsstrahlung and gamma conversions in the tracker material will adversely affect the detection and identification of electrons and photons in the electromagnetic calorimeter outside the tracker. Thus, the distribution of materials in the tracker, also referred to as the material budget of the tracker, balances the need for mechanical support for the detector, and the

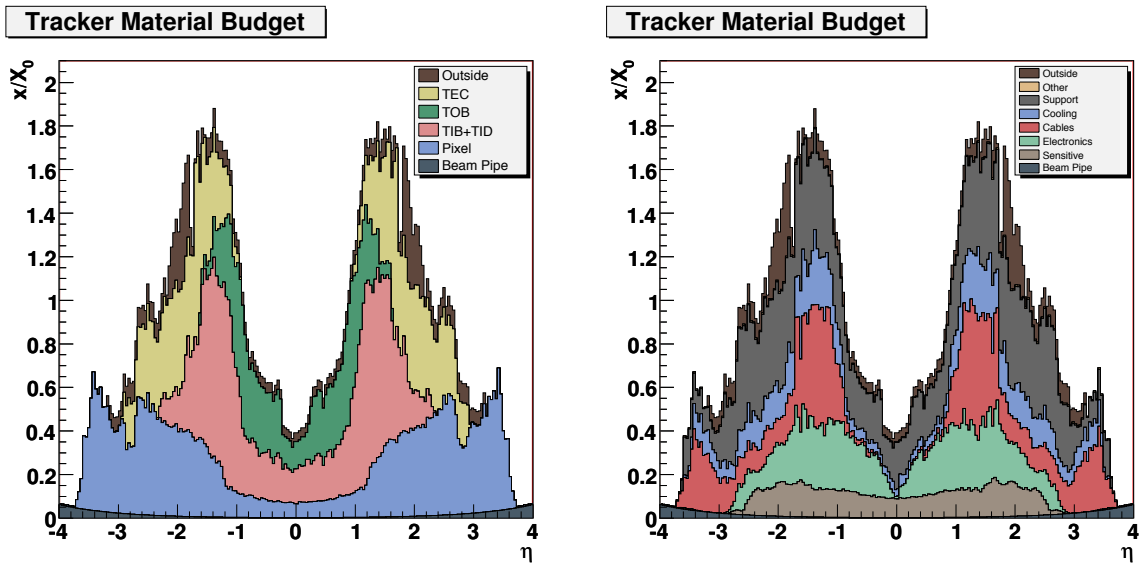

Figure 2.4. Material budget of the tracker detector, displayed in terms of fractional radiation length versus η, (Left) broken down by contribution from the individual subsystems comprising the tracker detector, and (Right) broken down by contributions from eight different classes of material [25].

need to minimize the material budget. Figure 2.4 displays the tracker material budget in terms of fractional radiation length as a function of η, broken down on the one hand by contributions from the individual subdetectors making up the tracker, and on the other hand by contributions from the different categories of material in the tracker volume (outside, support, cooling, cables, electronics, sensitive material, beam pipe, and other). From these plots, it is clear that the material budget of the tracker does not exceed two radiation lengths in any direction.

2.2.1.1 Pixel detector

The pixel detector is the innermost layer of the tracker detector, with three concentric cylindrical barrel layers complemented by two endcap disks on either side of the interaction point. The barrel layers are located at radii 4.4, 7.3, and 10.2 cm from the beam line, extending out to 2.9 m from the interaction point in the ±z directions. Two endcap disks are located at z = ±34.5 cm and ±46.5 cm, with inner radius 6 cm and outer radius 15 cm. Both the barrel cylinders and endcap disks are split in halves along the y axis for ease of extraction and access.

The pixel sensors consist of 52×80 high-dose n-type pixels implanted in a high-resistance n-type substrate of 320 μm thickness on a sensor plate. Each pixel is bump-bonded onto a readout chip (ROC) connected to the sensor plate.

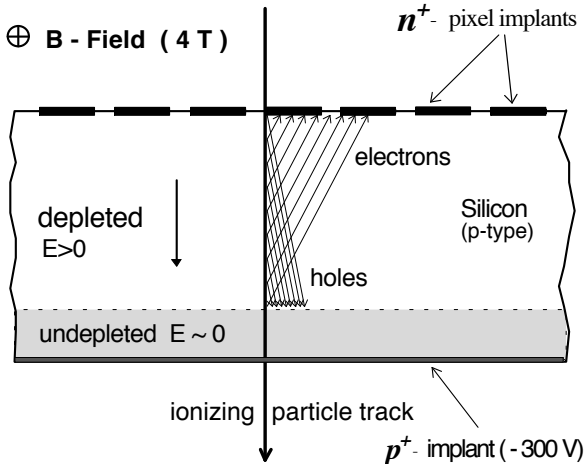

Figure 2.5. Illustration of charge sharing in a pixel sensor [25].

When a charged particle passes through a silicon sensor, it induces charge carriers in the n-doped silicon substrate; traveling towards the pixels to be collected, electrons undergo a signficant Lorentz drift (roughly 32°) due to the 4 T magnetic field along the z axis through the tracker volume, and thus the signal current ends up being spread over adjacent pixels. This process, known as charge sharing, is illustrated in Figure 2.5.

The ROC collects and amplifies the analog signals from the pixels, storing them in a buffer until the arrival of the appropriate readout control and clock signals causes it to pass the sensor signal further down the readout system. The analog signal from the ROC is then digitized by the pixel front end digitizer (pxFED) and passed to the CMS data aquisition system for storage. Analog signals above a certain threshold – to screen out background noise – are clustered into "hits" by pattern recognition algorithms.

Each pixel sensor has an area of 100 μm × 150 μm. The nearly square design is intended to provide good cluster size and hit resolution in both rϕ and z for the barrel and rϕ and r for the endcaps, as both of these coordinates are needed for measuring track impact parameters. In the barrel, the sensors are arranged 2 × 8 on rectangular modules (or 1 × 8 along the edges of the half-cylinders). In each endcap half-disk, the sensors are

arranged into groups called plaquettes (there are five different sizes: 1×2, 2×3, 2×4, 1×5, and 2×5), and the plaquettes themselves are arranged into groups of 3 or 4 on a wedge-shaped panel. The panels are mounted onto trapezoidal blades on the endcap disk, with two panels per blade (one on each side); the arrangement of the plaquettes ensures full coverage of the area of the panel. An illustration of the components and shapes of barrel modules and endcap blades can be found in Figure 2.7. Altogether, the pixel detector covers a pseudorapidity range of $|\eta| < 2.5$.

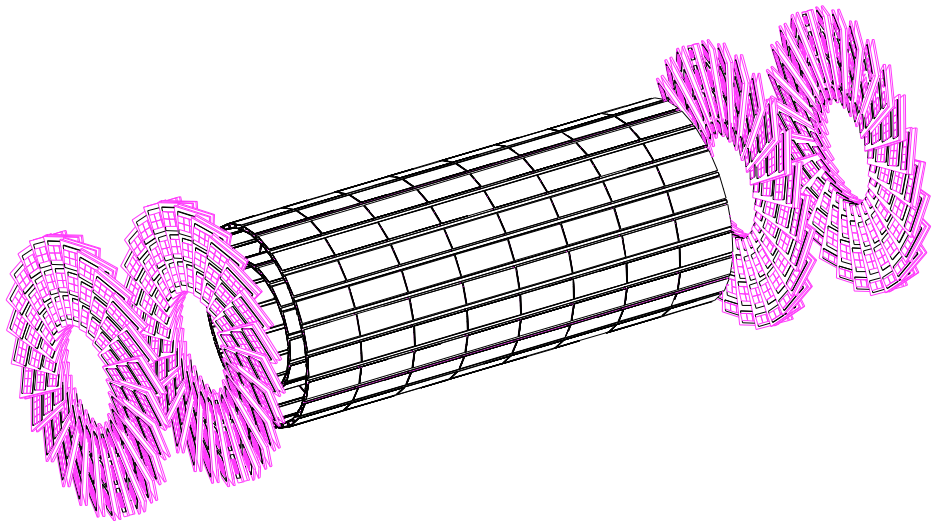

Figure 2.6. Layout of the pixel detector: barrel layers and endcap disks [26]. The magenta wedges on the endcap disks are carbon-fibre blades, which hold plaquettes (rectangular arrangements of pixel sensors that come in five different sizes). The black rectangles in the barrel layers are the barrel modules (2×8 rectangular arrangements of pixel sensors), mounted on rectangular carbon-fibre blades.

In the barrel, the normal direction of the pixel cells points along the radial direction, perpendicular to the magnetic field, so the Lorentz drift is along the rϕ direction. To induce charge sharing in the pixel endcaps, whose normal points along the magnetic field axis, the blades are rotated at a 20° angle about their radial axis in a turbine-like geometry. Hit resolution is improved by interpolating between signals from multiple neighbouring pixels due to charge sharing. All in all, resolutions of 10-20 μm are achieved in the pixel detector. The hit efficiency of each sensor is the probability of reconstructing a cluster if that sensor has been traversed by a charged particle; Figure 2.8 shows that the average

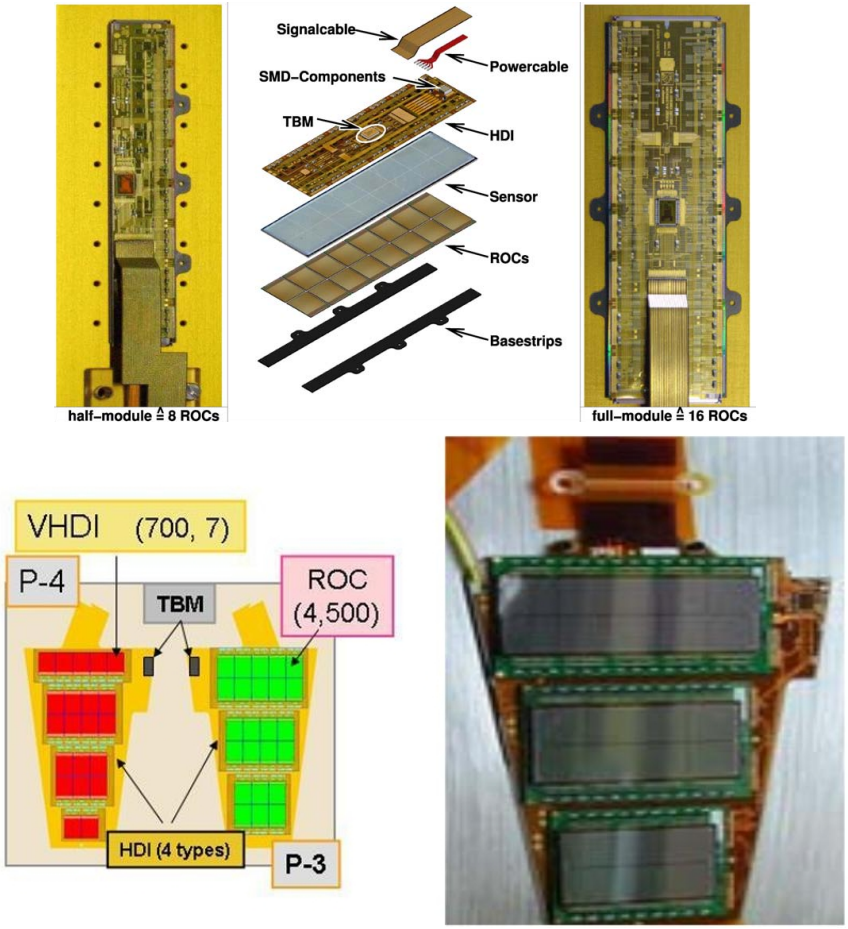

Figure 2.7. (Top) Components and shape of a barrel pixel module. (Bottom) Components and shape of an endcap pixel blade, showing the five different plaquette sizes. [24]

hit efficiency of each barrel layer and endcap disk, measured in proton-proton collisions at 7 TeV, is well above 90%.

2.2.1.2 Silicon strip detector

Outside the pixel detector, at $r = 20$ cm to 116 cm about the beam line and extending 118 cm in the +z and -z directions, lies the silicon strip detector. It is composed of 3 main subsystems: the tracker inner barrel and disks (TIB and TID), the tracker outer barrel (TOB), and the tracker endcaps (TEC).

To optimize coverage, the silicon microstrip sensors in any given barrel layer or endcap disk are positioned to partially overlap with one another, thus resulting in a nonzero pitch

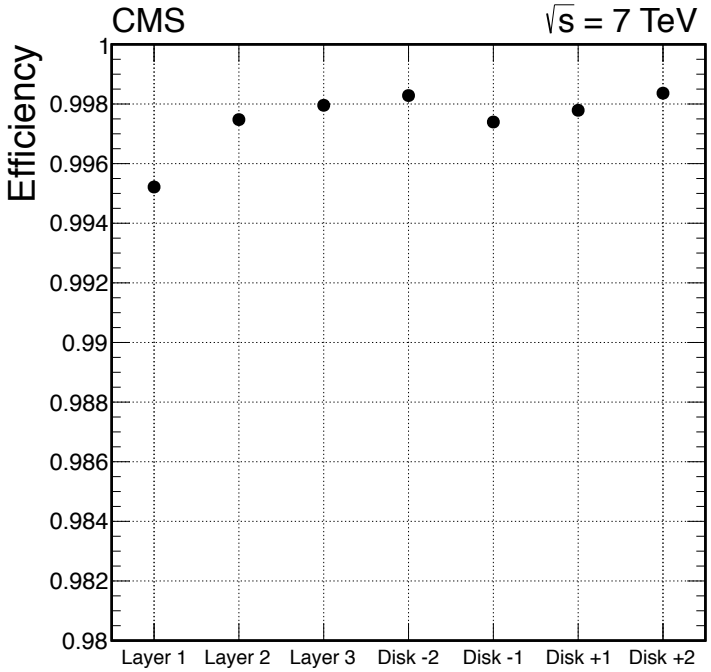

Figure 2.8. Average hit efficiency in CMS pixel barrel layers and endcap disks, recorded in 7 TeV data from the LHC [27]. Defective modules were excluded from the efficiency calculations.

with respect to the surface to which they are attached. The silicon strip tracker contains a total of 15,148 modules, each bearing one 320 μm-thick sensor or two 500 μm-thick sensors. All in all, there are 29 different module designs, differing in the size of their active area based on the number and size of the sensors. Figure 2.9 shows an example of a module from the TEC, housing two sensors. The silicon microstrip sensors have a single-sided p-on-n design. Signal currents are amplified and stored by a custom integrated circuit called an APV25 before being transmitted via optical fibers through the readout system for digitization and storage.

The TIB consists of 4 barrel layers 140.0 cm in length, with radii of 255.0 mm, 339.0 mm, 418.5 mm, and 498.0 mm, altogether providing up to 4 r-ϕ measurements per particle trajectory. The sensors are silicon microstrips with a thickness of 320 μm, lying parallel to the beam axis, with a pitch of 80 μm on the inner two layers and 120 μm on the outer two layers, yielding a single-point resolution of 23 μm for the inner two layers and 35 μm for the outer two layers.

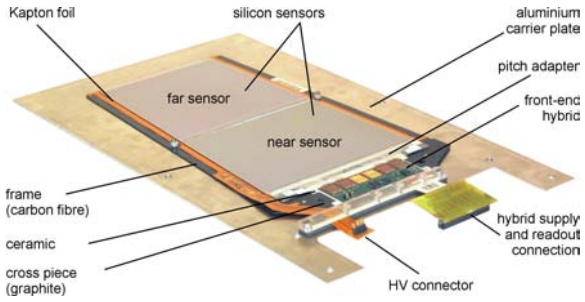

Figure 2.9. Photo of a TEC module, composed of 2 sensors [24].

The TID is comprised of six endcap disks, three at each end of the TIB between ±80.0 cm and ±90.0 cm on the z axis; each disk consists of three support rings from r = 200 μm to 500 μm. Silicon microstrips, similar to the ones used in the barrel, lie radially on the disks, with a pitch varying from 100 μm to 141 μm. While the pixel detector is split down the y axis into half-cylinders for ease of installation, access, and independent testing, the TIB and TID are split into half-shells along the x axis for similar reasons. Two carbon-fibre service cylinders are coupled to the ±z ends of the TIB, providing a route to the TIB shells for service cables originating from a service distribution disk called a margherita, and also housing the TID.

The TOB surrounds the TIB with 6 barrel layers, reaching to an outer radius of 116 cm and spanning 118 cm in the ±z directions. The silicon microstrip sensors here are 500 μm thick and have a pitch of 183 μm in the inner four layers and 122 μm in the outer two layers, yielding a single-point resolution of 53 μm and 35 μm respectively in those layers. On either end of the TOB, the TEC extends radially from 220 mm to 1135 mm and in the ±z direction from ±1240 mm to ±2800 mm; each side consists of an assembly of 9 disks with up to 7 rings bearing silicon microstrip sensors, plus 2 extra disks that serve as front-back termination. The microstrip sensors used here have a thickness of 320 μm in the four rings closest to the TOB and 500 μm in the remaining outer rings, with a pitch varying from 97 to 184 μm.

The average hit efficiencies for the various subsystems of the strip detector is shown in Figure 2.10. These efficiencies are close to 100%, and even when defective modules are

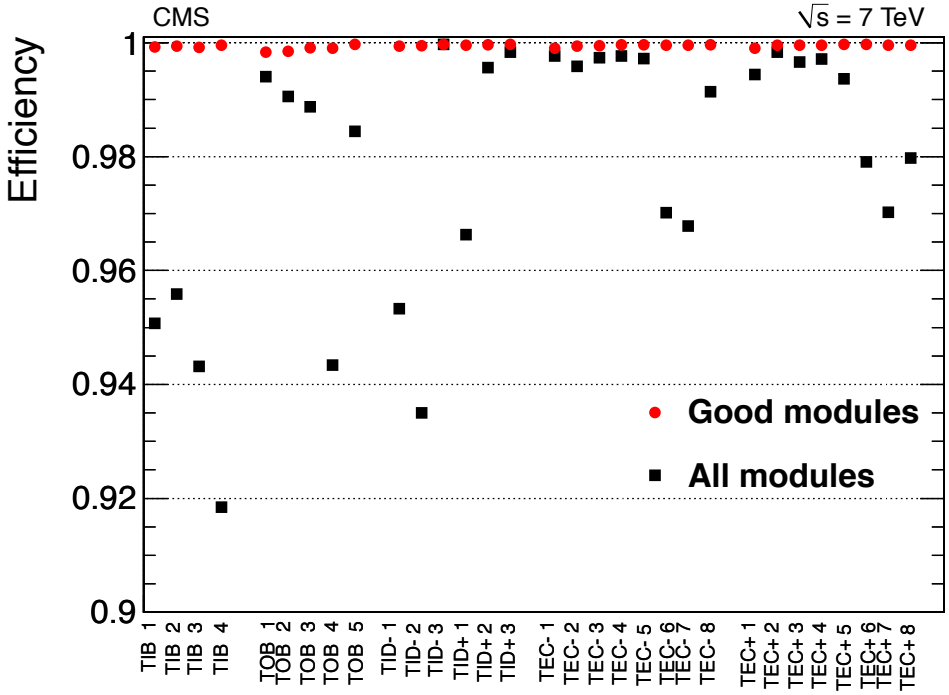

Figure 2.10. Average hit efficiency for the layers and endcap disks of the CMS strip detector, recorded in 7 TeV data from the LHC [27]. The red points are efficiencies measured without considering defective modules, while the black efficiency points were measured with defective modules included.

considered, the efficiencies are still above 90%.

2.2.2 Electromagnetic calorimeter

One of the primary goals of the CMS experiment is the elucidation of the Higgs mechanism and measurement of the properties of the Higgs boson. The diphoton final state is an important search channel for the Higgs boson, as such a resonance would show up clearly above the background spectrum in the mass range constrained by the latest measurements from the LEP experiment. Other possible search channels include final states involving electrons or positrons, such as $H \rightarrow WW/ZZ$. Standard Model measurements and searches for new physics signals can also profit from efficient electron and photon identification. Thus, the need for the CMS detector to provide good photon reconstruction and resolution has driven the optimization of the electromagnetic calorimeter.

The hermetic, homogeneous CMS electromagnetic calorimeter (ECAL) is made up of lead tungstate ($PbWO_4$) crystals. Like the inner tracking system, it is composed of

a cylindrical barrel system (with a total of 61,200 crystals) and an endcap system (one endcap disk on each side of the barrel, with 7324 crystals per disk). Its layout is illustrated in Figure 2.11.

The total thickness of the ECAL is larger than 25 radiation lengths χ_0. When electrons or photons pass through, the resulting electromagnetic showers that they generate excite the scintillator atoms, causing them to emit blue-green scintillation light (up to 420-430 nm in wavelength) that is collected, amplified, and read out by photodetectors glued to the ends of the crystals. Adjacent energy deposits above a background-rejecting threshold are sorted by pattern recognition algorithms into groups called clusters, which may be sorted further into larger agglomerations called superclusters.

High granularity is needed for good energy resolution in the ECAL. The choice of $PbWO_4$, with a density of 8.28 g/cm^3, a radiation length of 0.89 cm, and a Moliere radius of 2.2 cm, allows for compact, granular crystals that can resist radiation damage under the high particle fluences that the detector is subjected to throughout its operation.

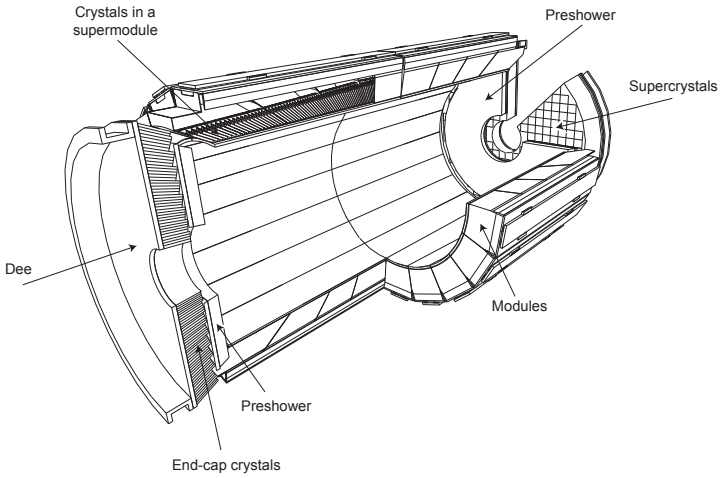

Figure 2.11. Layout of the CMS electromagnetic calorimeter, showing the barrel, endcaps, and preshower [24].

The barrel extends from an inner radius of 1.29 m to an outer radius of 1.77 m. In the barrel, the detector granularity is 360-fold in ϕ and 170-fold in η. The crystals have a truncated pyramidal shape that varies with their position in η, with a cross-sectional area of approximately 22×22 mm^2 at the front face (closest to the beam line) and 26×26 mm^2

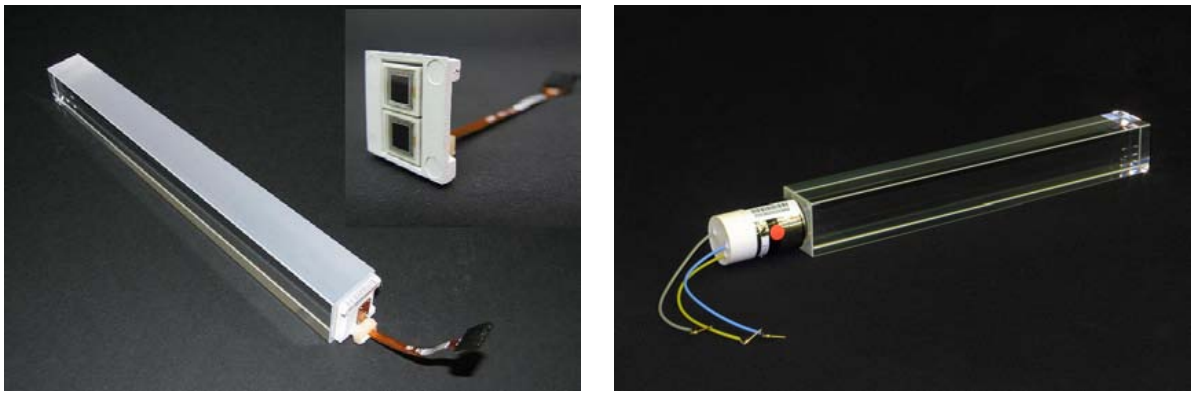

Figure 2.12. (Left) ECAL barrel crystal with attached APD. (Right) ECAL endcap crystal with attached VPT. [24]

at the rear face, and a length of 230 mm (corresponding to 25.8 radiation lengths). Groups of 2×5 barrel modules are referred to as sub-modules, encased within a 0.1-mm thick wall consisting of an aluminium layer facing the crystal, followed by two layers of glass fibre-epoxy resin. Submodules are grouped together into larger groups called modules (see an example in Figure 2.13) containing 400 or 500 crystals, whose shape depends on the position in η of the module. Four modules form a supermodule, and 18 supermodules form half a barrel.

On either side of the barrel, the ECAL endcaps cover the pseudorapidity range $1.479 < |\eta| < 3.0$. Endcap crystals have a truncated pyramidal shape similar to barrel crystals, with a front face cross-sectional area of about 28.62×28.62 mm^2 and a rear face cross-sectional area of 30×30 mm^2, and a length of 220 mm (24.7 radiation lengths). They are arranged in groups of 5×5 known as supercrystals, enclosed in a carbon-fibre alveolar wall. Each half of an endcap disk, called a Dee, is composed of 138 full supercrystals and 18 partial supercrystals; Figure 2.13 shows a Dee made up of supercrystals.

Sandwiched between the barrel and each endcap disk is a disk-shaped sampling calorimeter with a thickness of 20 cm; these two disks make up the preshower detector. Its main purposes are neutral pion identification and rejection in the fiducial region $1.653 < |\eta| < 2.6$, identify electrons against the background of minimum ionizing particles, and improve the position resolution of electrons and photons in the ECAL. The preshower is two layers

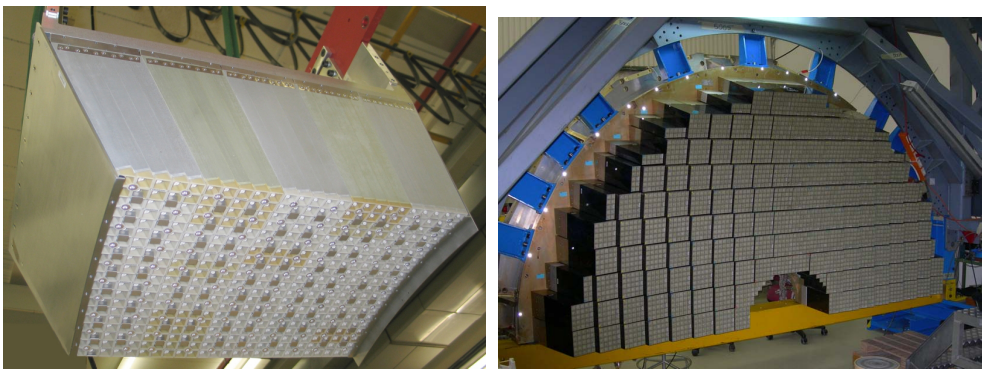

Figure 2.13. (Left) Module from the ECAL barrel. (Right) Supercrystals mounted on an ECAL endcap Dee. [24]

thick; each layer consists of a layer of lead radiators for initiating electromagnetic showers, followed by a layer of silicon strip sensors for detecting the showers.

To take advantage of total internal reflection for the collection of scintillation light, all the ECAL crystals (Figure 2.12) have been precisely polished during production. The crystals have a truncated pyramidal shape, which tends to cause nonuniform light collection along their length; thus, in the barrel, one crystal face is left unpolished to compensate for this effect. This technique is not used in the endcaps because the endcap crystal faces are nearly parallel to one another, resulting in more uniform light collection. The different magnetic field configuration and particle flux in the barrel and endcaps led to different choices of photodetectors for those two systems: avalanche photodiodes (APDs) for the barrel and vacuum phototriodes (VPTs) in the endcaps. At 18° C, approximately 4.5 photoelectrons per MeV of deposited energy are collected by both APDs and VPTs.

As the number of scintillation photons produced by the crystals and the amplification provided by the photodiodes both tend to decrease with increasing temperature, the temperature of the ECAL detector needs to be maintained at a steady operating temperature during its operation. This is done by a cooling system that uses water as a coolant, maintaining the ECAL at a stable temperature of 18° C with an uncertainty of ±0.05° C.

For electron or photon energies below 500 GeV, the ECAL energy resolution can be

estimated as follows:

$$(\frac{\sigma}{E})^2 = (\frac{S}{\sqrt{E}})^2 + (\frac{N}{E})^2 + C^2 \tag{2.4}$$

In this equation, S is the stochastic term, N is the noise term, and C is the constant term. The stochastic term S comes from stochastic fluctuations in electromagnetic shower containment, a photostatistics contribution of 2.1% coming from the uncertainty on the number of primary photoelectrons produced per MeV of deposited energy, and fluctuations in the energy deposited in the preshower absorber. The term N accounts for noise due to electronics, the signal digitization process, and energy deposited by pileup particles. Finally, the constant term C covers various sources of systematic error including the nonuniformity of light collection due to crystal shape, intercalibration errors, and leakage of energy from the back of the crystal.

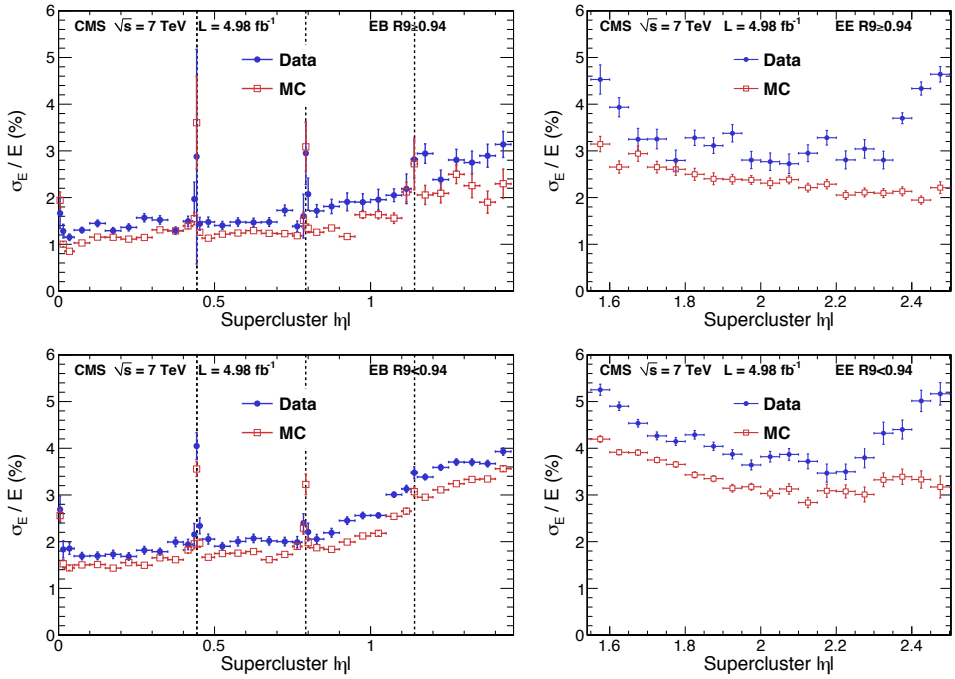

Figure 2.14. Relative energy resolution in the CMS ECAL detector as a function of supercluster $|\eta|$, measured from $Z \to e^+e^-$ events [28]. Blue points: data. Red circles: MC events. The term R9 is an ECAL cluster shape parameter used to differentiate between photons that convert in the tracker material (R9<0.94) and those that do not (R9≥0.94). (Top Left) Barrel, R9<0.94. (Top Right) Endcap, R9<0.94. (Bottom Left) Barrel, R9≥0.94. (Top Right) Endcap, R9≥0.94.

Studies done on 7 TeV proton-proton collision events in LHC data and Monte Carlo simulation have measured energy resolutions of less than 2% for $|\eta| < 0.8$ in the ECAL barrel, and between 2-5% everywhere else. These results are shown in Figure 2.14.

2.2.3 Hadronic calorimeter

Because proton beams are collided in the LHC and therefore hadronic processes feature in every event, the identification of jets is a crucial part of event selection and reconstruction at CMS. The main goal of the CMS hadronic calorimeter (HCAL) is to measure the energy and direction of hadronic jets. It is also useful for measuring the missing energy in each event, which is needed to identify final states involving not only neutrinos but also possibly exotic particles that do not interact significantly with the detector material, such as neutralinos or dark matter candidates.

The HCAL is a sampling calorimeter extending radially from 1.77 m to 2.95 m from the beam line. Its main goals are to measure the energy deposited by hadronic jets in each event, and to measure indirectly the missing energy in an event. Jet identification is an important part of event selection and reconstruction, and the determination of missing energy is useful not only for final states with neutrinos but also for in the search for final states with exotic particles that can pass through the detector without interaction, such as neutralinos and some dark matter candidates. The HCAL is split up into four subsystems, of which the innermost are the HCAL barrel and endcaps. To provide enough calorimeter material to absorb hadronic showers for jet and missing energy measurement while respecting the spatial constraints from ECAL within and the superconducting solenoid without, the barrel and endcaps are supplemented respectively by the outer calorimeter outside the solenoid and the forward calorimeters outside the muon endcaps. Altogether, these four subsystems provide coverage up to $|\eta| < 5.2$; their positions are illustrated in Figure 2.15.

The barrel system (HB) covers up to $|\eta| < 1.3$. It is subdivided longitudinally into two half-barrels, each of which consists of 18 identical wedges, making a total of 36 HB wedges. The wedges are subdivided into 4 azimuthal sectors staggered for optimum coverage, and they hold the absorber plates and scintillator material. The HB absorber is made up of an

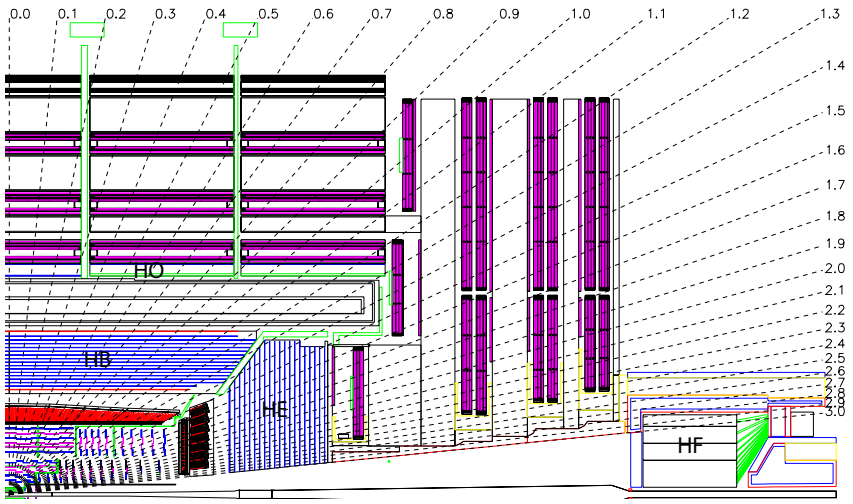

Figure 2.15. Cross-sectional view in the r-z plane of one quadrant of the CMS detector, showing the layout of the CMS hadronic calorimeter and the positions of the HB, HE, HO, and HF [24].

innermost 40-mm-thick steel front plate, eight 50.5-mm-thick brass plates, six 56.5-mm brass plates, and an outermost 75-mm-thick steel back plate. Passing hadrons interact with the nuclei of the absorber, producing hadronic showers of quarks and gluons. The energy in these showers is deposited in and measured by 17 layers of scintillator material that are interspersed between the absorber plates; layers in azimuthal sectors 1 and 4 fit into slots in the edges of a wedge, while layers in azimuthal sectors 2 and 3 fit into slots in the middle of a wedge. Besides the division into 18 sectors in ϕ, the scintillator material is also split into 16 sectors in η, resulting in a granularity of (0.087, 0.087) in (η, ϕ).

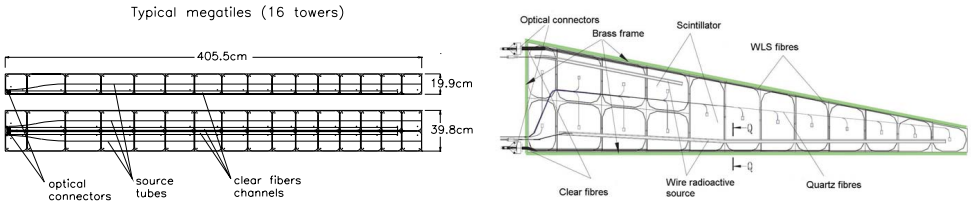

Figure 2.16. (Left) HB scintillator tray. (Right) HE scintillator tray. [24].

The smallest scintillator unit is called a tile; tiles in layer 0 are made of Bicron BC408 and primarily serve to sample hadronic showers that develop between the HB and the ECAL barrel, while tiles in layers 1-16 are made of of Kuraray SCSN81 plastic scintil-

lator. Tiles in a given azimuthal layer are grouped into units called trays for ease of assembly, access, and individual testing; each barrel layer contains a total of 108 trays. The scintillation light produced in each tile is collected by a Kuraray Y-11 wavelength-shifting (WLS) fibre. Each WLS fibre is spliced to a clear Kuraray double-clad fibre that serves all the tiles in the tray and delivers the collected light to an optical decoding unit, which arranges the clear fibres into readout towers and transmits their signals to a hybrid photodiode (HPD) for amplification and readout. Multipixel HPD's were chosen because of the wider range of wavelengths that they can detect and their low sensitivity to magnetic fields; those used in the HB, HE, and HO have a gain of roughly 2000.

The endcap system (HE) disks on either side of the HB are mounted onto the iron yoke of the muon endcap system and span the pseudorapidity range $1.3 < |\eta| < 3$. The disks are subdivided azimuthally into 36 identical wedges and 18 layers, with 79-mm-thick brass absorber plates. The HE contains a total of 20916 trapezoidal scintillator tiles, arranged into 1368 trays, with a granularity of (0.087, 0.087) in (η, ϕ) for $|\eta| < 1.6$ and (0.087, 0.087) for $|\eta| \geq 1.6$. The collection of scintillation light by WLS fibres and the transmission of the collected signal from tray to HPD for readout is similar to the design of the HB. Figure 2.16 shows schematic diagrams of both HB and HE trays.

The cylindrical outer calorimeter (HO) is located outside the solenoid, taking advantage of the solenoid coil as an additional sampling layer for late-starting hadronic showers. The thickness in hadronic interaction lengths λ_I of the HE + HB ensembles varies with η between 7-11 λ_I; when the HO is considered, the total thickness varies from 10-15 λ_I. The HO layers are mounted within the iron yoke that returns the magnetic field of the solenoid; they are the first sensitive layer in each of the 5 rings of the yoke. Each layer of the HO is divided into 12 sectors in ϕ, each separated by the dead space of 75-mm-thick steel beams that are part of the support structure of the return yoke and the muon system. The HO layers are 40 mm thick, of which 16 mm corresponds to the detector layer and the rest is occupied by aluminium support structures. The tiles in the HO (see Figure 2.17) are grouped into towers with a granularity of (0.087, 0.087) in (η, ϕ) and have a similar WLS scintillation light collection and HPD readout to the HE and HB

trays. Each HO tray spans 5° in ϕ and the entire range of a muon ring in η. Studies done to assess the effect of the HO on pion energy measurement by the HCAL have shown that the inclusion of energy measurements in the HO significantly recovers the effects of hadronic shower energy leakage and improves energy resolution (see Figure 2.18), which also directly improves the accuracy of missing transverse energy (MET) measurement in an event.

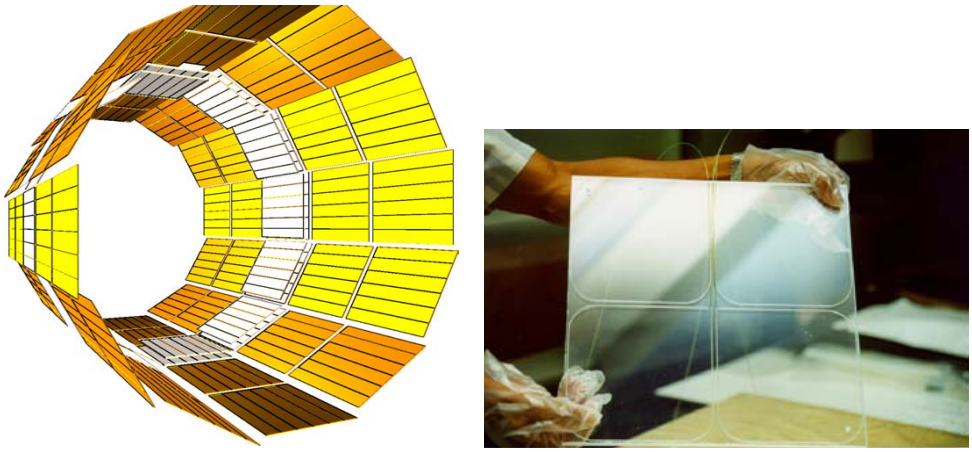

Figure 2.17. (Left) Positions of all the HO trays. (Right) Photo of a typical HO tile. [24]

The forward calorimeters (HF) are located far down the beamline, with the front face of each cylinder at 11.2 m from the interaction point of the CMS detector; a cross-sectional view is shown in Figure 2.19. The inner radius is 12.5 cm from the beam line, and the outer radius is at 130 cm. Unlike the HB, HE, and HO, the HF calorimeters use quartz fibres rather than plastic as the scintillation medium, so as to withstand the heavy particle fluxes (an average deposited energy of 760 GeV per proton-proton interaction, compared with 100 GeV for the rest of the HCAL) in the forward region that they occupy. Each HF cylinder is divided azimuthally into eighteen 20° wedges; the absorber is composed of 5-mm-thick steel plates with grooves into which the scintillation fibres are inserted. The absorber is subdivided into two parts in η; in one half, the fibres run through the full depth of the absorber, while in the other half, the fibres start at a depth of 22 cm from the front of the detector. With this design, the HF can distinguish between electromagnetic and hadronic showers, since the former tend to deposit most of their energy within the

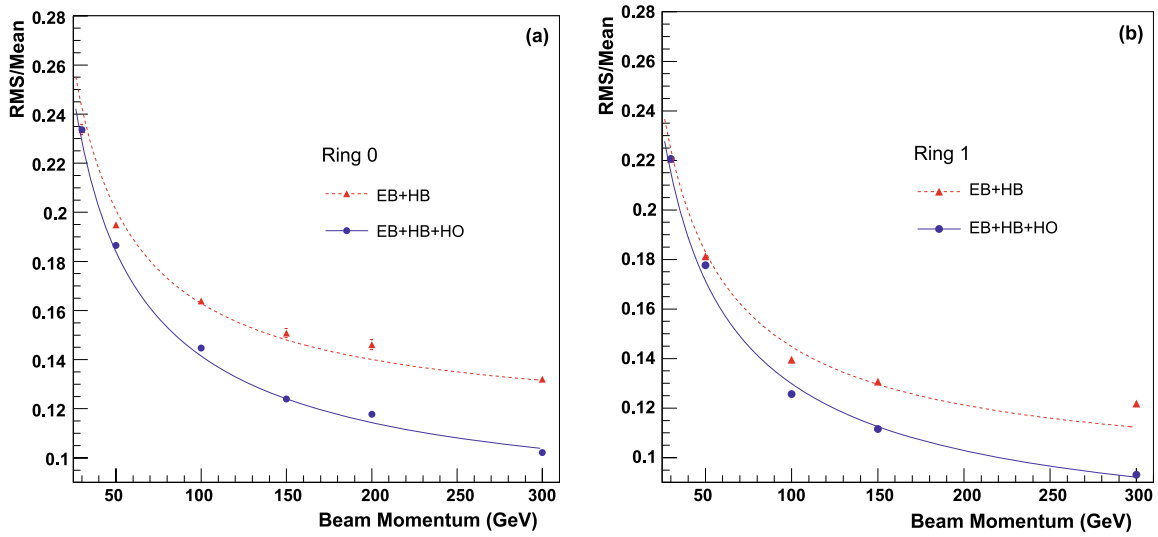

Figure 2.18. Energy resolution for pions as a function of pion test-beam energy, measured for EB and HB only (red curve) and for EB, HB, and HO (blue curve) [29]. The resolution visibly improves when HO energy measurements are included in the calculation. (Left) Resolution measured for a pion beam fired at $\eta = 0.22$. (Right) Resolution measured for a pion beam fired at $\eta = 0.56$.

first 22 cm and the latter deposit nearly equal energy in both segments.

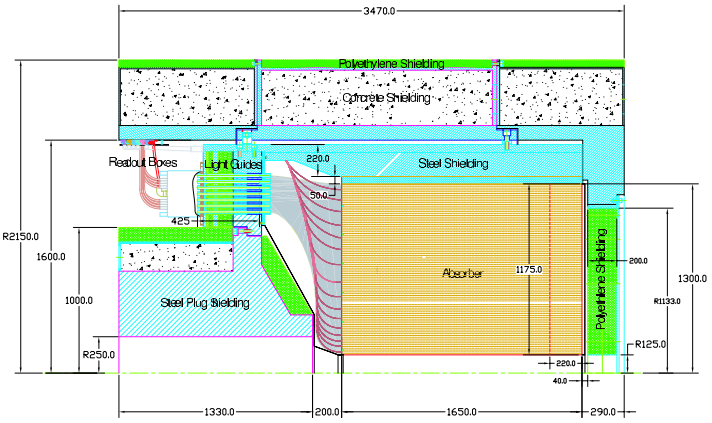

Figure 2.19. Cross-sectional view of the HF and surrounding support structures [24].

The quartz fibres are made of a fused-silica core with polymer hard-cladding. Charged particles passing through the crystals with energy above the Cherenkov threshold (roughly 190 keV for electrons) will generate Cherenkov radiation, which is captured by the quartz fibres. The fraction of light f_{trap} captured by a crystal is equal to $\mathrm{NA}/2\mathrm{n}^2_{core}$, where NA

is the numerical aperture and n_{core} is the refractive index of the quartz core. The light that hits the core-cladding interface at an angle larger than the critical angle of 71° can contribute to the calorimeter signal, while the rest is lost. The fibres run in bundles parallel to the beam line, forming towers with a granularity of (0.175, 0.175) in (η, ϕ). Bundles of fibres emerging from the back of the calorimeter are routed to air-core light guides protected by a shielding matrix of steel, lead, and polyethylene. The interior of a light guide is coated with highly reflective metal and guides the received light from the quartz fibres to a photomultiplier for readout; one readout box holds 24 photomultipliers and services half of a wedge.

In addition to contributing to hadronic shower energy measurement, the HF is also used to monitor the luminosity of the LHC. The average number of empty towers in the HF is related to the mean number of interactions per bunch crossing, and there is also a linear relationship between luminosity and the average transverse energy per tower.

2.2.4 Magnet

The superconducting solenoid magnet of CMS encloses the tracker detector, ECAL, and part of the HCAL inside a cylindrical bore 6.3 m in diameter and 12.5 m long, weighing 220 tonnes. The solenoid is composed of a 4-layer winding of coils made of NbTi, which is mechanically reinforced by being mixed with an aluminium alloy, an innovative method that makes the coils serve both as conductors and as their own self-stabilizing structural support.

At full current, the solenoid carries 19.14 kA, producing a nearly uniform 3.8 T magnetic field through the volume enclosed by the bore, containing 2.6 GJ of stored electromagnetic energy at full current. The magnetic flux is returned through an iron yoke, consisting of a system of barrel wheels and endcap disks arranged outside the volume of the solenoid in a cylindrical pattern and interleaved with the muon tracking system; the yoke is segmented longitudinally into 5 rings, each 2.536 m long in the z direction, with their centers positioned at z = -5.342 m, -2.686 m, 0 m, +2.686 m, and +5.342 m (rings number -2, -1, 0, 1, and 2 respectively).

The uniformity of the magnetic flux density inside the detector volume is illustrated

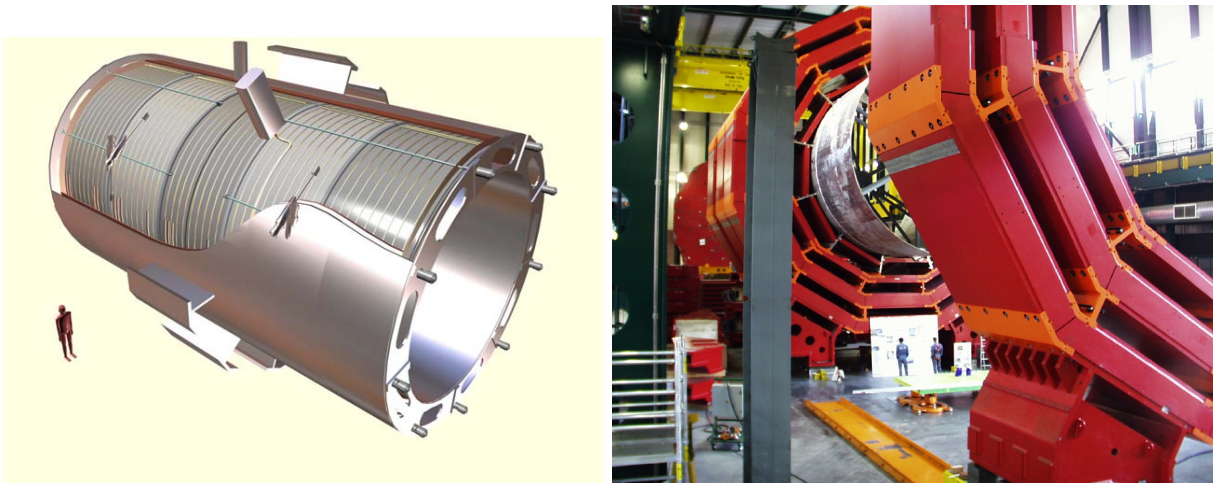

Figure 2.20. (Left) Cold mass of the solenoid magnet, showing the five longitudinal segments. (Right) Image of the iron yoke of the solenoid magnet, also showing the central barrel that supports the vacuum chamber of the superconducting coil. [24]

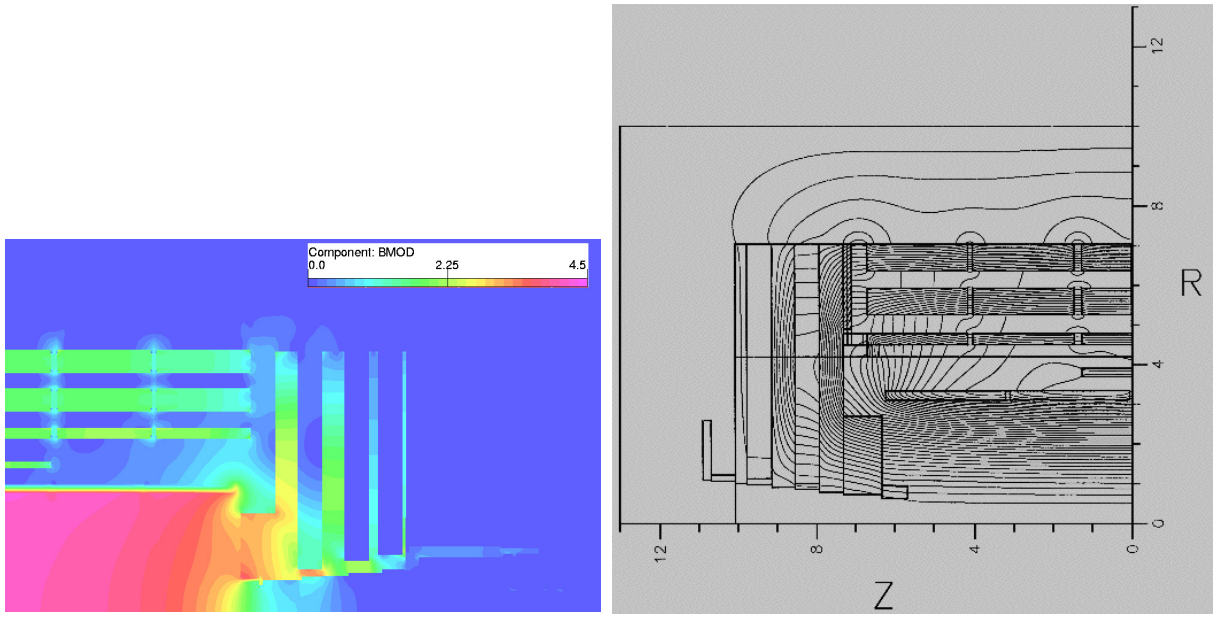

Figure 2.21. (Left) Longitudinal cross-sectional view of one quarter of the CMS detector, showing a color plot of magnetic flux density (Tesla) in simulation [30]. (Right) Magnetic flux lines in the return yoke and muon system [31].

in Figure 2.21, which was obtained via simulation [30]. Within the solenoid volume, the field is nearly uniform at 3.8 T. The field lines becomes nonuniform at the end faces of the cylindrical solenoid volume, and the return flux is mostly confined to the iron yoke, though there is also a residual nonzero magnetic flux in the muon system volume.

2.2.5 Muon system

Muons tend to undergo far lower radiative losses in detector material than electrons with similar kinetic energy and are consequently cleaner to reconstruct. Because they are expected to be the only particles (not counting neutrinos) that survive past the tracker and calorimeter layers, the muon detector system is the outermost layer of the CMS detector, located outside the solenoid. Muons are part of many of the physics signatures sought at the LHC, such as the decay $H \rightarrow ZZ* \rightarrow 4\mu$ or $H \rightarrow WW/ZZ$, and so the experiment benefits from having a muon system with wide angular coverage and efficienct reconstruction capabilities.

Like the rest of the CMS subsystems, the muon system is comprised of a barrel region and two endcap regions. The muon detector chambers are mounted on the flux return yoke system. The barrel is made of four layers called stations, and its longitudinal segmentation is dictated by the segmentation of the yoke into 5 rings, while the azimuthal segmentation comes from the iron ribs of the yoke that serve as support structures for the solenoid coil cryostat and the HCAL outer calorimeter [32].

The goals of the muon system are to provide reliable reconstruction of muon momentum and charge over the entire kinematic range of the LHC, as well as to provide the capability to trigger on events in the detector (the concept of triggering is explained in Section 2.3). To these ends, the muon system employs three different kinds of gas ionization detector technology: drift tubes (DTs), cathode strip chambers (CSCs), and resistive plate chambers (RPCs).

The DT system is located only in the barrel of the muon detector. The choice to use DTs was based on the low expected particle flux and low magnetic field in the barrel. The smallest sensitive element is the drift cell, with a transverse dimension of 21 mm and an anode wire of roughly 2.4 m in length. Figure 2.22 shows a cross-section of a DT cell.

The cells contain an 85%:15% mixture of Argon and CO_2 gas; when a charged particle passes through, the ionization of the gas results in a charge carrier cascade that is picked up by the anode and cathode wires and amplified by the readout system.

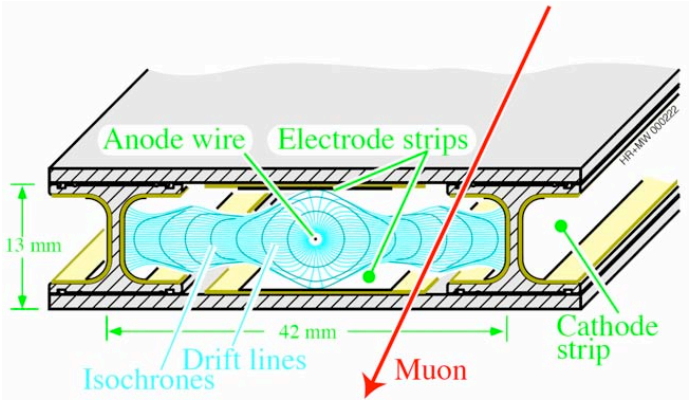

Figure 2.22. Cross-sectional view of a DT cell, with drift lines indicated [24].

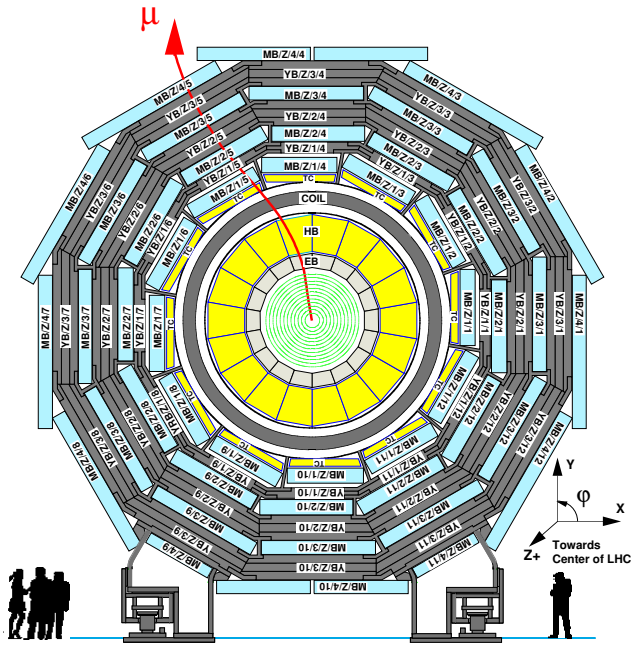

Figure 2.23. Layout of the DT chambers (light blue) in the iron yoke [24].

Drift cells are grouped into four layers staggered by half a cell; these groups of four layers are independent units known as superlayers (SLs). SLs are further grouped into groups of either 2 or 3 to form a drift chamber. Drift chambers with only SL's have

wire anodes running only in the r-ϕ direction and provide only a measurement of the ϕ coordinate of a hit; these chambers are only located in the fourth muon station. In the inner three muon stations, the chambers are composed of 3 SL's, of which the outermost two have wire anodes running in the r-ϕ direction and the other has the wire running in the z direction, thus yielding both r-ϕ and z coordinates of a hit. The inner three muon stations have 60 drift tube chambers each, while the fourth station has 70; Figure 2.23 depicts the layout of the DT chambers mounted in the iron yoke.

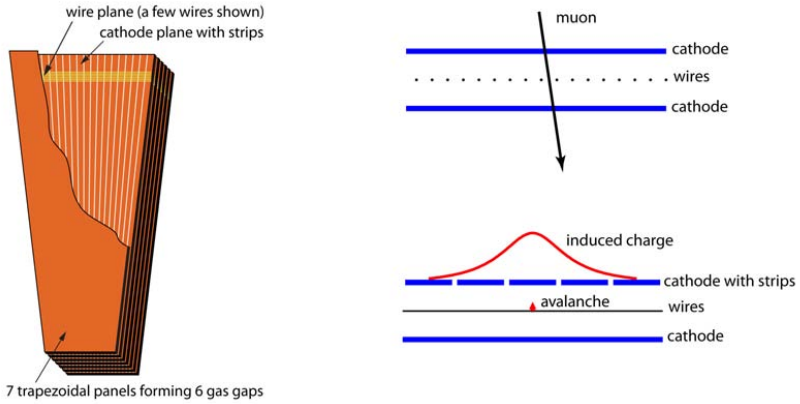

Figure 2.24. (Left) Layout of a CSC, with the top panel peeled back to show anode wires and cathode strips. (Right) Diagram of charge induction in a CSC gap by a passing high-energy particle. By the same logic as for charge sharing in the pixel system, the interpolation of charges induced on cathode strips by an avalanche of positive charge carriers near a wire leads to better resolution of the avalanche position along the wire direction [24]

CSCs were chosen for the endcap region of the muon detector where the particle flux is higher. The cathode strip chambers are trapezoidal multi-wire proportional counters, with 6 anode wires running along the azimuthal direction and 7 cathode strips running perpendicular to the wires in the radial direction. Thus, the passage of a particle through a CSC will yield a measurement in both r and ϕ coordinates (from the charges induced in the anode and cathode strips respectively), with spatial resolution within 80 μm. The shape and structure of a CSC are shown in Figure 2.24, as well as a depiction of charge induction on cathode strips in a CSC.

There are 468 cathode strip chambers in total, providing coverage over the pseudora-

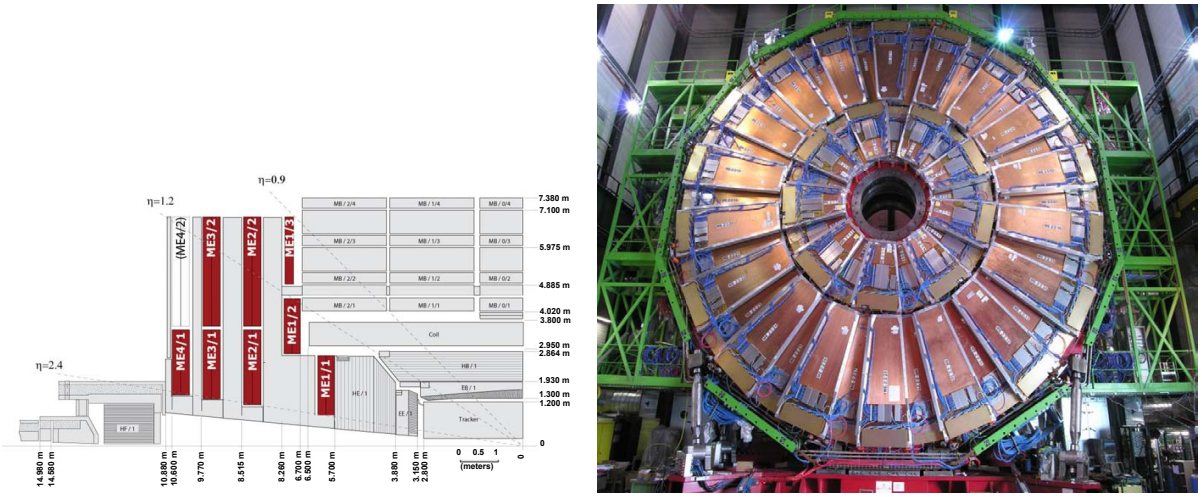

Figure 2.25. (Left) View of the CSC chambers (dark red) in the CMS detector. (Right) Photo of one of the CSC stations. The chambers in each ring (inner and outer) overlap to provide continous azimuthal coverage. [24]

pidity range $0.9 < |\eta| < 2.4$. The gas used in the chambers is a mixture of argon, CO_2, and CF_4 in a ratio of 40%:50%:10%, with the CF_4 included to prevent polymerization on the wires. The CSC system has two purposes: to yield a precise measurement of muon tracks, and to provide information for triggering. Local charged track (LCT) boards sample all the anode and cathode strip readouts of all chambers in synchronization with every LHC bunch crossing. Algorithms seek patterns of hits in the output read out by the anode and cathode LCT boards; if patterns are found that are consistent with muon tracks, they serve as trigger primitives and are used for building muon track candidates.

Finally, there is the RPC system, which consists of gaseous parallel-plate detectors arranged into six barrel layers and three layers per endcap. The gas used in the chambers is a mixture of $C_2H_2F_4$, iC_4H_{10}, and SF_6 in a ratio of 96.2%:3.5%:0.3%, with water vapour added to achieve 45% relative humidity in order to avoid changes in the resistivity of the bakelite plates. The six RPC barrel layers (see schematic layout in Figure 2.25) are embedded in the iron yoke, with two in each of the first and second muon stations and one in each of the third and fourth stations. With the current endcaps, the RPC system extends out to $|\eta| = 1.6$, with future plans to install an additional fourth endcap layer that would provide coverage out to $|\eta| = 2.1$.

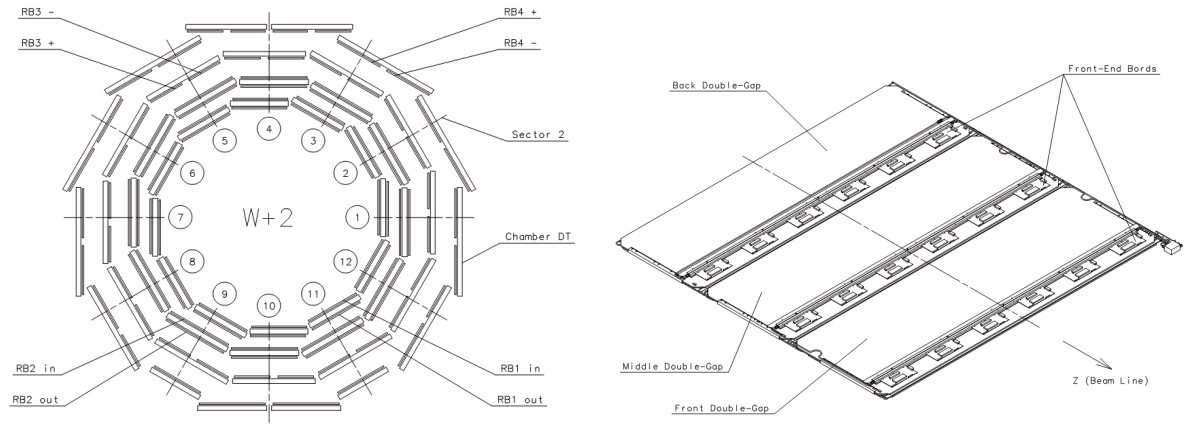

Figure 2.26. (Left) Layout of the RPC barrel (dark grey) in the iron yoke. (Right) Barrel RPC module with 3 double-gaps. [24]

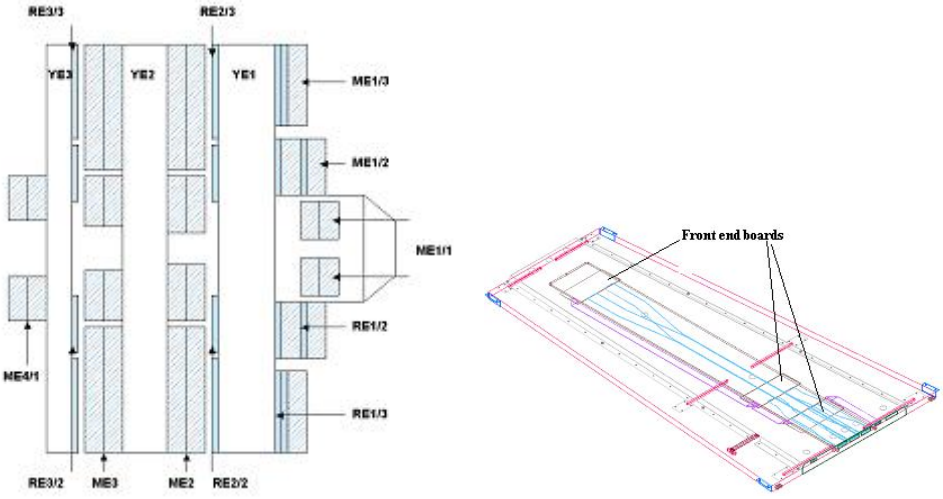

Figure 2.27. (Left) Layout of the RPC endcap. (Right) Endcap RPC chamber. [24]

An RPC chamber in the barrel has a rectangular shape and consists of 2 or 3 double-gap modules with up to 96 strips per double-gap, in which the strips run parallel to the beam direction; there are a total of 480 barrel chambers. In the endcap, chambers also have a double-gap structure but are trapezoidal in shape; each chamber has 32 strips running radially and covers 20° in ϕ in the innermost ring and 10° in ϕ in the outermost ring. Figure 2.26 shows the arrangement of the barrel RPCs in the iron yoke, as well as an example diagram of a 3-gap barrel RPC module. Figure 2.27 shows an analogous layout of the RPC endcap system, and a diagram of a typical RPC chamber.

Front-end electronics boards are always located at the strip end of both barrel and endcap chambers, to minimize the signal arrival time with respect to the interaction point. The RPCs provide timing resolution on the order of less than the 25 ns between bunch crossings, allowing the precise assignment of muon tracks to bunch crossings as well as transverse momentum measurement.

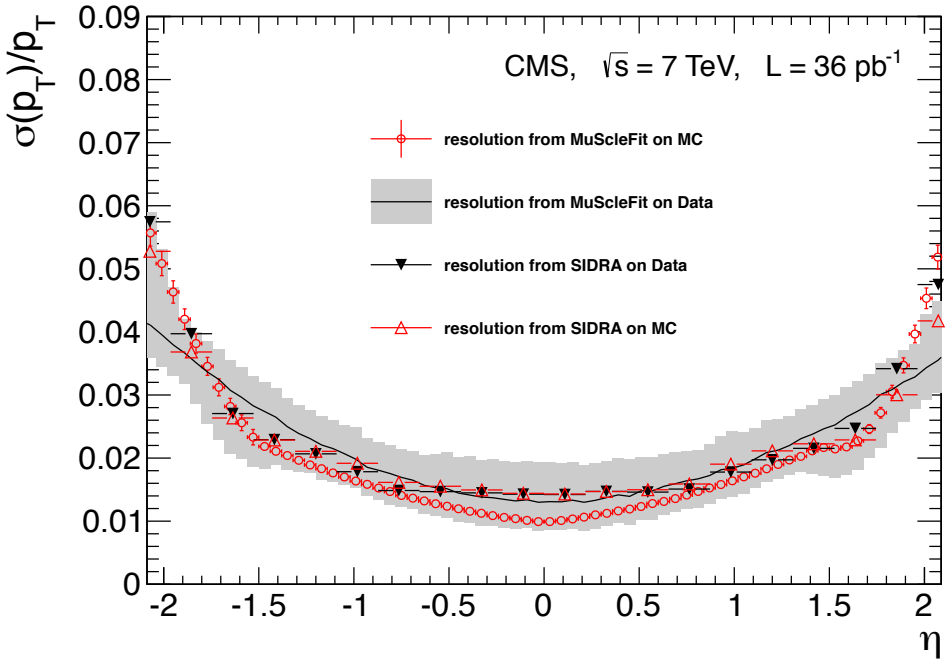

Figure 2.28. Relative p_T resolution for muons from Z decays, measured in 7 TeV LHC data (black curves) and in MC events (red curves), using two different algorithms, MuScleFit and SIDRA [33].

All in all, the muon system achieves a high relative p_T resolution, measuring 1.3-2.0%

in the barrel and less than 6% in the endcaps [33]. The results of a study done in 7-TeV data and MC events are displayed in Figure 2.28; in this plot of relative muon p_T resolution versus η, the relative p_T resolution is shown not to exceed 6%.

2.3 Triggering and data acquisition

Every collision of proton bunches at the CMS interaction points produces a large number of particles traversing the volume of the CMS detector. With protons colliding in the LHC tunnel at a rate of 40 MHz (1 bunch crossing every 25 ns) and about 20 proton-proton collisions per bunch crossing at the design luminosity of 10^{34} cm^{-2}s^{-1}, a huge amount of data is generated at a high frequency. In order to save information from only the most important events and thus respect the limitations of the electronics and computer systems that process and store this data, a method for drastically reducing the rate of data acquisition is required. The trigger system of the CMS experiment performs a rapid assessment of each event recorded, storing the data from the event only if it passes certain criteria that mark it as an event of interest. Triggering is performed in two main steps, the Level 1 (L1) Trigger and the High Level Trigger (HLT). Together, these two steps provide a combined rate reduction factor of at least 10^6.

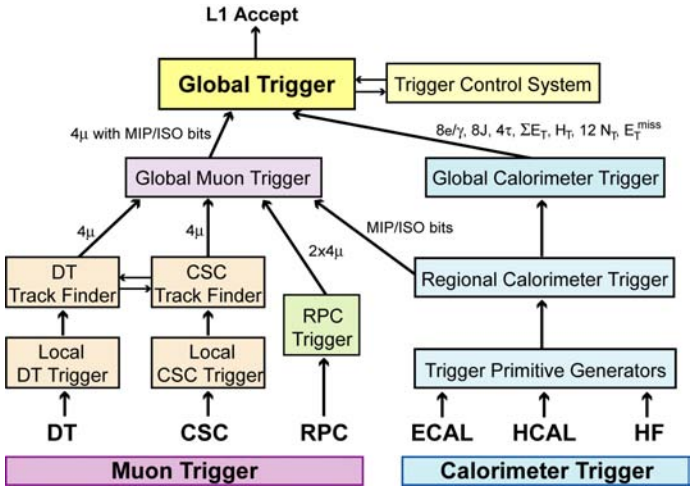

Figure 2.29. Flowchart of the steps and components in the L1 Trigger sequence [24].

The L1 Trigger consists of a system of programmable electronics, located partly in the

CMS detector and partly in the underground control room near the experimental cavern. Energy deposits in calorimeter trigger towers serve as input to the local calorimeter trigger component of the L1, while track segments and hit patterns from the muon chamber serve as input to the local muon trigger component. Regional calorimeter triggers and DT and CSC track finders take the information from the local triggers and search for patterns in order to rank and sort trigger objects based on their energy or momentum and quality. Next, the global calorimeter trigger and global muon trigger take the output of their respective regional triggers and pick out the highest-ranking calorimeter and muon objects across the whole experiment. This information is finally passed to the global trigger at the end of the L1 decision chain, which uses various algorithm calculations to assess whether to reject the event completely or to accept it to be passed on to the HLT; if the event passes the L1 criteria and all the subdetectors and data acquisition (DAQ) systems are ready, then the event data are passed to the HLT.

The above steps are illustrated by the flowchart in Figure 2.29. The L1 Trigger analyzes every bunch crossing, with a latency period of 3.2 μs between one bunch crossing and the distribution of the trigger decision to the front-end electronics. Its design output rate limit without data recording is 100 kHz; in practice, the final rate is on the order of 30 kHz.

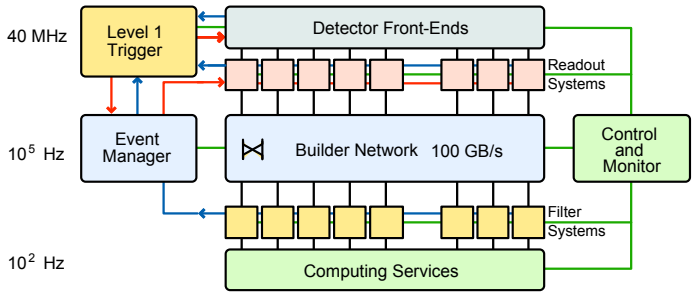

Figure 2.30. Flowchart of the CMS DAQ architecture [24].

When an event passes the L1 Trigger, the data is passed to the CMS DAQ system, which is schematically depicted in Figure 2.30. The subdetector front-end systems store event data in buffers until the L1 Trigger decision allows them to release this data to the

front-end drivers (FEDs) of the DAQ system. Event fragments from the FEDs are merged by a set of processors called the Event Builder to produce a data structure containing the complete event information, which is then passed to the Event Filter, a computer farm of about one thousand PC's. The Event Filter submits the event data to the HLT, a flexible software system implemented in the computers of the Event Filter, and also performs data quality monitoring (DQM) to assess the goodness of the data.

The HLT [34] runs reconstruction and filtering algorithms on the event data. The reconstruction process builds physics objects (C++ data structures) from the raw data using a streamlined version of the CMS offline reconstruction software (which is described later in Section 3.1. The filtering process selects whether or not to keep an event, based on the criteria that classify it as having interesting physics content; these criteria define what is known as the HLT path, and they vary widely depending on the physics object or combinations of objects being selected. Many different HLT paths are used for the numerous physics searches carried out by the CMS experiment. One example of an HLT path is a single-muon trigger that stores events if they contain at least one muon with p_T $>$ 24 GeV and whose isolation – measured by summing the pixel tracks and calorimeter energy deposits in a cone of fixed radius about the direction of the muon's reconstructed four-momentum – is lower than some maximum cutoff. Other examples of HLT paths include selections on the p_T, isolation, and multiplicity of electrons and muons, the energy and multiplicity of jets, and the amount of missing transverse energy in an event.

From the input rate of around 100 kHz from the L1 Trigger, the HLT ultimately reduces the rate of event processing to $\mathcal{O}(100)$ Hz. Datasets of raw events passing the HLT are written to permanent storage and undergo further offline processing before they are ready for use in physics studies or for calibration and data monitoring.

2.4 Luminosity measurement

The CMS experiment has two different systems for measuring the luminosity delivered by the LHC. One relies on the forward calorimeter HF, while the other uses the pixel detector [35]. These two methods have different strengths and can be used as cross-checks

for one another.

The HF zero counting method calculates the average luminosity per bunch crossing to 1% statistical accuracy by estimating the mean number of interactions per bunch crossing from the average number of empty HF towers. Because of its high accuracy and the fact that HF can operate even during unstable beam conditions, the official online luminosity measurement (i.e., measured while the detector is running) during Run I has been taken from the HF zero counting method. Figure 2.31 shows a plot of the total integrated luminosity per day delivered by the LHC and recorded by CMS during the proton-proton collisions at center-of-mass energy 8 TeV during the year 2012.

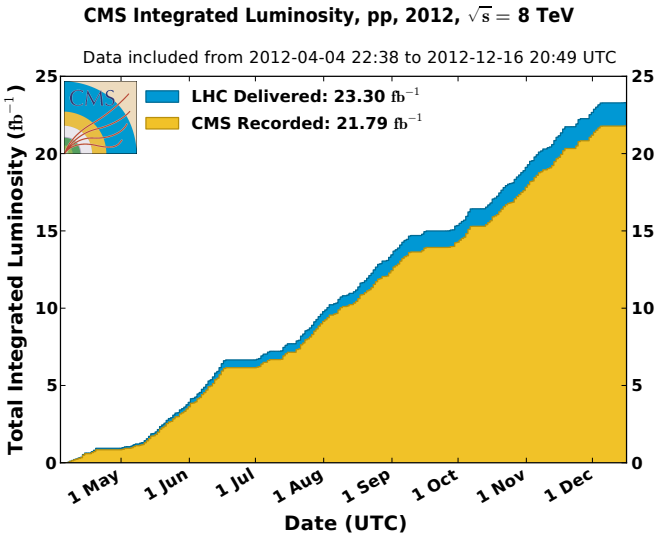

Figure 2.31. Total integrated luminosity per day delivered by the LHC (blue) and recorded by CMS (yellow) for proton-proton collisions at 8 TeV in 2012 [36].

Offline measurements of luminosity are performed using information from the pixel system, which can only operate during stable beam conditions. The pixel-cluster counting method uses the average number of pixel clusters in a bunch crossing to calculate the instantaneous luminosity per bunch L according to this formula:

$$L = \frac{\nu < n >}{\sigma_{vis}} \tag{2.5}$$

where $\nu = 11246$ Hz is the beam frequency, $< n >$ is the average number of pixel clusters per event, and σ_{vis} is the visible cross-section for inelastic collisions. σ_{vis} is calibrated via

the Van der Meer scan technique, in which the two LHC beams are scanned across each other along the horizontal and vertical planes to measure the shape and overlap of the beams. Unlike the HF-based luminosity measurement, which is susceptible to calibration drifts due to various factors, such as changes in the PMT gains and the HF detector's nonlinear response with respect to pileup, the pixel-based luminosity measurement is stable over time. Figure 2.32 illustrates the stability of the luminosity measurement as a function of pileup.

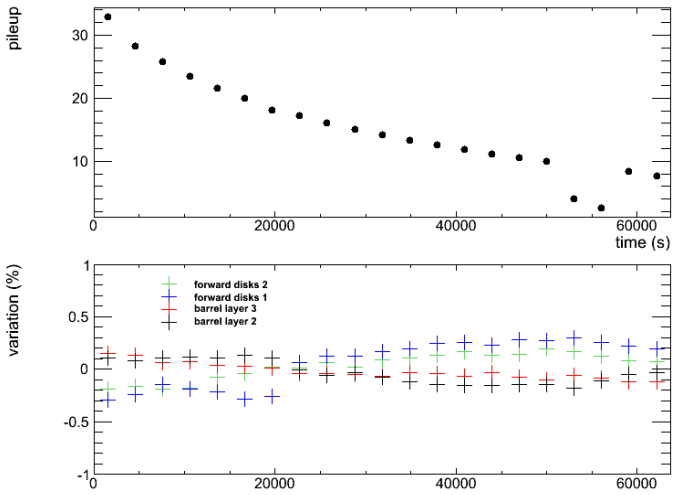

Figure 2.32. (Top) Average rate of pileup events as a function of time during an LHC proton beam fill in November 2012. The dip towards the end is due to an end-of-fill Van der Meer scan. (Bottom) Percent variation in luminosity measurement as a function of time (or pileup), showing stable behaviour as the variation is bounded between $\pm 0.5\%$ over the fill period for all layers and disks of the pixel system. [36].

Chapter 3

Event reconstruction and simulation

Once all the detector hits and deposits produced in a bunch collision have been recorded and stored, a sophisticated series of pattern-recognition algorithms takes these disparate data and assembles them into a picture of what actually occurred in that collision. Particle energies and trajectories are reconstructed, and particle types are identified.

Particle reconstruction algorithms are used not only for experimental data, but also for datasets of simulated physics events. Event simulation is an important component of particle physics research, as it provides a way to test data against theoretical predictions.

In this chapter, an overview is given of the standard CMS particle reconstruction procedure, and of how event reconstruction is typically performed.

3.1 Particle reconstruction

Before reaching a relatively stable final state, the immediate products of the proton-proton collisions in the LHC tunnel undergo various interactions such as radiation, decays, and/or hadronization. The resulting particles travel through and interact with the material of the CMS detector, which records their passage and reconstructs their paths and energies, which are then used to identify the particles and deduce the interactions that they underwent. The set of algorithms used predominantly for particle reconstruction and identification in the CMS experiment is called Particle Flow [37], often abbreviated as PF.

The basic building blocks passed to the PF algorithm are tracks and clusters. Charged

particle tracks are reconstructed from hits in the silicon tracker by an iterative tracking algorithm; a similar process is done to reconstruct muon tracks in the muon detector. Track resolution is crucial for accurately determining the trajectory as each track, as inaccuracies can lead to large discrepancies in reconstructed energies. Clustering algorithms search for patterns in energy deposits in the ECAL and HCAL, to reconstruct the energies and trajectories of photons, neutral hadrons, and charged hadrons.

Once tracks and clusters are reconstructed, linking algorithms make tentative associations between these elements, based on criteria such as their closeness in distance and in energy. Track paths are extrapolated from the tracker into the ECAL and HCAL and matched to clusters if their reconstructed momenta and positions are compatible; clusters in the ECAL and HCAL are linked if the extrapolated trajectory of the ECAL cluster lies within the HCAL cluster envelope; tracks from the silicon and muon trackers are linked by performing a global χ^2 fit between the two types of tracks.

Particle Flow combines all this track and cluster information from all subdetectors to build reconstructed particle objects. The abbreviation "reco" will often be used in this dissertation to refer to reconstructed objects. The first objects to be constructed are stable individual particles: electrons, muons, photons, and charged and neutral hadrons.

If the momentum from the combination of a linked muon track and silicon tracker track is compatible with the silicon tracker track momentum alone, the linked object is a PF muon candidate. Charged particle tracks linked to ECAL clusters are used to seed electron and charged hadron candidates, while ECAL and HCAL clusters that cannot be matched to any track are used to seed photon and neutral hadron candidates respectively; the effect of the HCAL granularity's coarseness on reconstructed hadron energy resolution is improved by combining information from the tracker system in the case of charged hadrons, and from the ECAL system (whose granularity is 25 times finer than that of the HCAL) in the case of neutral hadrons. As an example, the reconstruction of the individual hadrons that make up a hadronic jet is illustrated in Figure 3.1.

The collections of individual reconstructed stable particles are then used to reconstruct more complicated objects. Jet clustering algorithms group electron, muon, photon, and

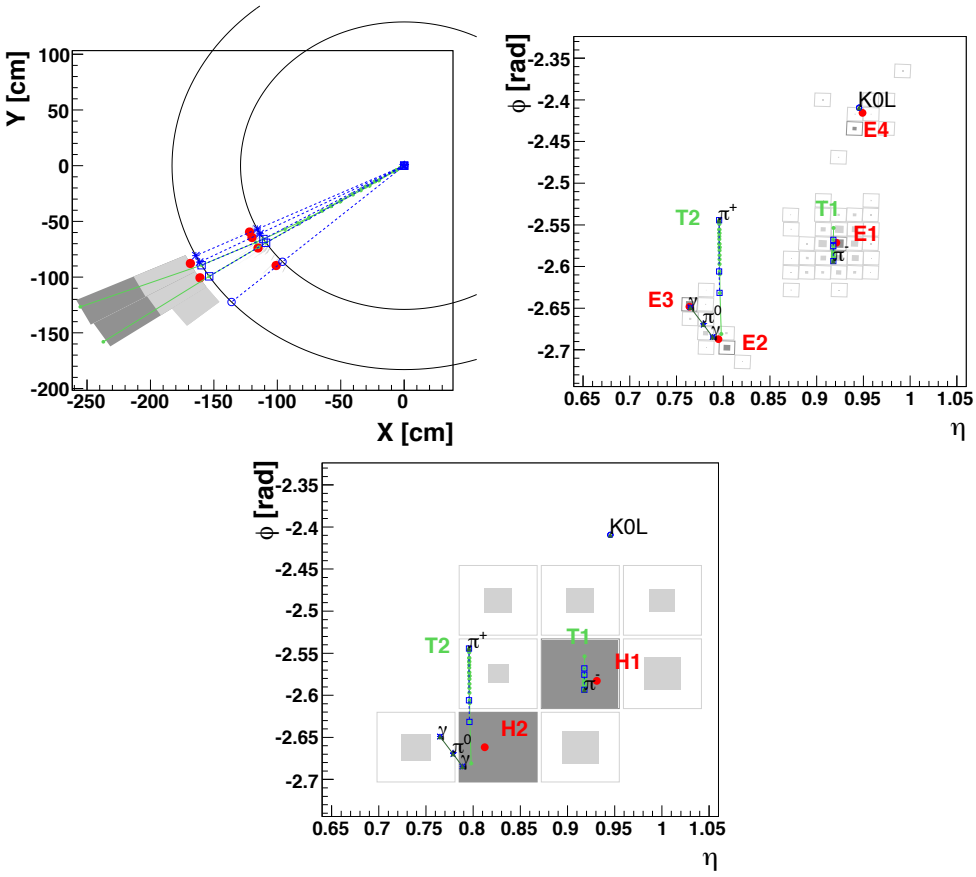

Figure 3.1. Event display of a hadronic jet consisting of a K_L^0, π^+, π^-, and π^0, reconstructed via the Particle Flow algorithm from tracks and calorimeter deposits. The π^0 is detected via its decay to a pair of photons in the ECAL. Figures copied from [37]. (Top Left) View in the xy plane, showing the tracks (green arcs), the ECAL and HCAL (represented by the two concentric circles), and calorimeter towers (dark and light grey for HCAL and ECAL respectively). The positions of impact of each of the four particles on the ECAL and HCAL are represented by open blue markers. (Top Right) View in the $\phi\eta$ plane for the ECAL, showing clusters reconstructed from ECAL deposits. The K_L^0, π^-, and $\pi^0 \to \gamma\gamma$ leave four well-separated clusters E1-E4 in the ECAL, while the π^+ passes through without leaving an energy deposit. (Bottom) View in the $\phi\eta$ plane for the HCAL. The two charged pions are reconstructed as charged tracks T1 and T2 (green lines), pointing to HCAL clusters H1 and H2; Particle Flow associates these tracks with the respective HCAL cluster that they point to. Cluster positions are indicated by solid red dots in all three views.

charged or neutral hadron candidates together into jets. iterative cone techniques [38] use the hardest particle or calorimeter tower in the event as a seed and builds a jet from the PF candidates in a cone around it, removes all of the jet candidates from consideration, and then moves on to find the next jet seed from the remaining candidates in the event, proceeding thus until no seeds are left. Hadronic tau decays are reconstructed from PF jets; currently the approved algorithm used by CMS is the Hadrons Plus Strips (HPS) algorithm [39], which reconstructs the decay mode of a tau based on the number of charged hadrons, neutral hadrons, and photons among the tau decay candidates. Finally, since the initial-state colliding protons have zero transverse energy, the missing transverse energy in the event is reconstructed by calculating the vector sum of the transverse energies of all reconstructed particles and taking the negative, based on the conservation of momentum.

3.2 Event simulation with Monte Carlo

Simulating particle collisions is an essential part of high-energy collider experiments. For instance, in order to determine whether actual data indicate the detection of new physics, one must know what data to expect if nothing new is occurring, so as to be able to compare collected data with predictions from the old theory; a reliable program for simulating physics processes based on known theory can provide a convenient means of obtaining a prediction for expected backgrounds. Other examples of uses for physics event simulation include calibrating the detector and testing the efficiency of its hardware and software.

The evolution of a simulated event in a collider can be broken down into two main steps: 1) modeling the particle collision and subsequent particle production, and 2) modeling the interaction of the final-state particles with the various parts of the detector, decays in flight, and the detector response. The principles behind these two aspects of event simulation will now be discussed. The physics event generation package PYTHIA [22] and the detector simulation package GEANT4 [40], both used in the Compact Muon Solenoid (CMS) experiment, will be described as concrete examples of event simulation software. Many other packages exist that have somewhat different mechanisms, but the general principles are essentially the same.

For any particle interaction, there exists a spectrum of final states, each with its own particular amplitude for occurring. The phase space describes all the possible final states that the system can achieve; the relative probabilities for these final states are a function of the momenta and relative trajectories of the particles. The evolution of the system involves an element of randomness due to quantum mechanics; the most common technique for simulating this is is the Monte Carlo (MC) method, which uses a random number generator to sample the phase space for each simulated physics process and thus evolve of the event in a probabilistic fashion.

PYTHIA generates physics events in a series of steps. The first step is the hard scattering of the colliding initial-state particles (protons, in LHC event simulations). To generate events with the relative frequencies with which one would expect them to occur in actual colliders, the various possible reaction channels need to be weighted according to their cross-sections, which PYTHIA calculates. The initial-state particles are characterized by parton distribution functions (PDFs); even leptons, which are not partons, are described with a PDF that reflects the likelihood of photon emission by the lepton before it enters the initial hard process with a fraction x of its original momentum.

Photon or gluon radiation can occur before and after the hard scattering process. PYTHIA is optimized to model $2 \to 1$ and $2 \to 2$ processes (where the numbers indicate the number of particles in the initial and final states), for which it can compute the cross-sections. However, a challenge arises in simulating radiation processes, which begin with one particle and end with two or more. Gluon radiation is especially problematic because it is governed by QCD, and for soft radiative processes – where the radiated gluons are roughly collimated with the final-state parton – the momentum transfer values involved are so low that strong coupling constant is large enough for processes higher than tree level to be significant, and thus amplitudes for these processes cannot be calculated perturbatively. This makes the computation of amplitudes for radiation processes extremely complicated, and PYTHIA does not perform such calculations. Instead, it uses the parton showering method to estimate the higher-order matrix elements for initial and final state radiation. Parton showering simulates the branching of one randomly-chosen "shower ini-

tiation" parton (not necessarily the parton involved in the hard process) into a number of other partons and combining the results. PYTHIA estimates the branching probabilities for quarks, leptons, and gluons with a simplified kinematic model that is a function of the 4-momentum fraction z for the branching; $1 \rightarrow 2$ decays are simulated with this model until the final-state particles in the shower reach a certain cutoff energy, below which no further radiation is simulated. The formation of jets from the beam remnants (the partons not involved directly in the hard scattering) is modeled similarly. Other event generation packages used by CMS, such as MADGRAPH [41], BlackHat, SHERPA [42], and POWHEG [43], do calculate some of the matrix elements for higher-order processes. However, due to the inherent non-perturbativeness of QCD, these calculations necessarily involve their own approximations and inaccuracies.

After the hard scattering and final-state showering, the resultant quarks and gluons must hadronize; in PYTHIA, this process is referred to as fragmentation. Often, the hadrons produced are unstable and will radiate and decay further; the series of fragmentation and decays that occur before the final state is reached are collectively termed hadronization. Again because of the non-perturbative QCD diagrams involved in the matrix elements, PYTHIA relies instead on simplified models of fragmentation based on the Lund string scheme to approximate the process.

The next step in event generation is to model the response of the detector to the particles produced in simulated collisions. At CMS, GEANT4 is used to model the trajectories of final-state particles and their interactions with the parts of the detector in their path, and the way in which the detector elements register and record the particles that pass through them.

The geometry of the detector – the material composition, positions, dimensions, etc. of its components (both sensitive elements and structural supports) – must be specified in the simulation program. Monte Carlo methods are used to model the interaction of particles with the detector components. The rate of energy loss per unit distance is determined by the medium's chemical composition and by the type of interaction involved in the energy deposition, which depends on on the particle's energy; for different particle energies,

processes such as ionization (governed by the Bethe-Bloch equation) and bremsstrahlung occur with different relative probabilities, characterized by a mean free path (radiation length for bremsstrahlung, and hadronic interaction length for strong interactions) [44].

When final-state particles emerge from a simulated collision, GEANT4 simulates each particle's trajectory step-by-step through the detector volume, evolving it under the influence of electromagnetic fields and also via its interactions with the materials that it passes through. For each step that a particle makes, all possible types of interactions with the detector material are considered, and cross-sections are computed for all of them; the probability and spacetime step-lengths for each interaction are then calculated, and the smallest step-length is selected; the spacetime position and kinematic properties of the particle are then updated and the particle is ready for the next step to be simulated. Simulated hits in the detector are interpreted by algorithms and clustered together to reconstruct the kinematics and paths of particles in the event.

The detector's efficiency is the frequency with which it correctly records and reconstructs events. To have a realistic picture of the detector's performance, its finite resolution and inefficiencies must be included in the simulation. For any given physics search, the detector efficiency can be considered as a combination of two main factors. The first is the acceptance – the probability that a simulated particle passing through the simulated detector will be reconstructed. This depends on the geometry of the detector – where its components and dead space are located – and on the physics that produced that particle, which determines the probability that it will be produced with the right kinematic properties (e.g., momentum and scattering angle) to pass through the active part of the detector and be recorded. The other factor – the reconstruction efficiency – is the probability that the track will be reconstructed in a way that accurately represents the actual particle and its path. This depends on whether the particle satisfies the triggering criteria, and on the selection criteria used to filter out background events. The total efficiency is the product of the acceptance and the reconstruction efficiency [45] [46].

Detector acceptance can be modelled in the GEANT simulation by passing it a database of the calibration constants and detector element efficiencies measured at CMS

in calibration studies; these detector conditions are used to correct the representation of particle interactions with individual detector components and better simulate their efficiencies or inefficiencies.

Modelling reconstruction efficiencies is best illustrated by an example. The efficiency for simulated muons to pass a particular HLT path may differ from that of actual muons in data, due to the imperfection of modelling the detector acceptance. The discrepancy may vary with the trigger muon's p_T, η, and other kinematic parameters. Thus, in order to make the simulated HLT efficiency agree better with actual data, official studies are done at CMS to measure data/MC scale factors as a function of trigger muon p_T, η, and other important kinematic parameters that these discrepancies may depend on. Then, when an HLT filter is applied to a dataset of MC events, each event is weighted by the appropriate scale factor, depending on the p_T, η, etc. of the trigger muon. Scale factors are calculated for basic ID selections on reconstructed objects at CMS and are thus used for weighting the events used in MC datasets in order to represent the effect of reconstruction inefficiencies.

Chapter 4

Physics signature and search strategy

This chapter discusses the physics signature being sought and the theoretical motivations for this search, followed by a description of the datasets used in this search.

4.1 Target signature

In this dissertation, I present a search for an NMSSM physics signal in a boosted four-tau final state (cf. Section 1.3 for a brief overview of the NMSSM Higgs sector). The signature process sought is the production of an SM-like Higgs boson H followed by its decay to a pair of lighter pseudoscalar Higgs bosons a, each of which decays to a pair of taus. Due to the large mass difference between H and a, the a's are produced with a large boost.

Four production channels (Figure 4.1) for the H are considered: W and Z associated production (WH and ZH), where a high-p_T isolated muon from the vector boson decay provides a convenient trigger, gluon fusion (ggH), and vector boson fusion (VBF). The search was originally optimized for the WH mode but is sensitive to the ggH+VBF mode due to its large cross section. Since no forward jet tagging is done, the search is only sensitive to the sum of ggH and VBF, not each mode individually.

One of the $\tau\tau$ pairs is identified via the $\tau_\mu\tau_{\text{had}}$ decay topology, while no selection is made on the other $\tau\tau$ pair. The most significant backgrounds to the signal are expected to be SM W and Drell-Yan production, where the W and Z decay to muons; $t\bar{t}$ with one

or two muons in the final state; and production of heavy flavor (c and b) mesons. In all of these backgrounds, a jet is misidentified as a boosted $\tau_\mu \tau_{\text{had}}$ pair.

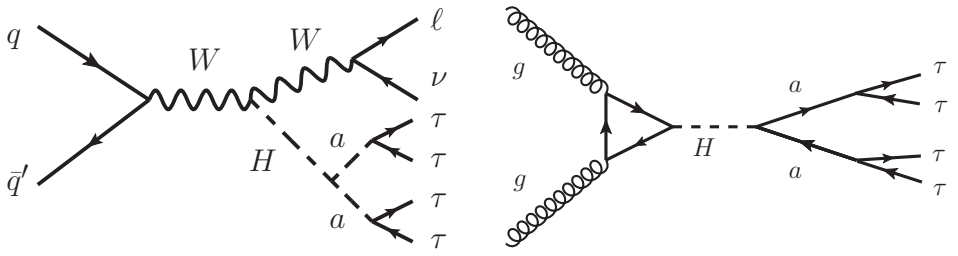

Figure 4.1. Feynman diagrams of signal processes. (Left) W associated production channel. (Right) Gluon fusion production channel.

4.2 Motivations

4.2.1 Light pseudoscalars

Following the discovery by the CMS and ATLAS experiments at the LHC [12, 11] of a Higgs-like particle H, additional measurements of its properties using the full data sets at $\sqrt{s} = 7$ and 8 TeV reveal that the observed state with a mass near 125.5 GeV is quite consistent with the standard model (SM) [47, 48, 49]; this agreement is illustrated well by the results of the most recent measurements of the Higgs production cross-sections and branching ratios shown in Figure 4.2.

It is thus clear that models with an extended Higgs sector are significantly constrained by the data. Consequently, it is interesting to explore the important possibility [51, 52] that decays of the type $H \to aa$ (where a is a lighter pseudoscalar) or $H \to hh$ (where h is a lighter scalar) are present (for reviews, see [53, 54]). Such decays are certainly possible in the context of various extensions of the SM, including two Higgs doublet models (2HDM), the next-to-minimal supersymmetric standard model (NMSSM), and purely Higgs-sector models containing additional singlet Higgs fields, but notably are not possible in the (CP-conserving) minimal supersymmetric standard model (MSSM) because of the tightly constrained nature of its Higgs sector. 2HDM studies that consider, at least briefly, the possible decays of the observed SM-like Higgs to a pair of lighter Higgs bosons include [55, 56, 54, 57, 58, 59, 60, 61, 62, 63]. Studies in the NMSSM or NMSSM-

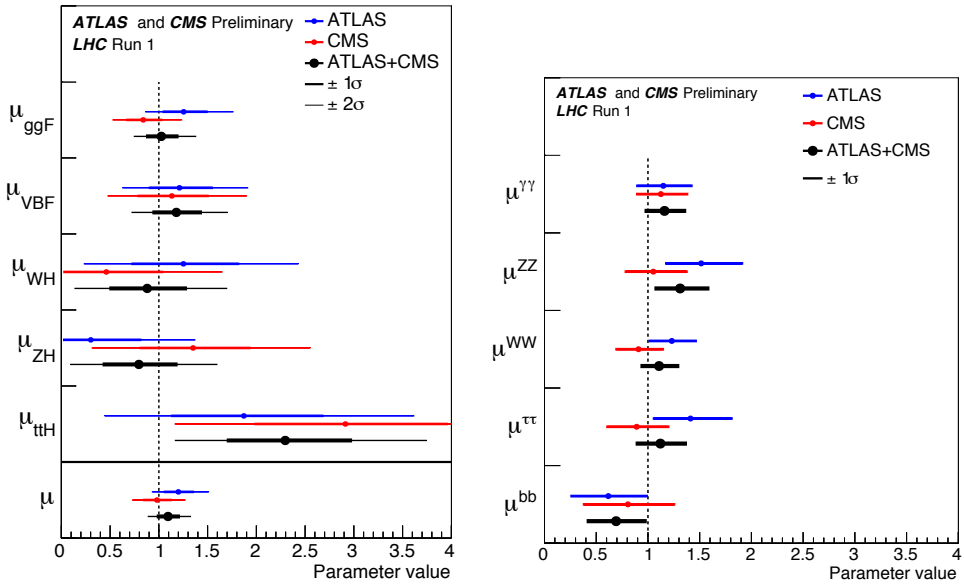

Figure 4.2. Best-fit results for signal strengths for the Higgs production cross-sections (Left) and branching ratios (Right) by the ATLAS and CMS experiments, normalized to Standard Model predictions [50].

like context include [54, 64, 65, 66, 67] and studies in the general case of adding a singlet field to the SM or the 2HDM can be found in [54, 68, 69].

The branching ratio for the H to decay to two lighter Higgs bosons is limited by the apparently SM-like nature of the H. An often-studied option is that of H decays to invisible states. However, branching ratio limits obtained under the assumption of invisibility do not apply to light Higgs pair states, which should rather be thought of as simply unseen, U, decay modes. The most thorough studies for this case are [70, 71], which combine CMS and ATLAS data. According to their latest results, it is found at 95.4% C.L. that: Br($H \to U$) $\leq$ 0.09 for a Higgs with completely SM-like couplings; Br($H \to U$) $\leq$ 0.23 for a SM Higgs if extra loop contributions to its $\gamma\gamma$ and gg couplings are allowed; and Br($H \to U$) $\leq$ 0.22 if the couplings to up quarks, down quarks and vector bosons are allowed to vary within a general model with only doublets and singlets in the Higgs sector (and no extra loop contributions to the gg and $\gamma\gamma$ couplings). If the up, down and vector boson couplings are allowed to vary completely freely (except for restricting vector boson couplings to be less than their SM values), then all LHC rates can be reproduced if all the couplings-squared are increased by a factor of 1/(1-Br($H \to U$)). The only limit in this

latter case arises from direct limits on the observed Higgs total width. At the moment, this is at the level of $\Gamma_{\rm tot} \leq 4\Gamma_{\rm tot}^{\rm SM}$ [72]. If the couplings-squared are all increased by 1/(1-Br($H\to U$)), the rates for $gg\to H\to U$ and other production mechanisms are all increased by a factor of Br($H\to U$)/(1-Br($H\to U$)), making such modes even more accessible. However, even if one adopts the more conservative approach of only considering doublets+singlets models, there is still an excellent prospect for seeing Higgs pair modes if Br($H\to U$) $\lesssim$ 0.22.

The 2HDM+singlets theory predictions for the values of Br($H\to aa$) and Br($H\to hh$) are constrained by the requirement that Br($H\to U$) $\lesssim 0.22$. This is because the required $H\to aa$ or $H\to hh$ couplings are inevitably present and are generically sizeable, and are sufficiently large that to avoid $\Gamma(H\to aa,hh) \gg \Gamma(H\to b\bar{b})$ requires significant parameter tuning (assuming m_a, $m_h < m_H/2$). For example, in the NMSSM (where the pseudoscalar mass eigenstate is defined by $a = \cos\theta_A\ a_{\rm MSSM} + \sin\theta_A\ a_S$, with $a_{\rm MSSM}$ being the MSSM-like pseudoscalar and a_S the singlet pseudoscalar of the NMSSM). $|\cos\theta_A| \ll 1$ is generically needed to keep the Haa coupling small by suppressing the doublet content of the a. In the 2HDM, fine-tuned relations among the parameters of the model are required for acceptably small ($\lesssim 0.2$) Br($H\to aa$) or Br($H\to hh$).

Direct constraints on the a or h play a role in assessing the possibilities. A previous CMS result [15] (based on [73]) places limits on $\sigma(pp\to a\to\mu\mu)$ on the order of 2$-$6 pb in the mass range from 5$-$14 GeV, excluding the upsilon resonance region. These limits, despite being based on just 1.3 fb^{-1} of 7 TeV data, can impact models. For example, [15] shows that they significantly constrain the $\cos\theta_A$ mixing angle factor defining the NMSSM. The constraints on $|\cos\theta_A|$ are especially strong at large $\tan\beta$, where $\tan\beta$ is defined as the ratio of the neutral Higgs vacuum expectation values of the two Higgs doublets in the NMSSM. In the case of the single pseudoscalar a of the CP-conserving 2HDM, points in the parameter space that are consistent with $m_H \sim 125.5$ GeV fits at 95% C.L. and other LHC and pre-LHC constraints can violate this limit in the case of Type II models (but not in the case of Type I) [74]. However, in general such constraints do not significantly restrict Br($H\to aa$) or Br($H\to hh$).

The techniques appropriate for detecting a Higgs-pair decay mode depend crucially on the mass of the lighter Higgs boson. One important possibility, particularly prominent in the NMSSM, is that the lightest CP-odd state a has mass below or not far above $2m_b$. Small m_a arises naturally in the limit of a so-called $U(1)_R$ symmetry of the model. However, a small mass for the light Higgs states is generically possible in all the models listed earlier. In addition, even if the light Higgs boson has mass above $2m_b$ (but, of course, below $m_H/2$) the $\tau\tau$ mode will still have a branching ratio of order 0.045 and will have smaller backgrounds than a purely $4b$ final state or the $bb\tau\tau$ final state. Thus, a generic exploration of the sensitivity in the 4τ final state is of considerable interest, especially as more and more integrated luminosity is accumulated in future running.

4.2.2 Semileptonic di-tau decays

This search explores the current level of sensitivity to the 4τ final state, and techniques are developed for isolating this final state from backgrounds. In particular, at least one of the tau pairs produced in the decays of the light Higgs bosons is required to decay semileptonically as $\tau_\mu\tau_{\text{had}}$. Requiring at least one hadronic tau decay is intended to maximize statistics, due to the higher branching ratio for hadronic tau decays – 64.76% compared to 17.41% and 17.83% for decays to muons or electrons respectively [75]. However, the choice for the other tau not to decay hadronically was motivated by issues with the reconstruction of boosted tau pairs, as well as the relative difficulty in discriminating fully hadronic tau pair decays from background processes.

Because of the large mass difference between the H and a, the final-state tau pairs are highly collimated, resulting in the overlap of their decay products. This spoils the isolation of the individual taus and renders their reconstruction difficult or impossible by standard means. In order to reconstruct boosted $\tau\tau$ pairs, a modified version of the standard hadron plus strips (HPS) [39] tau reconstruction procedure has been developed for this search. This method, described fully in Section 5.2.2, involves identifying and removing a leptonic tau decay candidate (muon or electron) from the PF jet used to seed the hadronic tau decay reconstruction. As this method has been successful in recovering hadronic tau ID efficiency, this search has thus focused on the reconstruction of semileptonic boosted

Dataset name	Run range
/SingleMu/Run2012A-22Jan2013-v1/AOD	190456-193621
/SingleMu/Run2012B-22Jan2013-v1/AOD	193833-196531
/SingleMu/Run2012C-22Jan2013-v1/AOD	198022-203742
/SingleMu/Run2012D-22Jan2013-v1/AOD	203777-208686

Table 4.1. Data samples.

tau decays – in particular, only $\tau_\mu \tau_{\text{had}}$ decays, due to the relative ease and cleanness with which low-p_T muons are reconstructed at CMS compared to electrons.

4.3 Datasets

4.3.1 Data samples and trigger

The datasets used in this search are the most recent SingleMu primary datasets collected by CMS in 2012 at $\sqrt{s} = 8$ TeV. The high level trigger (HLT) path HLT_IsoMu24_eta2p1 requires the presence of at least one isolated muon with $p_T > 24$ GeV found within the CMS muon coverage of $|\eta| < 2.1$. More details about the HLT muon reconstruction and isolation requirement can be found in Ref. [76]. These datasets, listed in Table 4.1, correspond to an integrated luminosity of 19.7 fb^{-1}.

4.3.2 Monte Carlo samples

The Monte Carlo (MC) samples used for the backgrounds outlined in Section 4.1 are listed in Table 4.2.

The Monte Carlo signal samples for associated WH production and gluon fusion production were generated with PYTHIA [79] and reconstructed with CMSSW version 5.3 using the S10 pileup scenario. The W in the associated W production sample is constrained to decay only leptonically. Since PYTHIA does not model NMSSM processes, the production and decay of the NMSSM scalar and pseudoscalar Higgs particles were approximated using PYTHIA's two Higgs doublet model instead. A separate sample of signal events was generated using MADGRAPH [41], which does contain methods for modeling NMSSM processes directly, and the kinematics of the PYTHIA and MADGRAPH event samples were shown to be compatible and equivalent. The benchmark for this search takes

Table 4.2. Monte Carlo background samples. Cross sections from [77] and [78].

Dataset name	Cross section (pb)
/W1JetsToLNu_TuneZ2Star_8TeV-madgraph/Summer12_DR53X-PU_S10_START53_V7A-v1/AODSIM	6601.5
/W2JetsToLNu_TuneZ2Star_8TeV-madgraph/Summer12_DR53X-PU_S10_START53_V7A-v1/AODSIM	2110.3
/W3JetsToLNu_TuneZ2Star_8TeV-madgraph/Summer12_DR53X-PU_S10_START53_V7A-v1/AODSIM	633.6
/W4JetsToLNu_TuneZ2Star_8TeV-madgraph/Summer12_DR53X-PU_S10_START53_V7A-v1/AODSIM	214
/DYJetsToLL_M-10To50_TuneZ2Star_8TeV-madgraph/Summer12_DR53X-PU_S10_START53_V7A-v1/AODSIM	14702
/DYJetsToLL_M-50_TuneZ2Star_8TeV-madgraph-tarball/Summer12_DR53X-PU_S10_START53_V7A-v1/AODSIM	3503.71
/TTJets_MassiveBinDECAY_TuneZ2star_8TeV-madgraph-tauola/Summer12_DR53X-PU_S10_START53_V7A-v2/AODSIM	245.8
/T_s-channel_TuneZ2star_8TeV-powheg-tauola/Summer12_DR53X-PU_S10_START53_V7A-v1/AODSIM	3.79
/T_t-channel_TuneZ2star_8TeV-powheg-tauola/Summer12_DR53X-PU_S10_START53_V7A-v1/AODSIM	56.4
/Tbar_s-channel_TuneZ2star_8TeV-powheg-tauola/Summer12_DR53X-PU_S10_START53_V7A-v1/AODSIM	1.76
/Tbar_t-channel_TuneZ2star_8TeV-powheg-tauola/Summer12_DR53X-PU_S10_START53_V7A-v1/AODSIM	30.7
/WW_TuneZ2star_8TeV_pythia6_tauola/Summer12_DR53X-PU_S10_START53_V7A-v1/AODSIM	54.838
/WZ_TuneZ2star_8TeV_pythia6_tauola/Summer12_DR53X-PU_S10_START53_V7A-v1/AODSIM	33.21
/ZZ_TuneZ2star_8TeV_pythia6_tauola/Summer12_DR53X-PU_S10_START53_V7A-v1/AODSIM	17.654

the masses of the NMSSM a, h_1, h_2, and h_3 to be 9, 125, 500, and 500 GeV respectively; a mass scan is performed over m_a from 5 to 15 GeV in increments of 2 GeV. The assumed cross sections for each signal sample are given in Table 4.3. These are the cross sections for SM 125 GeV Higgs production at 8 TeV [80] multiplied by BR($H \to aa \to 4\tau$) = 100%, which is why the cross sections are constant with pseudoscalar mass. The $W \to$ leptons branching ratio is included in the quoted cross sections for the WH signals.

Table 4.3. Assumed signal MC cross sections.

		Cross section (pb)
WH	$m_a = 5$ GeV	0.2296
	$m_a = 7$ GeV	0.2296
	$m_a = 9$ GeV	0.2296
	$m_a = 11$ GeV	0.2296
	$m_a = 13$ GeV	0.2296
	$m_a = 15$ GeV	0.2296
ggH	$m_a = 5$ GeV	19.27
	$m_a = 7$ GeV	19.27
	$m_a = 9$ GeV	19.27
	$m_a = 11$ GeV	19.27
	$m_a = 13$ GeV	19.27
	$m_a = 15$ GeV	19.27
ZH	$m_a = 5$ GeV	0.4153
	$m_a = 7$ GeV	0.4153
	$m_a = 9$ GeV	0.4153
	$m_a = 11$ GeV	0.4153
	$m_a = 13$ GeV	0.4153
	$m_a = 15$ GeV	0.4153
VBF	$m_a = 5$ GeV	1.578
	$m_a = 7$ GeV	1.578
	$m_a = 9$ GeV	1.578
	$m_a = 11$ GeV	1.578
	$m_a = 13$ GeV	1.578
	$m_a = 15$ GeV	1.578

4.3.3 Higgs transverse momentum reweighting for ggH

In gluon fusion Higgs production, the Higgs p_T spectrum can be significantly affected by next-to-leading logarithmic (NLL) and next-to-next-to-leading logarithmic (NNLL)

corrections, especially in the low-p_T range [81]. Thus, a set of weights binned in p_T, calculated with the Higgs p_T-reweighting HqT software [82], was applied to the Higgs p_T spectrum of signal MC events in the ggH production channel. The effect of this reweighting was observed to be quite small, as it produced a change of less than 2% in signal-to-background ratio for each signal and a change of less than 2.5% in the number of each type of signal event passing the final selection.

4.3.4 ZH and VBF production channels

A study was done to assess the contribution of ZH and VBF production channels to the expected signal significance. ZH and VBF signal samples were generated for pseudoscalar mass point $m_a = 9$ GeV and the numbers of events passing the full selection sequence were normalized to 19.7 fb^{-1} using the official SM production cross sections for ZH and VBF. Then, expected exclusion limits for the WH and ggH signal channels were calculated after the total expected yield of ZH and VBF events was distributed among the WH and ggH expected yields proportionally to their sizes, and these expected exclusion limits were compared to the nominal expected exclusion limits for WH and ggH without the added events. The combined presence of ZH and VBF signals changed the WH and ggH expected exclusion limits by at most 10%, and the change was always well within the 1σ error band of the nominal expected limits. Yet, in one of the search regions, the contribution of VBF was larger than that of WH, so ultimately it was concluded that the contributions of VBF and ZH should be considered too.

However, due to a shortage of time, ZH and VBF MC samples could not be generated for the other pseudoscalar mass points. Instead, the contribution of VBF for each pseudoscalar mass point was estimated by taking the number of surviving ggH events in the counting experiment bin after the full selection and normalizing this number to the expected SM VBF production cross-section, since the selection efficiency is expected to be the same for ggH and VBF topologies; effects due to the different H p_T spectra for the two topologies were found not to be significant. Also, since ZH and WH are expected to have similar selection efficiencies (with the ZH trigger efficiency being 1.1 times higher than for WH due to the decay of Z to two high-p_T muons rather than one), the contribu-

tion of ZH at each pseudoscalar mass point was estimated similarly by rescaling the WH contribution to the expected SM ZH production cross-section.

Chapter 5

Event selection

This chapter treats the selections used to isolate the signal process and reject background events. For the events passing the high-level trigger `HLT_IsoMu24_eta2p1` described in Section 4.3, a series of selection cuts has been developed to identify the most important physics objects in the signal – the high-p_T trigger muon, the tau decay muon τ_μ from one leg of the $a(h)$ decay, and the tau τ_{had} from the other leg of the $a(h)$ decay – and optimize the signal sensitivity. This set of cuts will be referred to as the preselection, and plays a role in the estimation of the background. The physics objects to which the selections are applied are reconstructed via the CMS particle flow (PF) algorithm.

A full list of the preselection cuts is as follows:

- Trigger μ p_T selection
- Trigger μ ID
- Trigger μ PF relative isolation selection
- τ_μ p_T selection
- τ_μ ID
- τ_{had} p_T selection
- τ_{had} HPS decay mode finding discriminator

- τ_{had} HPS isolation discriminator
- Charge requirement: $q(\text{Trigger}\mu) \cdot q(\tau_\mu) > 0$
- Charge requirement: $q(\tau_{\text{had}}) \cdot q(\tau_\mu) < 0$
- b-jet veto
- Neighbouring lepton veto around trigger muon
- Requirement of consistency with the primary vertex

Finally, events are classified into one of two bins: low transverse mass $M_{\text{T}} \leq 50$ GeV or high transverse mass $M_{\text{T}} > 50$ GeV, where $M_{\text{T}} = \sqrt{2p_T^{\text{Trig}\mu} \not{E}_{\text{T}}(1 - \cos\Delta\phi(\text{Trig}\mu, \not{E}_{\text{T}}))}$. The low-$M_{\text{T}}$ bin is sensitive to gluon fusion and VBF signal production, where there is no real W, while the high-M_{T} bin is optimized for WH production.

5.1 Trigger muon ID

Events are required to have at least one reco muon that satisfies the following criteria:

- $p_T > 25$ GeV (this is at the turn-on point for the HLT used in this search, as shown in [83])
- $|\eta| < 2.1$
- Tight muon ID:
  - The reco muon is reconstructed as a global muon and as a PF muon
  - The global muon track fit has $\chi^2/\text{ndof} < 10$ and at least one muon chamber hit
  - The reco muon has segments in at least 2 muon stations
  - The reco muon's tracker track has $d_{\text{xy}} < 2$ mm and $d_{\text{z}} < 5$ mm
  - Number of pixel hits > 0
  - More than 5 tracker layers with hits

- Relative isolation $I_{\text{rel}} < 0.12$, where the I_{rel} of a muon is defined as the pileup-corrected sum of the transverse energy of the photons and charged and neutral hadrons in a cone of radius $\Delta\text{R} = \sqrt{\Delta\eta^2 + \Delta\phi^2} = 0.4$ about the muon, divided by the p_T of the muon.

- Isolation from nearby leptons located within a cone of $\Delta\text{R} = 0.4$ around the trigger muon, where nearby lepton ID criteria are as follows:

  - **Electrons:** `reco::GsfElectron` passing PF reconstruction with $p_T > 7$ GeV and $|\eta| < 2.5$ (same as [84])
  - **Muons:** PF muon with $p_T > 5$ GeV and $|\eta| < 2.4$ passing the soft muon ID described in Section 5.2.1 and [85]
  - **Taus:** HPS PF tau with $p_T > 10$ GeV and $|\eta| < 2.3$, passing DecayModeFinding and MediumCombinedIsolationDBSumPtCorr discriminators (explained in Section 5.2.2), reconstructed from a PF jet that has been cleaned of the trigger muon with the same jet-cleaning algorithm described in Section 5.2.2. The p_T cut at 10 GeV rather than 20 GeV was chosen to make the veto more stringent.

The reco muon passing the above criteria (or, if more than one reco muon passed, the one with the highest p_T) is then matched to the object that fired the `HLT_IsoMu24_eta2p1` trigger. This is done by requiring $\Delta\text{R} < 0.1$ between the reco muon and the trigger object.

5.1.1 Neighbouring lepton veto for trigger muon

The nearby lepton isolation requirement is motivated by the desire to have a well understood trigger and PF relative isolation efficiency for the ggH and VBF production modes to which this search is sensitive. Unlike in the WH and ZH production channels, in which the particle firing the isolated muon trigger is an isolated muon from W or Z decay, the particle firing the isolated muon trigger in the ggH and VBF production channels is a muon coming from a tau decay. The difference is illustrated in Figure 5.1.

Due to the boost and low p_T of the pseudoscalar tau decay pairs in ggH and VBF events, most are rejected by `HLT_IsoMu24_eta2p1`. Those that are accepted fall into two

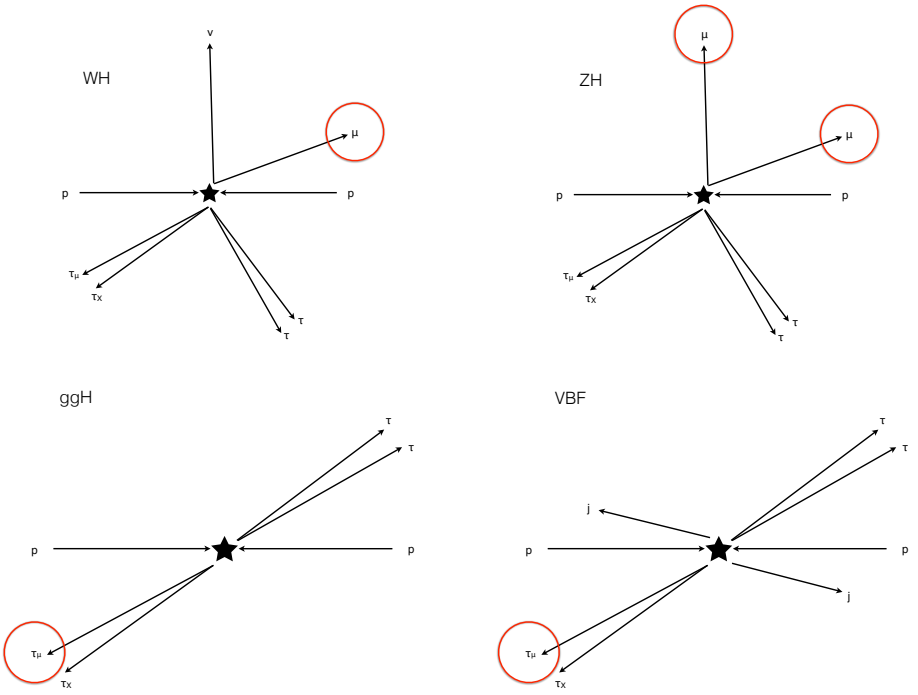

Figure 5.1. Diagrams of the four Higgs production modes considered in this search, with the triggering particle circled in red. (Top Left) WH. (Top Right) ZH. (Bottom Left) ggH. (Bottom Right) VBF.

categories:

1. $a \to \tau\tau$, one tau decays to a 24 GeV muon and the other tau decays to particles with p_T low enough to pass the HLT muon isolation cut

2. $a \to \tau\tau$, one tau decays to a 24 GeV muon and the other tau decays far enough away to not be counted in the HLT muon isolation sum

To avoid likely systematic effects in the MC description of ggH and VBF trigger and PF relative isolation efficiency due to the presence of low p_T particles around the trigger muon, events in category 1 are rejected by the nearby lepton isolation requirement. With this requirement, tau decay muons from the accepted category 2 events appear very similar to muons from Z's or W's, for which `HLT_IsoMu24_eta2p1` and PF relative isolation $<$ 0.12 efficiency measurements and standard scale factors for data-simulation differences are well understood. The following sections demonstrate that once the nearby lepton

isolation requirement is imposed, the efficiency of `HLT_IsoMu24_eta2p1` and PF relative isolation < 0.12 for ggH and VBF tau decay muons in MC is very similar to that of W decay muons in MC.

5.1.2 Study of the HLT efficiency for signal events produced via gluon fusion

The efficiency for ggH $a \to \tau \to \mu$ decay muons to fire `HLT_IsoMu24_eta2p1` is calculated for two reconstructed muon selections. The first selection is criteria described in Sec. 5.1, except that the nearby lepton isolation requirement is removed. The second selection is identical to the criteria described in Sec. 5.1. Trigger efficiency is compared for the two selections. Since the signature of pseudoscalar decays in the detector is similar between the ggH and VBF production modes, the results obtained for ggH simulation can be applied to VBF simulation as well.

The trigger efficiencies of the two selections are given by

$$\epsilon_{\text{HLT}}^{\text{no } l \text{ iso}} = \frac{\text{No. gen-matched reco'd muons passing no-lepton-isolation ID and HLT}}{\text{No. gen-matched reco'd muons passing trigger muon ID}} \quad (5.1)$$

$$\epsilon_{\text{HLT}} = \frac{\text{No. gen-matched reco'd muons passing trigger muon ID and HLT}}{\text{No. gen-matched reco'd muons passing trigger muon ID}} \quad (5.2)$$

where

- the gen-matching criteria is Δp_T(reco muon, gen $a \to \tau \to \mu$ muon) < 0.1 GeV and the gen muon is from the decay of a tau that is itself from the decay of a pseudoscalar;
- the trigger muon ID for $\epsilon_{\text{HLT}}^{\text{no } l \text{ iso}}$ is as described in Sec. 5.1 but with the nearby lepton isolation requirement removed;
- the trigger muon ID for ϵ_{HLT} is as described in Sec. 5.1; and
- "HLT" refers to firing `HLT_IsoMu24_eta2p1`.

Figure 5.2 shows $\epsilon_{\mathrm{HLT}}^{\mathrm{no}\ l\ \mathrm{iso}}$ as a function of ΔR(gen $a \to \tau \to \mu$ muon, gen τ_2), where the gen $a \to \tau \to \mu$ muon is matched to the reco'd muon as described in Eq. 5.1 above and the gen τ_2 is the other tau from the $a \to \tau\tau$ decay. The efficiency is calculated separately for each decay mode of the gen τ_2 (electronic, muonic, or hadronic). $\epsilon_{\mathrm{HLT}}^{\mathrm{no}\ l\ \mathrm{iso}}$ is $\sim$90% for $\Delta\mathrm{R} > 0.4$, when the two taus from pseudoscalar decay are separated enough that the tau decay muon appears isolated. This is similar to the efficiency of `HLT_IsoMu24_eta2p1` for W decay muons [86]. When the two taus are closer than $\Delta\mathrm{R} \sim 0.4$, the efficiency decreases because the non-triggering tau spoils the isolation of the tau decay muon that fires the trigger. The effect is worst in the $\tau_\mu\tau_e$ and $\tau_\mu\tau_\mu$ modes because electrons and muons contribute to isolation at the trigger level, but are not counted in the offline PF relative isolation.

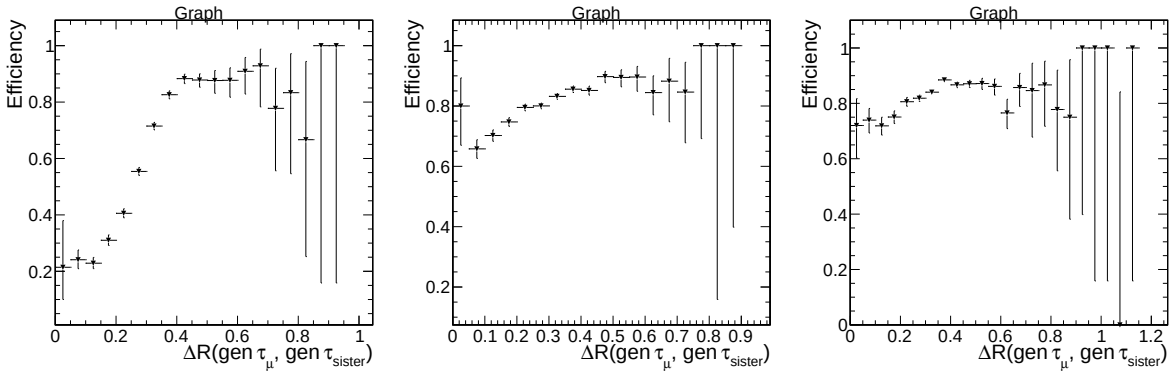

Figure 5.2. $\epsilon_{\mathrm{HLT}}^{\mathrm{no}\ l\ \mathrm{iso}}$ for the ggH signal as a function of the separation ΔR(gen $a \to \tau \to \mu$ muon, gen τ_2), where the gen τ_2 is a decay product of the same pseudoscalar as in the $a \to \tau \to \mu$. The $a \to \tau \to \mu$ muon is matched to the reco'd muon as described in the text. The reco'd muon is required to pass the trigger muon ID of Sec. 5.1, but with the nearby lepton isolation requirement removed. (Left) Gen τ_2 decays to an electron. (middle) Gen τ_2 decays to a muon. (Right) Gen τ_2 decays to hadrons.

In contrast, Figure 5.3 shows ϵ_{HLT} as a function of ΔR(gen $a \to \tau \to \mu$ muon, gen τ_2), where the gen $a \to \tau \to \mu$ muon is matched to the reco'd muon as described in Eq. 5.2 above and the gen τ_2 is the other tau from the $a \to \tau \to \mu$ decay. The efficiency is calculated separately for each decay mode of the gen τ_2 (electronic, muonic, or hadronic). The efficiencies are much flatter in ΔR when the nearby lepton isolation requirement is applied to the reconstructed trigger muon, because it ensures that events can pass the

selection sequence only if the two reconstructed taus from the pseudoscalar decay are well separated. The trigger efficiency for $a \to \tau \to \mu$ muons in these events is similar to that of W decay muons and is in the regime where the trigger muon is isolated and MC describes the data well.

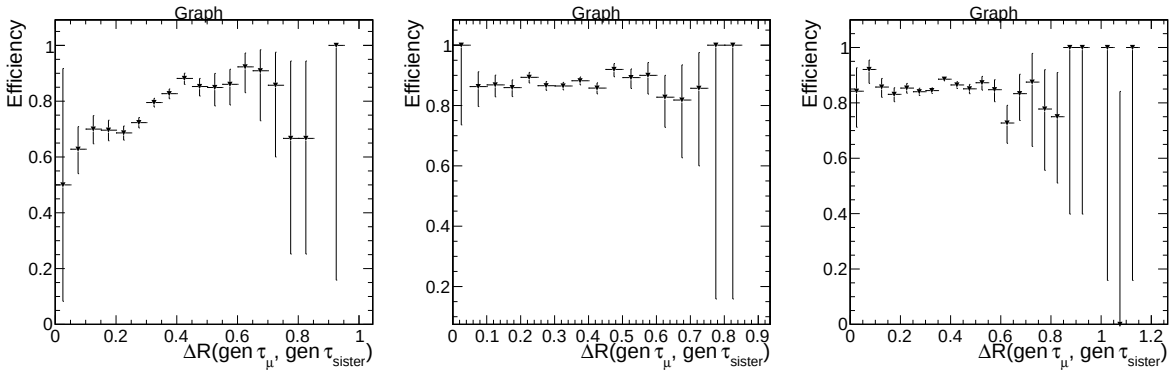

Figure 5.3. ϵ_{HLT} for the ggH signal as a function of the separation ΔR(gen $a \to \tau \to \mu$ muon, gen τ_2), where the gen τ_2 is a decay product of the same pseudoscalar as in the $a \to \tau \to \mu$. The $a \to \tau \to \mu$ muon is matched to the reco'd muon as described in the text. The reco'd muon is required to pass the trigger muon ID of Sec. 5.1. (Left) Gen τ_2 decays to an electron. (middle) Gen τ_2 decays to a muon. (Right) Gen τ_2 decays to hadrons.

Figure 5.4 shows the HLT efficiency for muons passing the trigger muon ID in both the WH and gluon fusion production modes. In both modes, the trigger muon ID includes the nearby lepton non-overlap requirement. The efficiencies are very similar for the reasons discussed above.

5.1.3 Study of the particle flow relative isolation efficiency for signal events produced via gluon fusion

The efficiency for ggH $a \to \tau \to \mu$ decay muons that pass the tight muon ID criteria to pass the PF relative isolation cut is calculated for two reconstructed muon selections. As a reminder, the PF relative isolation is defined as the p_T sum of all PF hadrons and photons within a cone of ΔR = 0.4 around the trigger muon divided by the trigger muon p_T. Both selections have the PF relative isolation requirement of Sec. 5.1 removed, since it is the efficiency of that requirement being studied here. Barring that, the first selection is identical to the criteria described in Sec. 5.1, except that in addition the nearby lepton

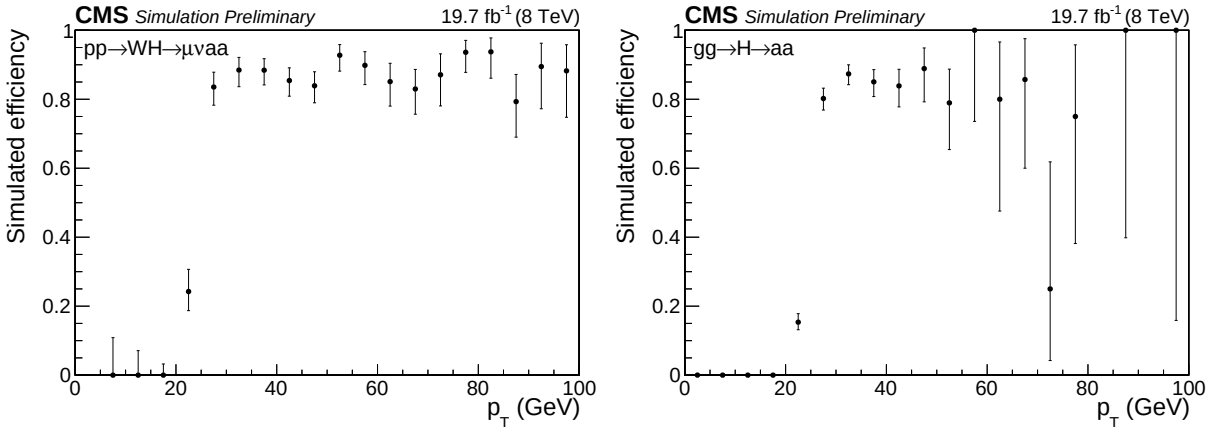

Figure 5.4. MC simulation prediction of efficiency for reconstructed muons passing the trigger muon ID to fire `HLT_IsoMu24_eta2p1`. Efficiencies were measured in MC events where the Higgs is produced via the (Left) WH and (Right) gluon fusion channels.

isolation requirement is removed. Similarly, barring the PF relative isolation requirement, the second selection is identical to the criteria described in Sec. 5.1. PF relative isolation efficiency is compared for the two selections. Since the signature of pseudoscalar decays in the detector is similar between the ggH and VBF production modes, the results obtained for ggH simulation can be applied to VBF simulation as well.

The PF relative isolation efficiencies of the two selections are given by

$$\epsilon_{\text{rel. iso}}^{\text{no } l \text{ iso}} = \frac{\text{No. gen-matched reco'd muons passing no-lepton-isolation ID and PF rel. iso.}}{\text{No. gen-matched reco'd muons passing trigger muon ID excl. PF rel. iso.}} \tag{5.3}$$

$$\epsilon_{\text{rel. iso.}} = \frac{\text{No. gen-matched reco'd muons passing trigger muon ID incl. PF. rel. iso.}}{\text{No. gen-matched reco'd muons passing trigger muon ID excl. PF rel. iso.}} \tag{5.4}$$

where

- the gen-matching criteria is Δp_T(reco muon, gen $a \to \tau \to \mu$ muon) < 0.1 GeV and the gen muon is from the decay of a tau that is itself from the decay of a pseudoscalar;

- the trigger muon ID for $\epsilon_{\text{rel. iso.}}^{\text{no } l \text{ iso}}$ is as described in Sec. 5.1 but with the PF relative isolation and nearby lepton isolation requirements removed;
- the trigger muon ID for $\epsilon_{\text{rel. iso.}}$ is as described in Sec. 5.1 but with the PF relative isolation requirement removed; and
- "rel. iso." refers to passing the cut PF relative isolation < 0.12.

Figure 5.5 shows $\epsilon_{\text{rel. iso.}}^{\text{no } l \text{ iso}}$ as a function of ΔR(gen $a \to \tau \to \mu$ muon, gen τ_2), where the gen $a \to \tau \to \mu$ muon is matched to the reco'd muon as described in Eq. 5.3 above and the gen τ_2 is the other tau from the $a \to \tau\tau$ decay. The efficiency is calculated separately for each decay mode of the gen τ_2 (electronic, muonic, or hadronic). $\epsilon_{\text{rel. iso.}}^{\text{no } l \text{ iso}}$ is $\sim$80%, independent of ΔR, for the $\tau_\mu\tau_\text{e}$ the $\tau_\mu\tau_\mu$ channels. Because PF electrons and muons are not counted in the PF relative isolation sum, the presence of a nearby τ_e or τ_μ does not significantly spoil the relative isolation of the main $a \to \tau \to \mu$ muon. The overall efficiency is lower than the efficiency for Z decay muons [86] by $\sim$15% due to the different kinematics of di-tau objects vs. isolated Z decay muons. Conversely, in the $\tau_\mu\tau_\text{had}$ channel, $\epsilon_{\text{rel. iso.}}^{\text{no } l \text{ iso}}$ is $\sim$80% only for $\Delta\text{R} > 0.4$, when the two taus from pseudoscalar decay are separated enough that the tau decay muon appears isolated. When the two taus are closer than $\Delta\text{R} \sim 0.4$, the efficiency decreases because the hadronically decaying non-identified partner tau spoils the relative isolation of the tau decay muon that is identified as a trigger muon.

In contrast, Figure 5.6 shows $\epsilon_{\text{rel. iso.}}$ as a function of ΔR(gen $a \to \tau \to \mu$ muon, gen τ_2), where the gen $a \to \tau \to \mu$ muon is matched to the reco'd muon as described in Eq. 5.4 above and the gen τ_2 is the other tau from the a $\to \tau\tau$ decay. The efficiency is calculated separately for each decay mode of the gen τ_2 (electronic, muonic, or hadronic). The efficiencies are much flatter in ΔR when the nearby lepton isolation requirement is applied to the reconstructed trigger muon, because it ensures that events can pass the selection sequence only if the two reconstructed taus from the pseudoscalar decay are well separated. The PF relative isolation efficiency for $a \to \tau \to \mu$ muons in these events is now similar to that of Z decay muons and is in the regime where the trigger muon is

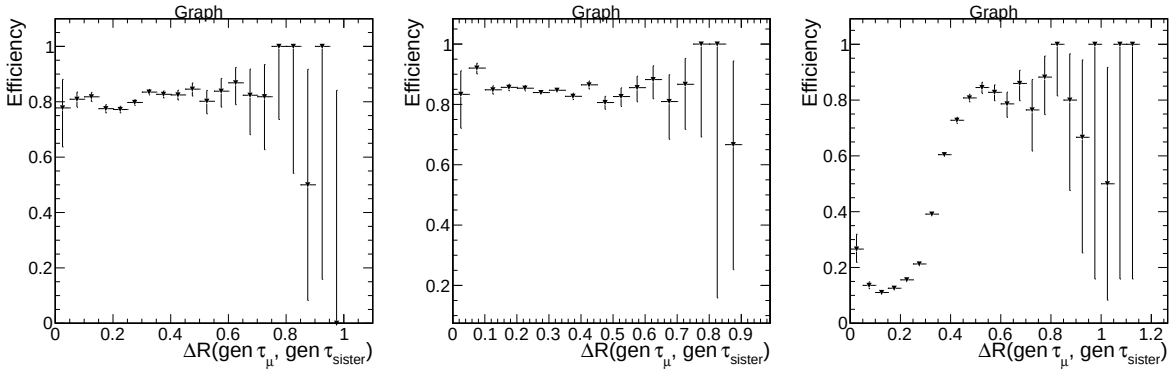

Figure 5.5. $\epsilon^{\text{no } l \text{ iso}}_{\text{rel. iso.}}$ for the ggH signal as a function of the separation ΔR(gen $a \to \tau \to \mu$ muon, gen τ_2), where the gen τ_2 is a decay product of the same pseudoscalar as in the $a \to \tau \to \mu$. The $a \to \tau \to \mu$ muon is matched to the reco'd muon as described in the text. The reco'd muon is required to pass the trigger muon ID of Sec. 5.1, but with the PF relative isolation (because this is the cut under study) and nearby lepton isolation requirements removed. (Left) Gen τ_2 decays to an electron. (middle) Gen τ_2 decays to a muon. (Right) Gen τ_2 decays to hadrons.

isolated and MC describes the data well.

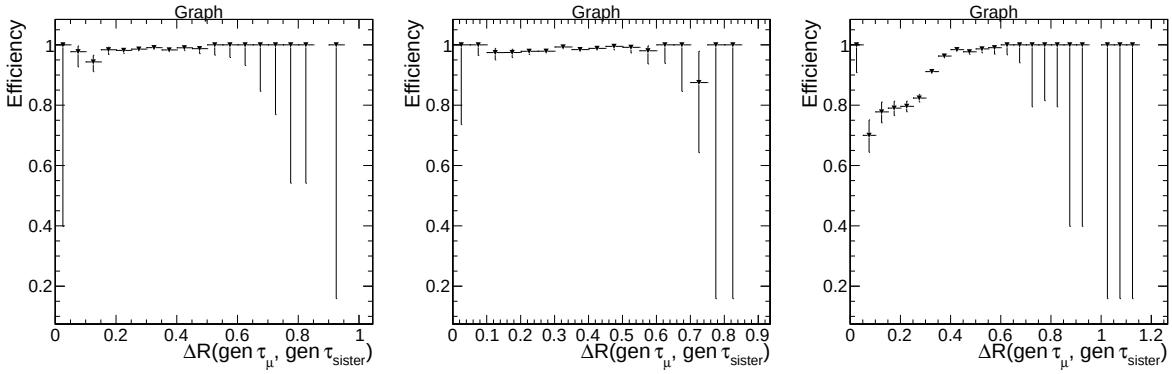

Figure 5.6. $\epsilon_{\text{rel. iso.}}$ for the ggH signal as a function of the separation ΔR(gen $a \to \tau \to \mu$ muon, gen τ_2), where the gen τ_2 is a decay product of the same pseudoscalar as in the $a \to \tau \to \mu$. The $a \to \tau \to \mu$ muon is matched to the reco'd muon as described in the text. The reco'd muon is required to pass the trigger muon ID of Sec. 5.1, but with the PF relative isolation requirement removed (because this is the cut under study). (Left) Gen τ_2 decays to an electron. (middle) Gen τ_2 decays to a muon. (Right) Gen τ_2 decays to hadrons.

After all other selection cuts, the acceptance of the nearby lepton isolation requirement ranges from 87% to 95% for ggH pseudoscalar masses 7, 9, 11, 13, and 15 GeV.

5.2 Boosted tau ID

The target signature in this search is one in which one of the a decays results in a $\tau_\mu\tau_{\text{had}}$ final state, while no constraints are placed on the decay of the taus coming from the other a. The hadronic tau identification algorithm employed in this search is the HPS algorithm.

As described in Section 4.2.2, the tau pairs produced in the pseudoscalar decays will be highly collimated, and their decay products will invade one another's isolation cones. In particular, for the $\tau_\mu\tau_{\text{had}}$ pair, the muon from τ_μ has been found to end up frequently among the constituents of the jet seeded by the τ_{had} decay and therefore among the isolation constituents of the τ_{had} reconstructed with HPS. Figure 5.7 shows the τ_{had} isolation efficiency for the standard HPS ID versus the boosted $\tau_\mu\tau_{\text{had}}$ ID described below. The isolation efficiency is about four times higher for the boosted $\tau_\mu\tau_{\text{had}}$ ID.

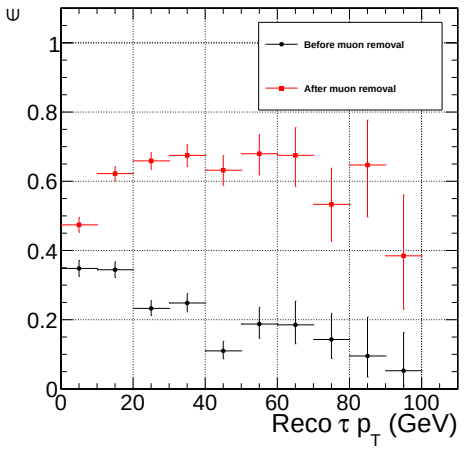

Figure 5.7. Hadronic tau isolation efficiency for the WH signal using the standard tau identification algorithm (black) and the boosted ID developed for this search (red).

To recover the correct reconstruction of the τ_{had}, soft muon candidates for the τ_μ are identified and removed from the constituents of any jet that contained them, while the remaining jet constituents are then reconstructed into a jet and passed to the HPS algorithm to reconstruct a τ_{had} decay. The soft muon ID used for identifying τ_μ candidates is described in [85].

5.2.1 Soft muon ID

In addition to the trigger μ requirement, events are selected that have at least one reco muon passing the following cuts:

- $p_T > 5$ GeV
- $|\eta| < 2.4$
- Distinct from the trigger muon
- Soft muon ID [85]:
  - Tracker muon track is matched with at least one muon segment in both x and y coordinates
  - More than 5 tracker layers with hits
  - Number of pixel layers > 1
  - The tracker muon track fit has $\chi^2/\text{ndof} < 1.8$
  - The reco muon's tracker track has $d_{\text{xy}} < 3$ mm and $d_{\text{z}} < 30$ mm

After all soft muons passing these requirements are collected, they are used as described in Section 5.2.2 to reconstruct τ_{had} objects.

5.2.2 Jet cleaning and hadronic tau ID

In this search, tau decays are reconstructed with the HPS algorithm from anti-k_{T} R = 0.5 [38] PF jets ("ak5PFJets"), where R is a parameter in the algorithm related to the radius of the jet cone. Before running HPS, jet constituents passing the soft muon ID (Sec. 5.2.1) are removed. In the majority of cases, only one soft muon is removed from a jet, but if more than one muon is removed, the highest p_T removed muon is identified as the τ_μ. A new ak5PFJet (henceforth referred to as the cleaned jet after the removal of the muon) is then reconstructed from the remaining PF constituents. These cleaned jets are submitted to the HPS algorithm and reconstructed as τ_{had} candidates.

The event is then selected to have at least one τ_{had} candidate reconstructed as above with $p_T > 20$ GeV, $|\eta| < 2.3$, and passing the HPS DecayModeFinding and MediumCombinedIsolationDBSumPtCorr discriminators. HPS is currently capable of reconstructing the following hadronic tau decay modes: single hadron (one prong decay, or one prong plus one low-energy π^0), single hadron plus one ECAL strip (one prong plus one π^0), single hadron plus two strips (one prong plus one π^0 in which the photons from the π^0 decay are well separated in the ECAL), and three hadrons (three prong decay). In addition, because no anti-muon or anti-electron discriminators are applied to the HPS object, some leptonic tau decays get counted as single hadron decays. The DecayModeFinding discriminator requires that the reconstructed HPS τ_{had} object have one of these four decay modes. Further selection on the isolation of the τ_{had} candidate helps to discriminate against fake taus reconstructed from quark or gluon jets, which tend to involve more hadronic activity and soft radiation and thus are less isolated. The HPS isolation energy, described in [87], is defined by the pileup-corrected E_T sum of the PF charged and neutral hadron and PF gamma candidates found within a $\Delta\text{R} = 0.5$ cone around the τ_{had} axis. The MediumCombinedIsolationDBSumPtCorr discriminator [88] imposes an upper limit of 1.0 GeV on the HPS isolation energy.

In the majority of events passing all these selection cuts, there is only one τ_{had} passing both HPS discriminators after the τ_μ is removed from the jet used to reconstruct it. If more than one τ_{had} object passes, then the one with the highest p_T is taken to be the τ_{had}. If more than one $\tau_\mu\tau_{\text{had}}$ object is found in the event, the one with the highest $\tau_\mu\tau_{\text{had}}$ invariant mass is chosen.

Only one $\tau_\mu\tau_{\text{had}}$ object is selected in this search. Requiring two such objects was tested, to see whether this could increase sensitivity to the WH signal, but this was observed to kill all MC background events, while only a handful of signal events and QCD control sample events survived. It would have been impossible to model the background meaningfully.

5.3 Opposite charge muon veto

The presence of two muons in the event, one of which is isolated and energetic, makes Drell-Yan di-muon production a large background. To combat this, the trigger μ and τ_μ are required to have the same electric charge. Figures 5.8 and 5.9 show the effect of this requirement on the Drell-Yan background, which is reduced by almost a factor of 20 in the low-M_T bin and 13 in the high-M_T bin, at a cost of reducing the signal by at most 50%.

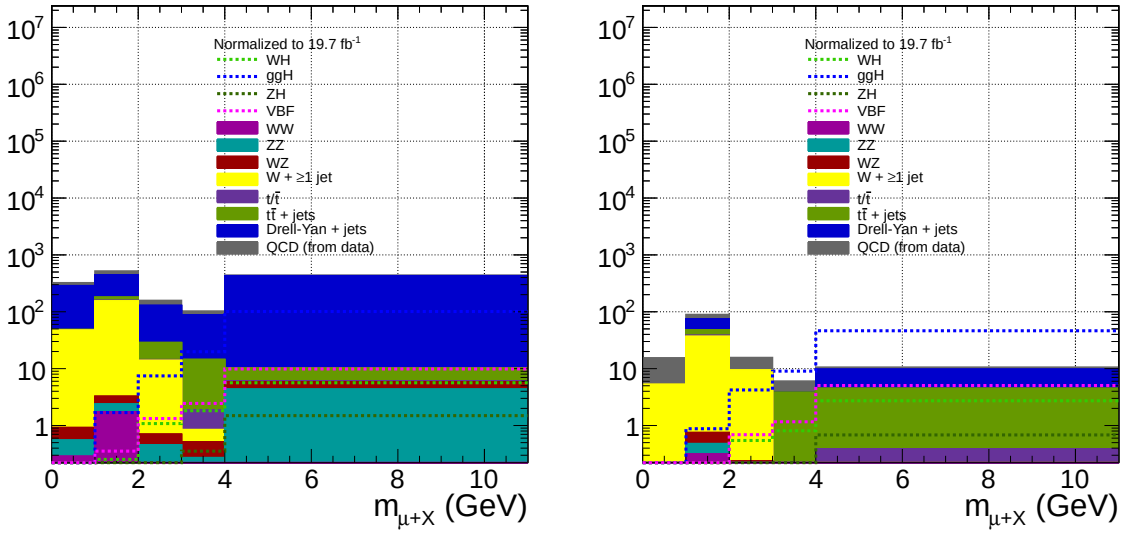

Figure 5.8. Invariant mass of the $\tau_\mu\tau_{\text{had}}$ pair for signals in the low-M_T bin with m_a = 9 GeV and all backgrounds before (Left) and after (Right) the (trigger μ)-τ_μ same charge requirement. All backgrounds except QCD are estimated from MC simulation.

5.4 Same charge tau veto

$\tau_\mu\tau_{\text{had}}$ pairs are reconstructed in background events when there is a poorly isolated real muon, either promptly produced or coming from a heavy flavor jet, or when one track in a light jet fakes a soft muon and the others fake an HPS tau. Fake $\tau_\mu\tau_{\text{had}}$ pairs rarely come from a real boosted di-tau decay, and therefore no correlation between the τ_μ and τ_{had} charge is expected. We therefore impose an opposite charge requirement on the τ_μ and τ_{had}, which reduces the background by about 20% while leaving the signal virtually unchanged (Figure 5.10).

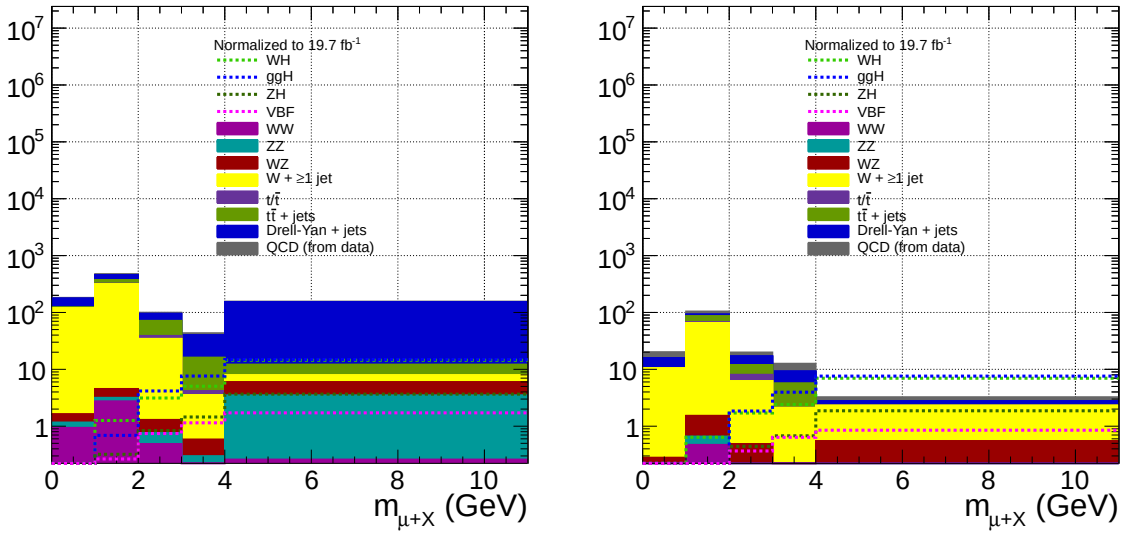

Figure 5.9. Invariant mass of the $\tau_\mu\tau_{\text{had}}$ pair in the high-M_T bin for signals with m_a = 9 GeV and all backgrounds before (Left) and after (Right) the (trigger μ)-τ_μ same charge requirement. All backgrounds except QCD are estimated from MC simulation.

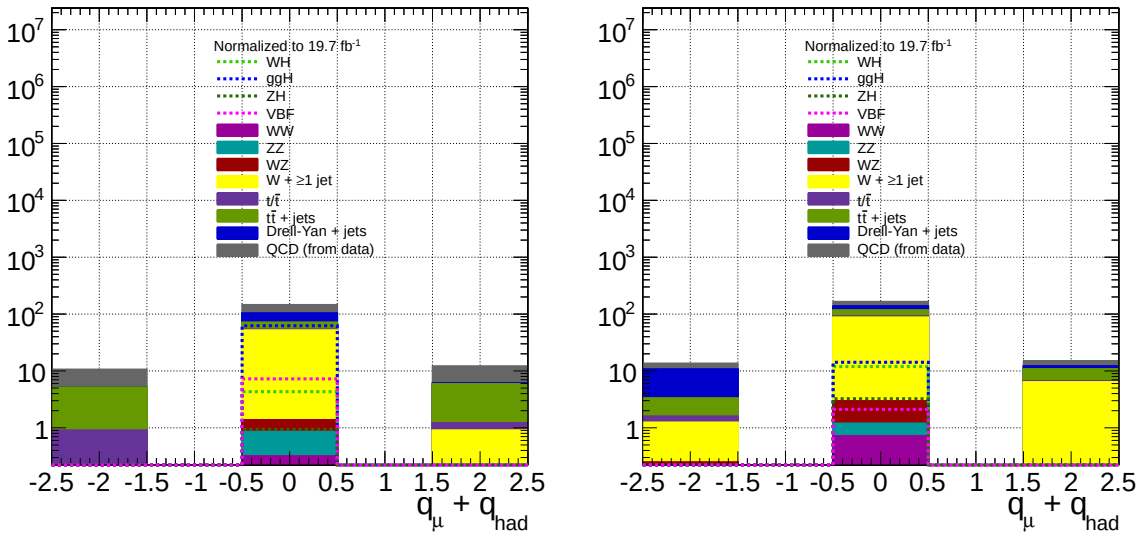

Figure 5.10. Sum of the τ_μ charge and τ_{had} charge for signals with $m_a = 9$ GeV and all backgrounds. All backgrounds except QCD are estimated from MC simulation. (Left) Low-M_T bin. (Right) High-M_T bin.

5.5 B-veto on tau jet

Heavy flavour jets, such as those from B meson or top decays, often contain a muon among the decay products which gets reconstructed as the τ_μ. Thus, the identification of b-jets can serve as a means to reject background from heavy flavour jets.

To optimize the identification of b-jets, b-tagging algorithms take advantage of the unique properties that distinguish b-jets from other kinds of jets produced at the LHC. One important property is the long lifetime of B mesons; when they decay, they will have travelled a significant distance (on the order of millimeters) from the primary vertex, resulting in displaced secondary vertices. Thus, the impact parameter and secondary vertex associated with such jets can be used as discriminating variables. The combined secondary vertex (CSV) algorithm [89], which is used in this search, employs a likelihood ratio that takes as input information about the primary vertex, secondary vertex, 2D and 3D impact parameters, track multiplicity, and track pseudorapidities of jets.

Figure 5.11 shows the distribution of the CSV discriminator for Monte Carlo signals and all backgrounds except QCD in the low-M_{T} and high-M_{T} bins. As can be seen from this figure, the τ_{had} jet in the signal final states generally does not have large values of CSV. By vetoing b-tagged jets, the single top and $t\bar{t}$ backgrounds can be cut down significantly. A b-tag veto was thus implemented by rejecting events in which the cleaned jet associated with the τ_{had} object had a CSV value greater than the medium CSV working point of 0.679 officially recommended by CMS for data at 8 TeV [90].

As shown in Figure 5.11, the distribution of the CSV discriminator for signal events tends to peak at low values of the discriminator, while single top and $t\bar{t}$ events with real b-jets peak at high values; thus, since no data/MC b-veto efficiency scale factors exist for tau jets, we make the approximation that the tau jets in our signal behave more like light jets with regard to the CSV discriminator, and so the expected signal is corrected for differences between data and MC b-veto efficiency using the BTV-provided scale factors [91] for light jets. The scale factors are binned in HPS tau parent (cleaned) jet η and p_T.

Two methods are used to assess the error on the expected signal due to the uncertainty

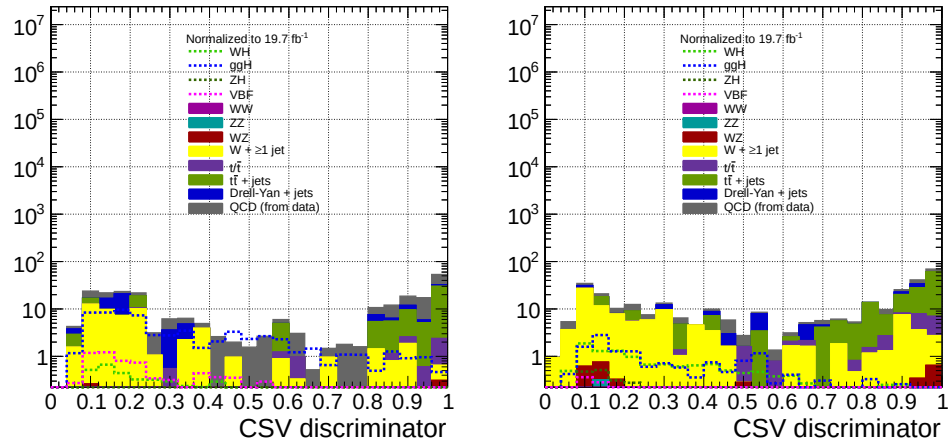

Figure 5.11. Distribution of the CSV discriminator for four signal models and all backgrounds, including data-driven QCD, after all the preselection cuts except the b veto have been applied. Normalized to 19.7 fb^{-1}. (Left) Low-M_{T} bin. (Right) High-M_{T} bin.

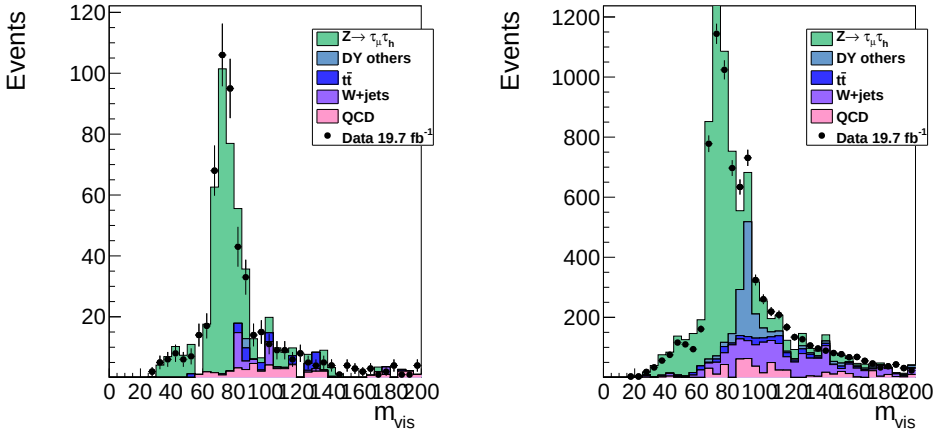

Figure 5.12. $\tau_\mu\tau_{\mathrm{had}}$ invariant mass plots showing Z peak in Run I data and MC, for events passing all Z peak selections including the medium combined isolation discriminator for τ_{had}, and passing or failing the medium CSV b-tag applied to the jet that seeded the τ_{had}. (Left) Events passing the medium CSV b-tag. (Right) Events failing the medium CSV b-tag.

on the scale factors. Firstly, the expected signal is recalculated for a coherent $+1\sigma$ shift in the scale factors in all simulated events, and then again for a -1σ shift; the difference between the nominal and $\pm1\sigma$-shifted expected signal is taken as the $\pm1\sigma$ error due to b-tag scale factors for light jets. Secondly, another systematic is calculated to account for the uncertainty of using light-jet scale factors for tau jets; the signal yields after the final selection are calculated using light-jet scale factors (the nominal method) and using b-jet scale factors (the logic being that the phase space for tau jets should be somewhere between the two extremes of light jets and b jets), and the percent difference in the yield is taken as a conservative uncertainty on the yield due to the usage of light-jet scale factors.

Since a b-veto is applied to a tau jet, the following cross-check has been performed to verify that the assigned systematic uncertainty for a potential data/MC discrepancy is adequate. First, a clean sample of of tau lepton candidates was obtained using $Z \to \tau_\mu \tau_{\rm had}$ selections in Run I data and MC and the Z peak was reconstructed as per the methods in [92], requiring that the $\tau_{\rm had}$ object pass the medium combined isolation discriminator. Then, additionally, a b-tag at the medium CSV working point was applied to the jets that seeded the $\tau_{\rm had}$, and two Z peaks were plotted – one for events passing the b-tag and one for events failing the b-tag. The Z peak plots are shown in Figure 5.12; these results suggest that the data/MC agreement is unaffected by whether the tau jets pass or fail the CSV b-tag, and also it can be seen that the percentage of events in data and MC that pass the medium CSV b-tag is in the neighbourhood of 10%, which is similar to the proportion of signal MC events observed to pass the medium CSV b-tag. Thus, these results lend confidence to the assumption that the requirement for the tau jet to pass the medium CSV b-veto does not significantly affect the known data/MC discrepancy covered by the present systematic uncertainty.

5.6 Primary vertex compatibility requirement

To reduce the background from $\tau_\mu \tau_{\rm had}$ pairs in which the τ_μ and $\tau_{\rm had}$ come from different pp interactions, we require $d_{\rm z}(\tau_\mu, {\rm PV}) < 0.5$ cm and $d_{\rm z}(\tau_{\rm had}, {\rm PV}) < 0.2$ cm, where PV refers to the hardest (primary) interaction in the event. The $d_{\rm z}(\tau_\mu, {\rm PV})$ and $d_{\rm z}(\tau_{\rm had}, {\rm PV})$

distributions for events passing all preselection cuts except those plotted are shown in Figures 5.13 and 5.14.

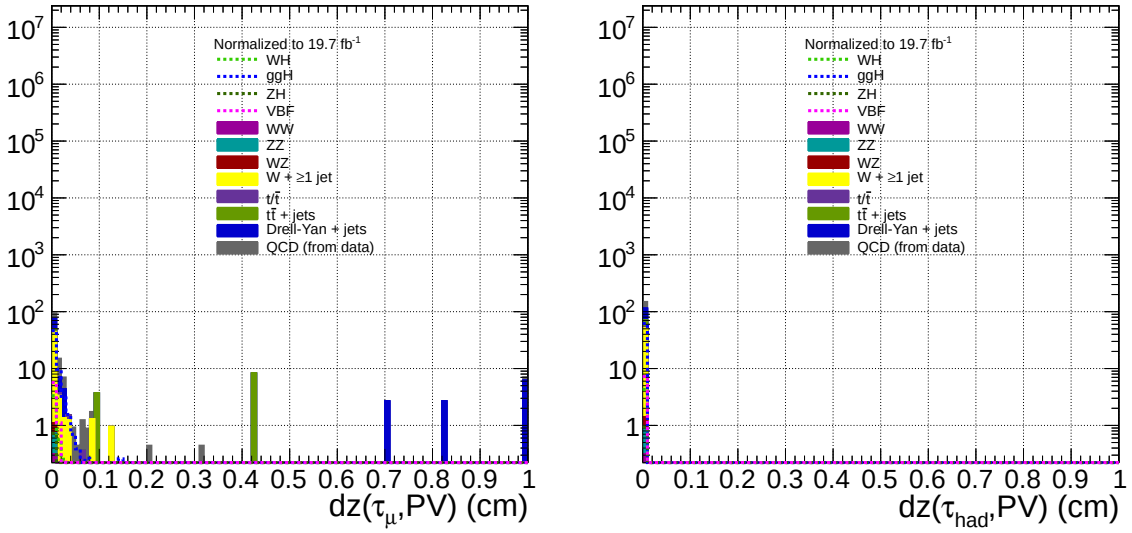

Figure 5.13. Distribution of (Left) dz(τ_μ,PV) and (Right) dz(τ_{had},PV) in the low-M_T bin for four signal models and all backgrounds including data-driven QCD, after all the preselection cuts except the dz cuts have been applied. Normalized to 19.7 fb^{-1}.

5.7 Transverse mass regions

Figure 5.15 shows the M_{T} distribution for four signal models and all backgrounds. The gluon fusion signal is clustered at low M_{T}, where Drell-Yan and QCD are the most important backgrounds. The WH signal can be found in the high M_{T} bin, where W+jets and $t\bar{t}$ dominate the background. We define $M_{\text{T}} \leq 50$ GeV as the low-M_{T} bin of this search, sensitive to ggH and VBF, and $M_{\text{T}} > 50$ GeV as the high-M_{T} bin, sensitive to WH. The cut was chosen to optimize $\text{S}/\sqrt{\text{S}+\text{B}}$ for the WH signal in the high-M_{T} bin. The M_{T} in this search is calculated using Type I-corrected [93] PAT $\not\!\!E_{\text{T}}$, which is also used in the calculation of $\not\!\!E_{\text{T}}$ systematics [94].

Following the JME approved procedure [94], uncertainty on the expected signal in each M_{T} bin due to $\not\!\!E_{\text{T}}$ scale is assessed by independently varying the e/γ, muon, tau, jet, and unclustered energy scales up and down by their approved 1σ errors for each event in the signal sample. $\not\!\!E_{\text{T}}$ and M_{T} are recalculated in each event, yielding an expected signal

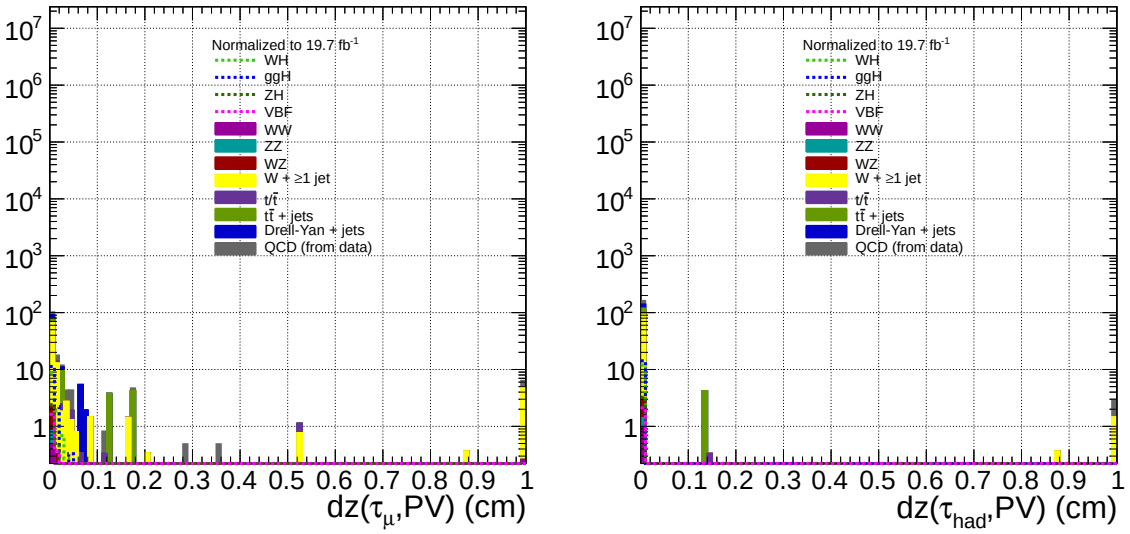

Figure 5.14. Distribution of (Left) dz(τ_μ,PV) and (Right) dz($\tau_{\rm had}$,PV) in the high-M_T bin for four signal models and all backgrounds including data-driven QCD, after all the preselection cuts except the dz cuts have been applied. Normalized to 19.7 fb^{-1}.

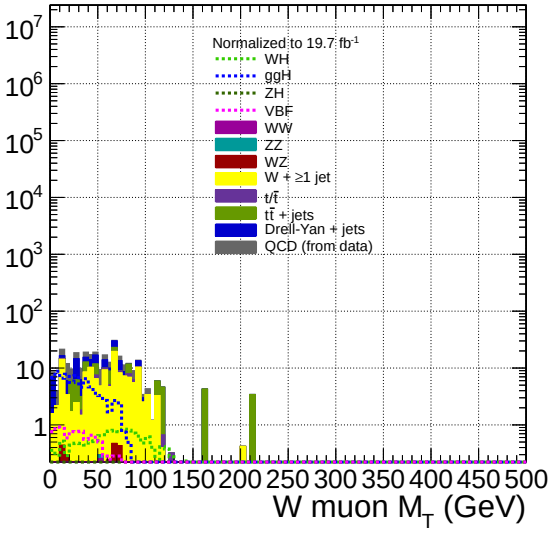

Figure 5.15. $M_{\rm T}$ distribution after the preselection (excluding the $M_{\rm T}$ cut) has been applied for four signal models and all backgrounds. The term "W muon" in the label refers to the trigger muon, not necessarily a muon from a W decay (as in the case of the ggH signal, for instance).

estimate in each of the $+1\sigma$ and -1σ scenarios. For each energy scale variation, the larger of the $\pm1\sigma$ deviations from nominal is taken as the error due to the uncertainty on that energy scale. The quadrature sum of these individual errors is taken as the total $\pm1\sigma$ error due to $\not\!\!E_{\mathrm{T}}$ scale.

For technical reasons, the $\not\!\!E_{\mathrm{T}}$ definition for the $\pm1\sigma$ varied $\not\!\!E_{\mathrm{T}}$ collections and the nominal from which deviations are measured is slightly different from the $\not\!\!E_{\mathrm{T}}$ definition used when quoting the expected signal. However, the deviations for the $\not\!\!E_{\mathrm{T}}$ uncertainty calculation are measured in a consistent way (same $\not\!\!E_{\mathrm{T}}$ definition for varied and nominal collections), and it is only the percent difference which is quoted as the $\not\!\!E_{\mathrm{T}}$ scale error.

5.8 Search region

Table 5.1 shows the number of events surviving each successive cut in the selection sequence for the $m_a = 9$ GeV WH and ggH signal samples and all background Monte Carlo samples (except for QCD Monte Carlo, due to poor statistics) used in the selection sequence optimization. Table 5.2 displays the selection efficiencies for the WH signal samples for each selection cut, expressed as the fraction of triggered signal events (i.e., events passing the HLT) surviving after each cut. Table 5.3 shows the analogous selection efficiencies for the ggH signal samples. The number of events is scaled to 19.7 fb^{-1} using the cross sections given in Tables 4.2 and 4.3.

Compared to WH, `HLT_IsoMu24_eta2p1` rejects a large fraction of ggH events (factor 10 *vs.* factor 4). However, the ggH cross section is 80 times larger than the WH cross section, making ggH an important signal in this search. Once the HLT selection has been applied, the acceptance of the trigger muon selection, $\tau_\mu\tau_{\mathrm{had}}$ selection, b-veto, and event-level cuts $q(\tau_\mu) \times q(\text{trigger } \mu) > 0$ and $q(\tau_\mu) \times q(\tau_{\mathrm{had}}) < 0$ is larger for the WH samples than for the ggH samples by factors 1.3-2 depending on pseudoscalar mass. A large portion of that difference is explained by the better acceptance of the trigger muon ID for W decay muons than for ggH $a \to \tau \to \mu$ muons.

In both the WH and ggH samples, the trigger muon + $\tau_\mu\tau_{\mathrm{had}}$ ID selects 1.7-4.9% of triggered events depending on pseudoscalar mass. The main contributors to this accep-

Table 5.1. Number of events in MC signal and background datasets remaining after each cut in the selection sequence. The signal samples have a pseudoscalar mass of 9 GeV. The number of events is scaled to 19.7 fb^{-1} using the cross sections given in Tables 4.2 and 4.3. For the rows labeled "d_Z to PV" and "$m_{\mu+\text{had}} > 4$ GeV", pileup reweighting has been applied, while for the other rows, no pileup reweighting has been applied.

Cut	WH	ggH	W+≥1 jet	Drell-Yan + jets	$t\bar{t}$ + jets	Single top	WZ	ZZ	WW
	4.522×10^3	6.721×10^4	1.883×10^8	3.587×10^8	4.842×10^6	1.825×10^6	6.542×10^5	3.478×10^5	1.080×10^6
`HLT_IsoMu24_eta2p1`	1.024×10^3	6.616×10^3	2.723×10^7	1.507×10^7	6.359×10^5	1.124×10^5	5.203×10^4	1.937×10^4	1.167×10^5
Trigger μp_T	1.001×10^3	6.083×10^3	2.633×10^7	1.459×10^7	6.238×10^5	1.091×10^5	5.090×10^4	1.913×10^4	1.138×10^5
Trigger $\mu\eta$, ID, and isolation	9.065×10^2	4.417×10^3	2.410×10^7	1.381×10^7	5.543×10^5	9.936×10^4	4.732×10^4	1.816×10^4	1.043×10^5
$\tau_\mu p_T$, η, and ID	2.702×10^2	2.600×10^3	2.950×10^5	3.705×10^6	1.449×10^5	1.430×10^4	7.300×10^3	6.842×10^3	6.138×10^3
HPS η and ID	1.270×10^2	9.838×10^2	7.518×10^4	9.074×10^4	6.805×10^4	6.508×10^3	7.541×10^2	6.257×10^2	7.861×10^2
HPS p_T and isolation	39.06	216.0	1145	4338	798.3	85.56	24.53	24.35	13.18
b-veto	34.28	194.8	903.6	4170	273.2	20.40	20.41	21.93	10.26
HLT matching	33.99	194.0	903.6	4088	266.1	20.40	19.95	21.33	10.05
$q(\tau_\mu) \times q(\text{trigger } \mu)$	16.70	92.56	171.0	93.21	81.60	8.365	3.075	1.455	1.404
$q(\tau_\mu) \times q(\tau_{\text{had}})$	16.55	91.80	154.3	68.84	67.41	6.901	2.159	1.171	1.296
Trigger μ nearby lepton filter	16.37	79.80	153.1	68.84	63.86	6.901	2.159	1.171	1.296
$M_{\mathbf{T}} >$ **50 GeV**	12.04	14.91	95.40	22.66	39.03	5.389	1.701	0.4968	0.8642
d_Z to PV	11.77	13.52	84.15	20.85	25.10	4.306	1.783	0.4766	0.7122
$m_{\mu+\text{had}} > 4$ GeV	6.972	7.598	1.761	0.4630	0	0	0.4038	0.1447	0
$M_{\mathbf{T}} \leq$ **50 GeV**	4.326	64.89	57.67	46.18	24.84	1.512	0.4580	0.6743	0.4321
d_Z to PV	4.225	60.79	50.30	32.03	17.62	1.839	0.5187	0.5394	0.3159
$m_{\mu+\text{had}} > 4$ GeV	2.723	46.35	0	5.248	4.104	0.2781	0	0.1082	0

Table 5.2. Selection efficiencies for MC WH signal samples, expressed as the fraction of triggered events surviving each selection cut.

Cut		$m_a = 5$ GeV	$m_a = 7$ GeV	$m_a = 9$ GeV	$m_a = 11$ GeV	$m_a = 13$ GeV	$m_a = 15$ GeV
Trigger μp_T		0.9780	0.9779	0.9769	0.9765	0.9754	0.9739
Trigger $\mu\eta$, ID, and isolation		0.8841	0.8841	0.8850	0.8801	0.8795	0.8802
$\tau_\mu p_T$, η, and ID		0.2898	0.2814	0.2638	0.2643	0.2620	0.2570
HPS η and ID		0.1688	0.1476	0.1240	0.1189	0.1051	0.0914
HPS p_T and isolation		0.0485	0.0439	0.0381	0.0379	0.0338	0.0282
b-veto		0.0407	0.0382	0.0335	0.0334	0.0305	0.0256
HLT matching		0.0404	0.0380	0.0332	0.0332	0.0303	0.0254
$q(\tau_\mu) \times q(\text{trigger } \mu)$		0.0204	0.0192	0.0163	0.0167	0.0152	0.0125
$q(\tau_\mu) \times q(\tau_{\text{had}})$		0.0202	0.0190	0.0162	0.0166	0.0150	0.0124
Trigger μ nearby lepton filter		0.0200	0.0188	0.0160	0.0165	0.0148	0.0124
$M_{\mathbf{T}} \leq$ **50 GeV**	d_Z to PV	0.00575	0.00494	0.00413	0.00480	0.00364	0.00342
	$m_{\mu+\text{had}} > 4$ GeV	0.00114	0.00148	0.00266	0.00403	0.00329	0.00335
$M_{\mathbf{T}} >$ **50 GeV**	d_Z to PV	0.0140	0.0136	0.0115	0.0111	0.0106	0.00835
	$m_{\mu+\text{had}} > 4$ GeV	0.000101	0.00379	0.00681	0.00837	0.00932	0.00774

Table 5.3. Selection efficiencies for MC ggH signal samples, expressed as the fraction of triggered events surviving each selection cut.

Cut		$m_a = 5$ GeV	$m_a = 7$ GeV	$m_a = 9$ GeV	$m_a = 11$ GeV	$m_a = 13$ GeV	$m_a = 15$ GeV
Trigger μp_T		0.9147	0.9226	0.9194	0.9169	0.9156	0.9155
Trigger $\mu\eta$, ID, and isolation		0.6093	0.6491	0.6675	0.6979	0.7328	0.7623
$\tau_\mu p_T$, η, and ID		0.4308	0.4082	0.3931	0.3973	0.4091	0.4173
HPS η and ID		0.1846	0.1679	0.1487	0.1319	0.1151	0.0918
HPS p_T and isolation		0.0364	0.0345	0.0327	0.0302	0.0250	0.0166
b-veto		0.0321	0.0303	0.0294	0.0280	0.0232	0.0153
HLT matching		0.0316	0.0301	0.0293	0.0279	0.0230	0.0152
$q(\tau_\mu) \times q(\text{trigger } \mu)$		0.0154	0.0144	0.0140	0.0133	0.0117	0.0073
$q(\tau_\mu) \times q(\tau_{\text{had}})$		0.0153	0.0143	0.0139	0.0131	0.0116	0.0072
Trigger μ nearby lepton filter		0.0102	0.0117	0.0121	0.0117	0.0110	0.0068
$M_{\mathbf{T}} \leq$ **50 GeV**	d_Z to PV	0.00816	0.00926	0.00914	0.00923	0.00787	0.00488
	$m_{\mu+\text{had}} > 4$ GeV	0.000089	0.00419	0.00701	0.00853	0.00763	0.00469
$M_{\mathbf{T}} >$ **50 GeV**	d_Z to PV	0.00112	0.00142	0.00204	0.00185	0.00234	0.00147
	$m_{\mu+\text{had}} > 4$ GeV	0	0.000381	0.00115	0.00139	0.00212	0.00130

tance are the $\tau_\mu\tau_{\text{had}}$ decay mode requirement, high tau p_T threshold of 20 GeV, and HPS isolation efficiency of $\sim$60%. For events with an identified trigger muon, the $\tau_\mu\tau_{\text{had}}$ ID accepts 4-5% of signal events but only 1 in $10^5 W$ + jets events, 1 in 10^4 Drell-Yan + jets events, and 1 in 1000 $t\bar{t}$ events. A drastic reduction in the Drell-Yan background comes from the requirement $q(\tau_\mu) \times q(\text{trigger } \mu) > 0$, and about 65% of the $t\bar{t}$ background is removed by the b-veto.

Signal versus background studies with Monte Carlo have shown that $m_{\mu+\text{had}}$, the invariant mass of the τ_μ and τ_{had}, provides good separation between the signal and the various backgrounds. The region $m_{\mu+\text{had}} < 2$ GeV is primarily background-dominated, while most of the signal distribution is found in the region $m_{\mu+\text{had}} > 4$ GeV.

The final selection consists of the preselection sequence followed by the requirement $m_{\mu+\text{had}} > 4$ GeV. Events passing the final selection constitute the signal region, where a counting experiment will be performed. The background $m_{\mu+\text{had}}$ distribution for events passing the preselection has been shown – using both Monte Carlo and QCD – to be modelled well by events that pass all preselection cuts up to and failing the HPS τ_{had} isolation cut. Thus, the expected background for events in the signal region will be estimated by normalizing the signal-poor $m_{\mu+\text{had}} < 2$ GeV sideband to match the background prediction and applying this normalization factor to the signal region.

Chapter 6

Muon and tau selection efficiency validation

The methods of identification for particle types such as muons and taus have their own inherent inefficiencies, related to factors such as acceptance and detector response (see Section 3.2 for a discussion of reconstruction efficiencies). The total efficiency for an ID algorithm in data may also differ from the efficiency in simulation, due to the fact that simulations can never perfectly represent physical processes or detector responses but are ultimately an approximation. To achieve better agreement between ID efficiency of a certain type of particle in data and simulation, scale factors are applied to MC events, so that the distributions of their relevant kinematic parameters agree with measured reconstruction efficiencies in data.

In this search, the data/MC scale factors used were the standard ones approved by CMS for the corresponding physics objects being identified. However, the uncertainties on the scale factors were often non-standard, due to some of the non-standard identification methods used (e.g., the jet-cleaning method for boosted hadronic tau ID). The uncertainties on the scale factors are nuisance parameters affecting the predicted signal yield, and also the predicted background yield used for the MC validation of the data-driven background-modelling method. This chapter presents the studies that were done to assess the uncertainties on the data/MC scale factors of the identification efficiencies for the most important physics objects in the selection: the trigger muon, the soft τ_μ

candidate, and the τ_h candidate.

6.1 Uncertainties on trigger muon data/MC scale factors

6.1.1 Tight muon ID

The tight muon ID used in the trigger muon selection is one of the standard ones used by the CMS experiment. A data/MC efficiency scale factor error of 0.5%, obtained from studies measuring the tight muon ID efficiency in $Z \to \mu\mu$ events [33], is used in the exclusion limit calculation (cf. Secs. 8.3 and 8.4).

6.1.2 HLT

For the WH and ZH signals, the error on the `HLT_IsoMu24_eta2p1` efficiency data/MC scale factor is taken as 0.2%, as measured from single-muon HLT efficiency studies in $Z \to \mu\mu$ events [83]. As the standard scale factors were computed for isolated Z decay muons, they are applicable to the isolated $W(Z)$ decay muons present in the WH(ZH) signal samples.

The trigger efficiency for the ggH and VBF samples is discussed in Sec. 5.1.2. With the nearby lepton isolation requirement on the reconstructed trigger muon, the trigger efficiency for ggH and VBF $a \to \tau \to \mu$ muons is in the regime where the trigger muon is isolated and MC describes the data well. A data/MC scale factor of 1 is applied to the ggH and VBF HLT efficiencies, with systematic uncertainty due to the small remaining inefficiency in the $\tau_\mu\tau_e$ mode taken as the difference ($\epsilon_{\mathrm{HLT}}(\Delta\mathrm{R} > 0.4)$ - $\epsilon_{\mathrm{HLT}}(\Delta\mathrm{R} > 0)$)/100 ($\epsilon_{\mathrm{HLT}}$ and ΔR are defined in Eq. 5.2 and Sec. 5.1.2, respectively). In the calculation of the error, all gen τ_2 (see Sec. 5.1.2) decay modes are integrated over. The uncertainty obtained is 4.2%.

6.1.3 Nearby lepton isolation

The nearby lepton isolation requirement is a veto on the presence of any reconstructed electron, muon, or tau within $\Delta\mathrm{R} = 0.4$ of the trigger muon, where the electron, muon, and tau selection criteria are summarized in Sec. 5.1. The selection criteria are standard

within CMS, so three additional uncertainties are used in the exclusion limit setting to cover the standard data-MC scale factor errors for the three lepton selections. They are 1% (electrons, cf. [95]), 1.5% (muons, cf. [96]), and 10% (taus). For the tau ID, Ref. [97] recommends an uncertainty of 6% for reconstructed taus with $p_T > 20$ GeV, which was increased to a conservative 10% to cover the difference in MC decay mode finding efficiency for 10 GeV $< p_T <$ 20 GeV between isolated hadronic taus from $Z \to \tau\tau$ and hadronic taus in $\tau_\mu\tau_{\text{had}}$ objects in the WH signal sample.

6.1.4 Particle flow relative isolation

For the WH and ZH signals, the error on the PF relative isolation efficiency data/MC scale factor is taken as 0.2%, a standard value measured in studies of $Z \to \mu\mu$ events in CMS [33]. As the standard scale factors were computed for isolated Z decay muons, they are applicable to the isolated $W(Z)$ decay muons present in the WH(ZH) signal samples.

The PF relative isolation efficiency for the ggH and VBF samples is discussed in Sec. 5.1.3. With the nearby lepton isolation requirement on the reconstructed trigger muon, the PF relative isolation efficiency for ggH and VBF $a \to \tau \to \mu$ muons is in the regime where the trigger muon is isolated and MC describes the data well. A data/MC scale factor of 1 is applied to the ggH and VBF HLT efficiencies, with systematic uncertainty due to the small remaining inefficiency in the $\tau_\mu\tau_{\text{had}}$ mode taken as the difference ($\epsilon_{\text{rel. iso.}}(\Delta R > 0.4)$ - $\epsilon_{\text{rel. iso.}}(\Delta R > 0)$)/100 ($\epsilon_{\text{rel. iso.}}$ and ΔR are defined in Eq. 5.4 and Sec. 5.1.3, respectively). In the calculation of the error, all gen τ_2 (see Sec. 5.1.3) decay modes are integrated over. The uncertainty obtained is 3.8%.

6.2 $\tau_\mu\tau_{\text{had}}$

The soft muon and HPS tau IDs used in this search are standard within CMS. However, they are used here in a nonstandard way, in particular for the special case where the soft muon and HPS tau are nearly overlapping. The soft muon and HPS tau efficiencies for the boosted tau signal have been studied in order to understand how these IDs perform in the signal environment. The HPS tau ID efficiency in the signal process is compared to the efficiency in the $Z \to \tau\tau$ process for which it was developed. As shown below,

the soft muon ID efficiency is quite high, so the data/MC ID efficiency scale factors and their errors are taken straight from official CMS recommendations, as measured from $Z/\gamma^* \to \tau\tau$ events in data [92]. The HPS tau ID efficiency for the signal is generally similar to the efficiency measured in $Z \to \tau\tau$ events, with some discrepancy in the lower-p_T region, an increased uncertainty on the data/MC efficiency scale factor is used.

All signal efficiency studies are performed with a Monte Carlo sample of WH signal events, with $m_a = 9$ GeV, generated as in Sec. 4.3. Signal events are required to pass the isolated muon trigger and have at least one reconstructed trigger muon according to the criteria in Sec. 5.1.

6.2.1 Soft muon ID efficiency

The soft muon efficiency ϵ_{soft} = (number of gen-matched muons with $p_T > 5$ GeV, $|\eta| < 2.1$, and passing the soft ID)/(number of gen-matched muons with $p_T > 5$ GeV, $|\eta| < 2.1$) is shown in Figure 6.1 for WH signal events. Gen-matching is done within a cone of ΔR = 0.3 around the reconstructed soft muon. The soft muon ID includes the requirement that the soft muon be distinct from the trigger muon, as described in Sec. 5.2.1.

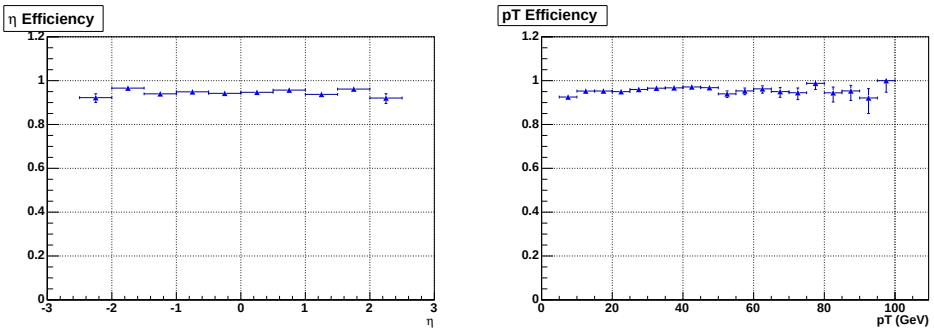

Figure 6.1. Soft muon efficiency as a function of η (Left) and p_T (Right) in WH signal events. Errors are statistical only.

The WH soft muon efficiency is ∼95% across a range of η and p_T and is in good qualitative agreement with soft muon efficiencies measured in CMS data $J/\psi \to \mu\mu$ events [86]. As the officially-measured soft muon efficiency agrees with the value from $J/\psi \to \mu\mu$ simulation [98] within the quoted error, the signal WH and ggH MC is not corrected for differences from data. Instead, recommended error of 1.5% is propagated to

the error on the expected signal.

6.2.2 HPS tau

The MC sample `/DYJetsToLL_M-50_TuneZ2star_8TeV-madgraph-tarball/Summer12_DR53X-PU_S10_START53_V7A-v1/AODSIM` is used to calculate HPS tau efficiency on $Z \to \tau_\mu \tau_{\text{had}}$ events. The τ_μ leg of the Z decay is required to fire `HLT_IsoMu24_eta2p1` and pass the trigger muon ID. HPS decay mode finding and isolation efficiency are measured on the τ_{had} leg. For WH signal events, in addition to the trigger muon ID described above, the HPS tau is required to be built from a jet cleaned of a soft muon (cf. Sec. 5.2).

The p_T distributions of gen-level taus matched to reconstructed HPS taus are shown in Figure 6.2 for the WH sample and Figure 6.3 for the $Z \to \tau\tau$ sample. Gen-matching is performed in a cone of $\Delta\text{R} = 0.3$ around the reconstructed HPS tau. Signal WH taus tend to be softer than Z decay taus, yet their ID and isolation efficiencies are similar as shown below.

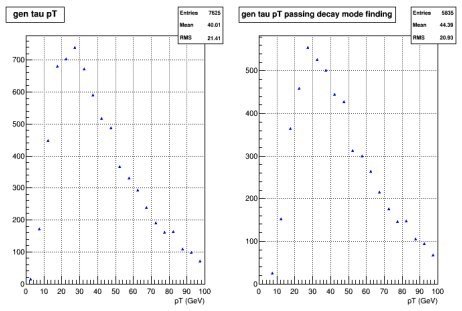

Figure 6.2. p_T of gen taus from WH $\tau_\mu \tau_{\text{had}}$ pairs matched to reconstructed HPS taus with associated soft muons (cf. Sec. 5.2). Errors are statistical only. (Left) No discriminator requirement. (Right) DecayModeFinding requirement.

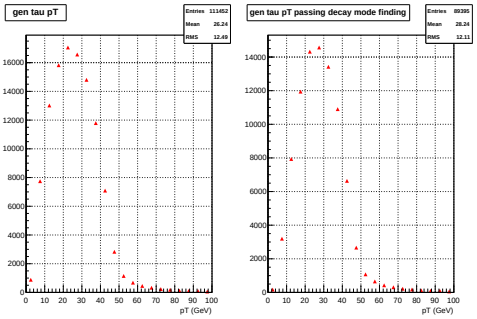

Figure 6.3. p_T of gen taus from $Z \to \tau\tau$ decay matched to standard reconstructed HPS taus. Errors are statistical only. (Left) No discriminator requirement. (Right) DecayModeFinding requirement.

The decay mode finding efficiency ϵ_{DMF} = (number of gen-matched HPS taus in $|\eta| < 2.4$ passing the DecayModeFinding discriminator)/(number of gen-matched HPS taus in $|\eta| < 2.4$) is shown for WHand $Z \to \tau\tau$ events in Figure 6.4 (Left). There is good agreement across a range of η and p_T between the simulated efficiency for signal boosted tau pair events reconstructed with the cleaning procedure described in Sec. 5.2.2 and $Z \to \tau_\mu \tau_{\text{had}}$ events. In particular, the agreement is good even for relatively low tau p_T ($<$ 20 GeV). Similarly, the decay mode finding and isolation efficiency $\epsilon_{\text{DMF+iso}}$ = (number of gen-matched HPS taus in $|\eta| < 2.4$ passing the DecayModeFinding and MediumCombinedIsolationDBSumPtCorr discriminators)/(number of gen-matched HPS taus in $|\eta| < 2.4$) is shown in Figure 6.5. There is qualitative agreement with publicly approved efficiencies for simulated $Z \to \tau\tau$ events [99].

To cover discrepancies of up to about 10% between the efficiencies in the signal and in $Z \to \tau_\mu \tau_{\text{had}}$, a conservative systematic of 10% is applied to the HPS tau ID efficiency data-MC scale factor when a p_T cut of 10 GeV is used. The p_T cut used for the τ_{had} from the $\tau_\mu \tau_{\text{had}}$ object is 20 GeV; however, this efficiency study is still important because a p_T cut of 10 GeV is applied to taus in the neighbouring lepton veto for the trigger muon (in order to have a better efficiency for the veto), and because future iterations of this search may explore the possibility of lowering the p_T cut on the τ_{had} from $\tau_\mu \tau_{\text{had}}$ as well, as was originally intended.

Since the HPS tau ID efficiencies and scale factors have been validated only

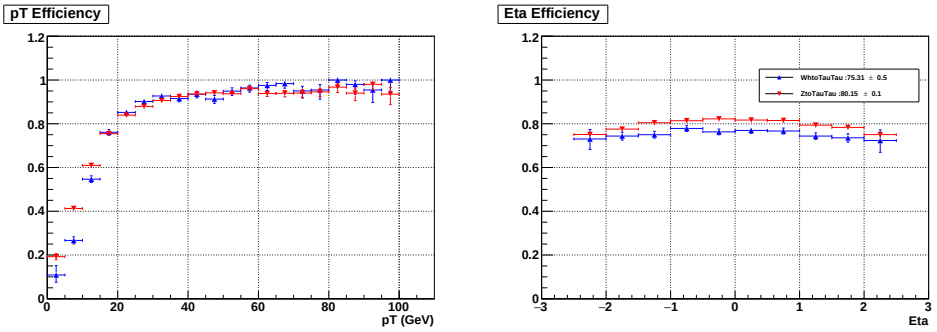

Figure 6.4. (Left) HPS decay mode finding efficiency as a function of matched gen tau p_T. (Right) HPS decay mode finding efficiency as a function of matched gen tau η. Signal HPS taus (blue) are reconstructed using the soft muon cleaning procedure described in this document, while taus from Z decay (red) are reconstructed with standard HPS. Errors are statistical only.

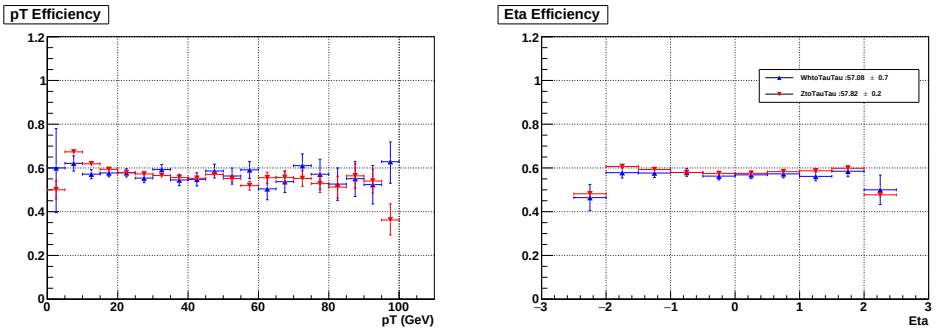

Figure 6.5. (Left) HPS decay mode finding + medium combined isolation efficiency as a function of matched gen tau p_T. (Right) HPS decay mode finding efficiency as a function of matched gen tau η. Signal HPS taus (blue) are reconstructed using the soft muon cleaning procedure described in this document, while taus from Z decay (red) are reconstructed with standard HPS. Errors are statistical only.

down to 20 GeV, a study was done to reconstruct the Z peak using HPS taus with p_T between 10 and 20 GeV and to compare it to the Z peak reconstructed from HPS taus with $p_T > 20$ GeV, to assess the reliability of using taus with p_T between 10 and 20 GeV. $Z \rightarrow \tau_\mu \tau_{\text{had}}$ events were selected in the MC sample `/DYJetsToLL_M-50_TuneZ2star_8TeV-madgraph-tarball/Summer12_DR53X-PU_S10_START53_V7A-v1/AODSIM`; the τ_μ leg of the Z decay wass required to fire `HLT_IsoMu24_eta2p1`, pass the trigger muon ID, and be gen-matched to the Z decay, while the gen-matched HPS tau was required to pass DecayModeFinding and MediumCombinedIsolationDBSumPtCorr discriminators and a p_T cut of either > 20 GeV for the standard case, or between 10 and 20 GeV for the low-p_T case of interest. As shown in Figure 6.6, the Z peak looks normal for HPS tau $p_T > 20$ GeV, and the Z peak shape for the low-p_T range looks normal aside from being biased to a lower mean due to the lower HPS tau p_T cut.

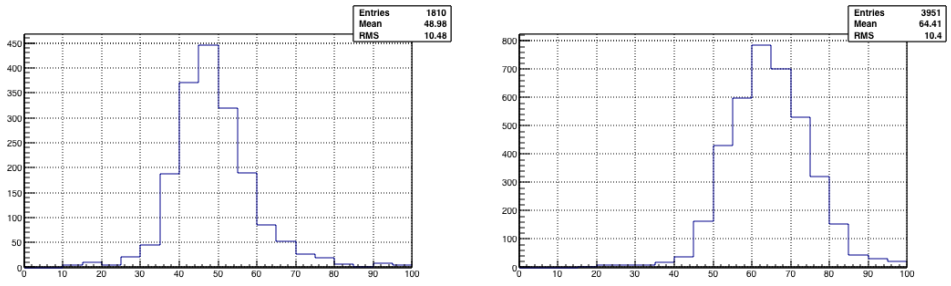

Figure 6.6. Z peak reconstructed in a sample of Drell-Yan MC events. (Left) HPS tau p_T between 10 and 20 GeV. (Right) HPS tau $p_T < 20$ GeV.

Chapter 7

Background modelling

Unlike the SM backgrounds, the new physics signal under study is characterized by the presence of a low mass $\tau_\mu\tau_{\rm had}$ resonance. Even though the τ decays cannot be fully reconstructed due to the neutrino decay products, the visible di-τ mass $m_{\mu+{\rm had}}$ distribution can still be used to discriminate the signal resonance from background τ fakes. Visible di-τ mass is defined as the invariant mass of the τ_μ and $\tau_{\rm had}$ objects described in Chapter 5. In this chapter, the method for estimating the background contribution from the most important sources will be explained. Cross-checks for assessing the significance of other backgrounds will also be presented.

7.1 Strategy

Figure 7.1 shows the distribution of $m_{\mu+{\rm had}}$ after the preselection has been applied for four signal models and all backgrounds. While the signals have broad peaks around 4 GeV, the backgrounds peak at lower values and fall off sharply.

This search is a blinded counting experiment in the signal region ("region A", cf. Figure 7.2) defined by the cuts described in Chapter 5 plus $m_{\mu+{\rm had}} \geq 4$ GeV. The value of 4 GeV has been roughly optimized for the signal-to-background ratio by eyeballing the distributions of $m_{\mu+{\rm had}}$.

The background shape is derived from a single control sample defined from the $\tau_{\rm had}$ isolation sideband ("region B", cf. Figure 7.2). It describes the shape of the background due to light jets or jets with single muon decays mis-reconstructed as $\tau_\mu\tau_{\rm had}$ pairs ("jet

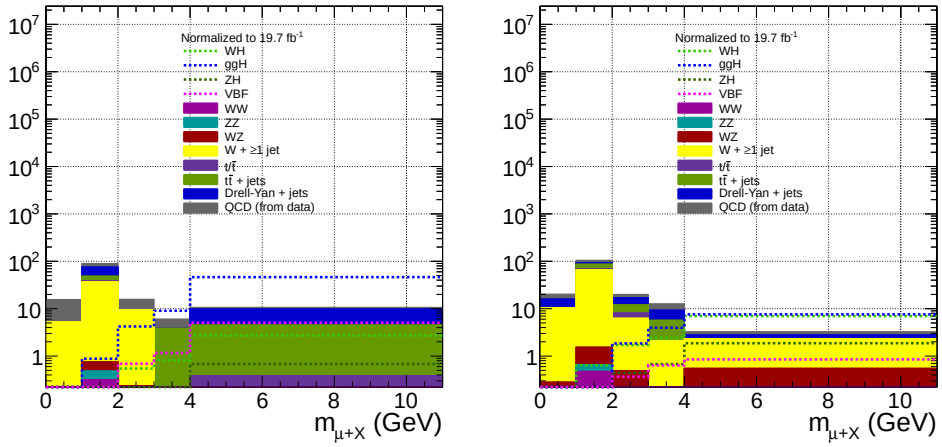

Figure 7.1. $m_{\mu+\text{had}}$ distribution after the preselection has been applied for four signal models and all backgrounds. Normalized to 19.7 fb^{-1}. (Left) Low-M_T bin. (Right) High-M_T bin.

fakes"). The signal-depleted sideband $m_{\mu+\text{had}} < 2$ GeV is used to normalize the $m_{\mu+\text{had}}$ distribution from the jet fake control sample to data passing all the cuts described in Chapter 5. This gives the nominal background prediction $N^{\text{A}}_{\text{fake; pred}}$ for signal region A; the final background prediction in the search region $m_{\mu+\text{had}} > 4$ GeV is obtained from the normalized region B distribution in the manner described later in Sec. 7.2.4, which factors in uncertainties in the background composition.

The expected SM backgrounds differ in the low and high M_T regions; in the high M_T region, trigger muons mostly come from W's, while in the low M_T region, they mostly come from jets (including heavy flavor). However, in both regions, the backgrounds are dominated by events with a $\tau_\mu\tau_{\text{had}}$ object. Thus, the assumption is made that the shape of the background $m_{\mu+\text{had}}$ distributions in regions A and B are similar. To account for possible differences in the scaling factors of different contributions between regions A and B, studies were done to investigate the changes in shape based on varying these assumptions.

7.2 Jet fake background estimation

The data control sample used for estimating the jet fake background ("region B") is defined by all of the cuts in Chapter 5, except that the tau isolation is required to be

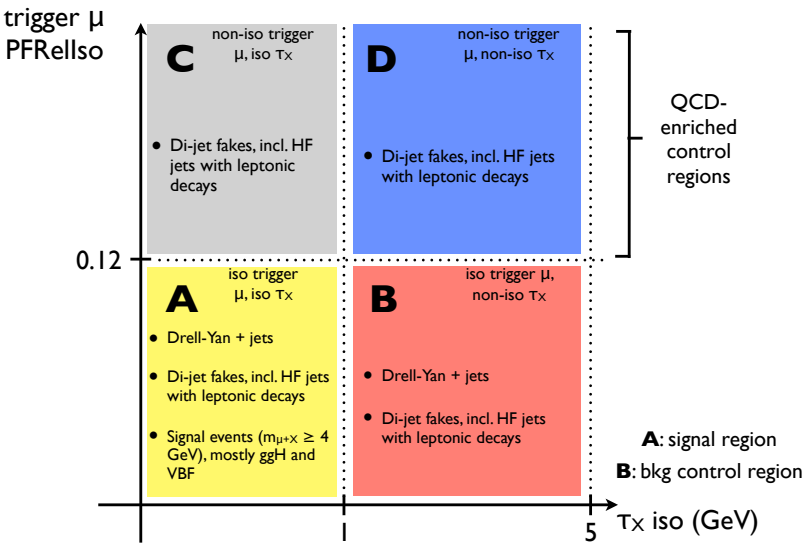

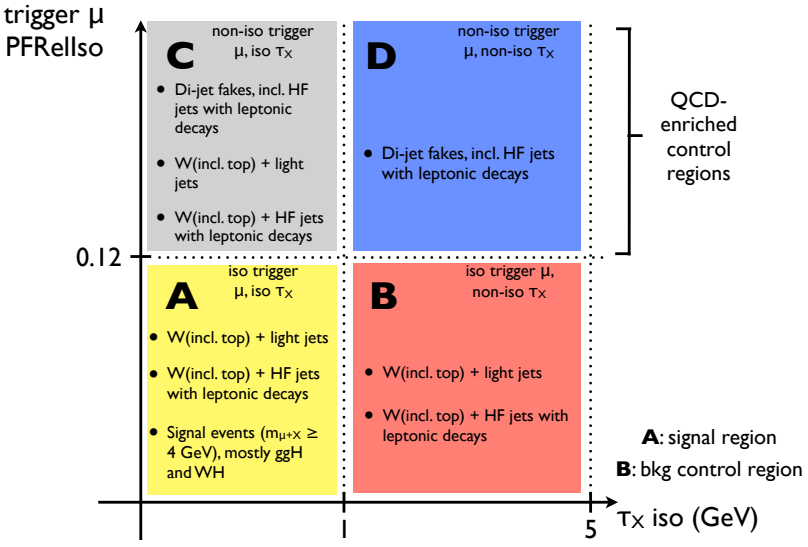

Figure 7.2. Schematic description of the signal and control regions in this search. Signal region A is defined by HPS τ isolation between 0 and 1 GeV, after events have passed all other preselection cuts detailed in Chapter 5. Likewise, jet fake control region B is defined by HPS τ isolation between 1 and 5 GeV, after events have passed all other preselection cuts. Regions C and D, which are enriched in QCD events, including those with double μ decays, are identical to regions A and B respectively, except that the trigger μ fails the tight isolation requirement and the neighbouring lepton filter is not imposed. (Top) Low M_{T}. (Bottom) High M_{T}.

between 1 and 5 GeV. Table 7.1 gives the cut values defining the search region (tau isolation $\leq$ 1 GeV); the cuts defining the jet fake control region are all identical except for the cut on tau isolation.

Table 7.1. Cut values defining the data search region A. The cuts defining the jet fake control region B are all identical, except that for having tau isolation >1 GeV && <5 GeV.

Variable	Search region
HLT	`HLT_IsoMu24_eta2p1_v[1-15]`
First muon p_T	>25 GeV
First muon $\|\eta\|$	<2.1
First muon ID (cf. Sec. 5.1)	Tight
First muon rel. iso. (cf. Sec. 5.1)	<0.12
ΔR(First muon, PF electron) (cf. Sec. 5.1.1)	>0.4
ΔR(First muon, soft muon) (cf. Sec. 5.1.1)	>0.4
ΔR(First muon, PF tau) (cf. Sec. 5.1.1)	>0.4
ΔR(first muon, HLT object)	<0.1
Second muon p_T	>5 GeV
Second muon $\|\eta\|$	<2.1
Second muon ID (cf. Sec. 5.2.1)	Soft
Second muon $\neq$ first muon	True
q(second muon) $\times$ q(first muon)	>0
Tau reco'd from jet cleaned of second muon (cf. Sec. 5.2.2)	True
Tau p_T	>20 GeV
Tau $\|\eta\|$	<2.3
Tau decay mode finding (cf. Sec. 5.2.2)	True
Tau isolation (cf. Sec. 5.2.2)	Medium ($\leq$1 GeV)
q(second muon) $\times$ q(tau)	=0
b jet veto (cf. Sec. 5.5)	CSVM
$d_z(\tau_\mu$,PV) (cf. Sec. 5.6)	<0.5 cm
$d_z(\tau_{had}$,PV) (cf. Sec. 5.6)	<0.2 cm

7.2.1 Validation of similarity of background shapes in isolated and non-isolated tau regions

For the Drell-Yan, W + jets, top, and di-boson backgrounds, MC is used to check how well the jet fake control sample (region B) defined in Table 7.1 is expected to predict

the shape of the $m_{\mu+\text{had}}$ distribution due to jet fakes in the search sample (region A). For each of the MC samples described in Table 4.2, comparisons of the shapes of the $m_{\mu+\text{had}}$ distributions between samples of events passing the search region (A) cuts and those passing the control region (B) cuts are shown in Figures 7.3 and 7.4 for the low-M_{T} bin and in Figures 7.5 and 7.6 for the high-M_{T} bin. These figures show only the shape comparisons for the individual background sources without any information about their relative ratios in regions A and B; to validate the similarity of the overall background shapes in regions A and B, a comparison of the sum of the MC backgrounds is shown in Figures 7.7 and 7.8.

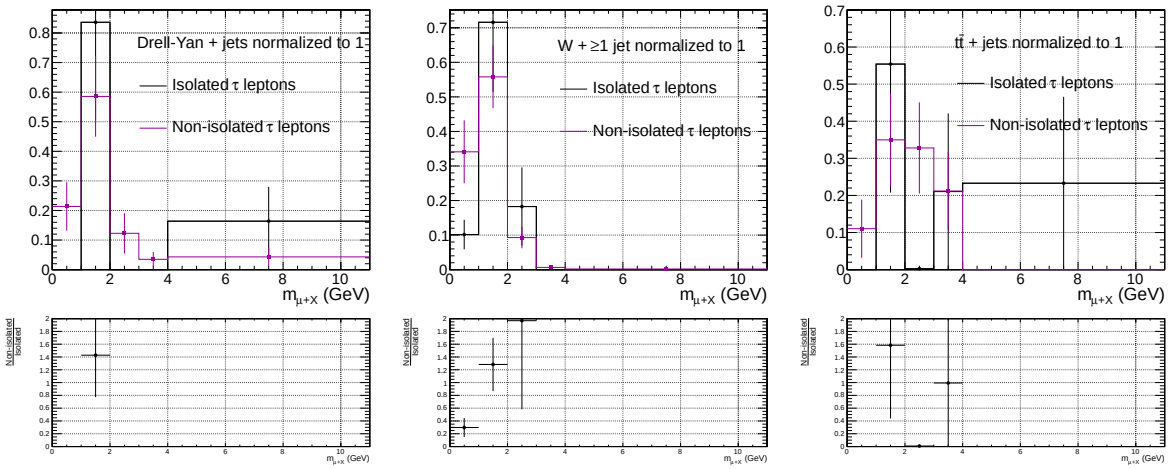

Figure 7.3. $m_{\mu+\text{had}}$ distributions in the low-M_{T} bin, normalized to one, for MC events passing the search region selection (black) and the jet fake control region selection (purple). The small plots beneath the main plots show the ratio of the control region distribution to the search region distribution. Errors are statistical only. (Left) Drell-Yan. (Middle) W + $\geq$1 jet. (Right) $t\bar{t}$.

In the high-M_{T} bin, where the main backgrounds are W + jets and $t\bar{t}$, the total $m_{\mu+\text{had}}$ shape from simulation is consistent with the total region B shape within statistical errors separately for the W + jets and $t\bar{t}$ samples (Figs. 7.5 (middle) and 7.5 (Right)). In the low-M_{T} bin, where the main backgrounds are QCD di-jets (no MC simulation available) and Drell-Yan + jets, the consistency between regions A and B separately for each MC sample is worse, but still the larger discrepancies are within statistical error.

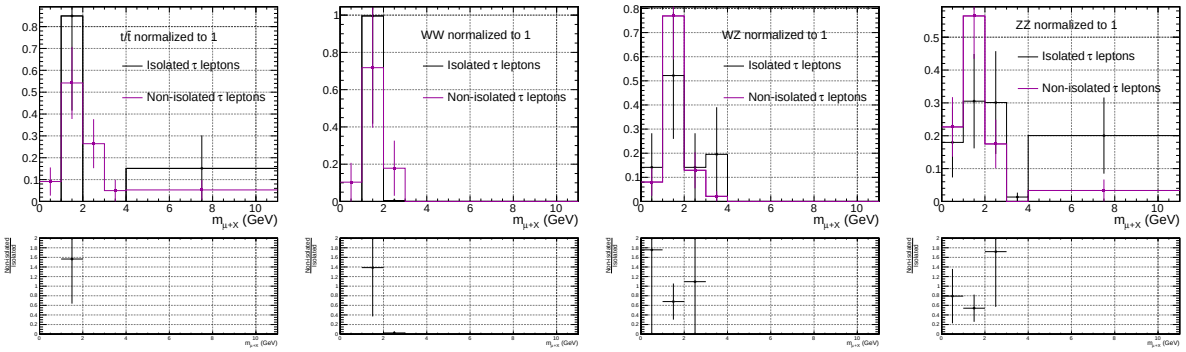

Figure 7.4. $m_{\mu+\text{had}}$ distributions in the low-M_{T} bin, normalized to one, for MC events passing the search region selection (black) and the jet fake control region selection (purple). The small plots beneath the main plots show the ratio of the control region distribution to the search region distribution. Errors are statistical only. (Left) Single top. (Middle Left) WW. (Middle Right) WZ. (Right) ZZ.

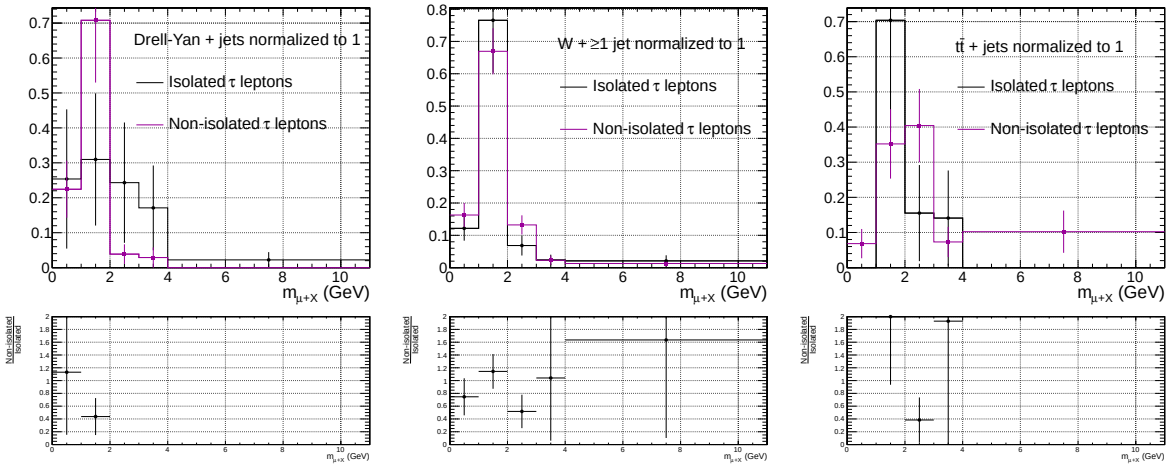

Figure 7.5. $m_{\mu+\text{had}}$ distributions in the high-M_{T} bin, normalized to one, for MC events passing the search region selection (black) and the jet fake control region selection (purple). The small plots beneath the main plots show the ratio of the control region distribution to the search region distribution. Errors are statistical only. (Left) Drell-Yan. (Middle) W + $\geq$1 jet. (Right) $t\bar{t}$.

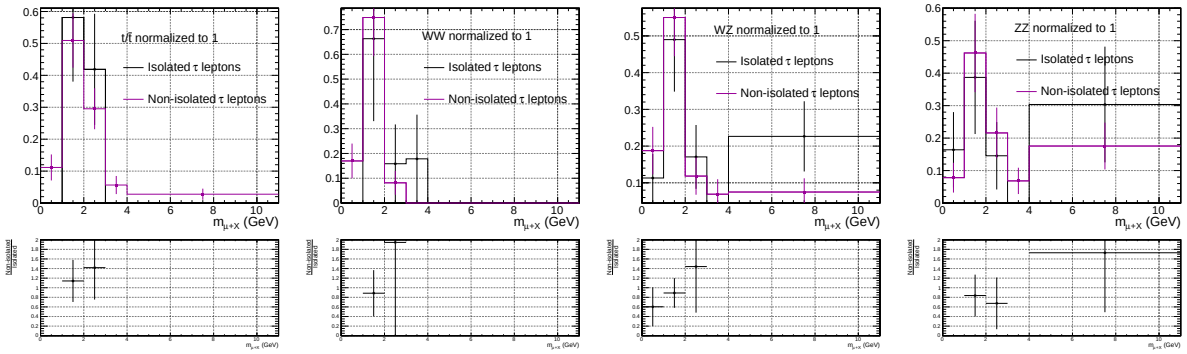

Figure 7.6. $m_{\mu+\mathrm{had}}$ distributions in the high-M_{T} bin, normalized to one, for MC events passing the search region selection (black) and the jet fake control region selection (purple). The small plots beneath the main plots show the ratio of the control region distribution to the search region distribution. Errors are statistical only. (Left) Single top. (Middle Left) WW. (Middle Right) WZ. (Right) ZZ.

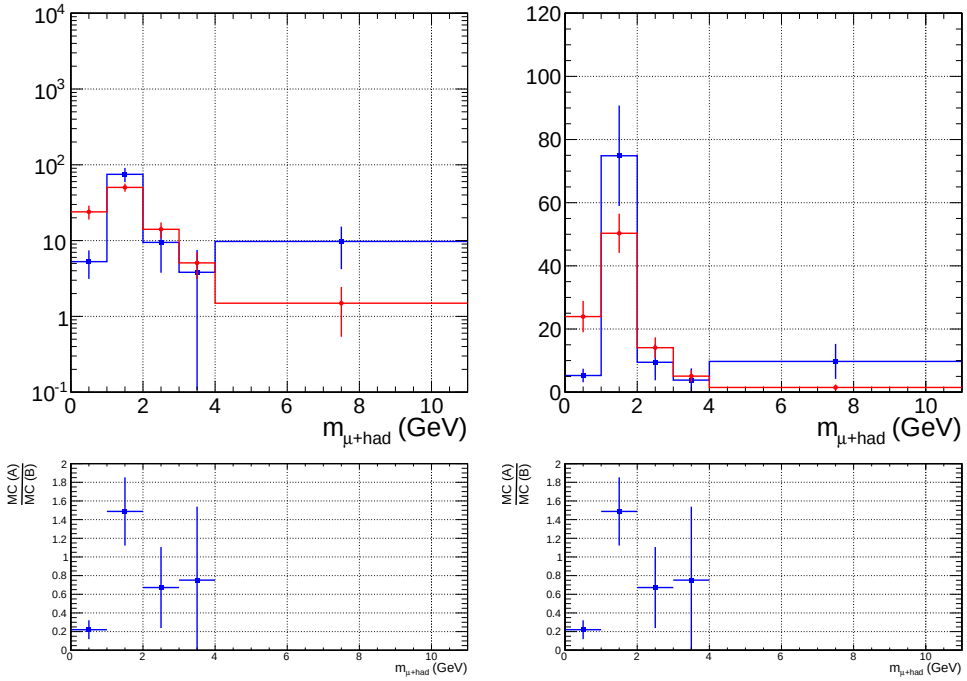

Figure 7.7. Comparison of the $m_{\mu+\mathrm{had}}$ distribution for the sum of the simulated backgrounds (all except QCD multi-jets) in the low-M_{T} bin in the signal region A (blue) with the same distribution in the sideband control region B (red), which is used to model the total background distribution in region A. The search region distribution is normalized to 19.7 fb^{-1}, while the control region distribution is normalized to the area of the search region distribution. The small plots beneath the main plots show the ratio of the search region distribution to the control region distribution. Errors are statistical only. (Left) Log scale for y axis. (Right) Linear scale for y axis.

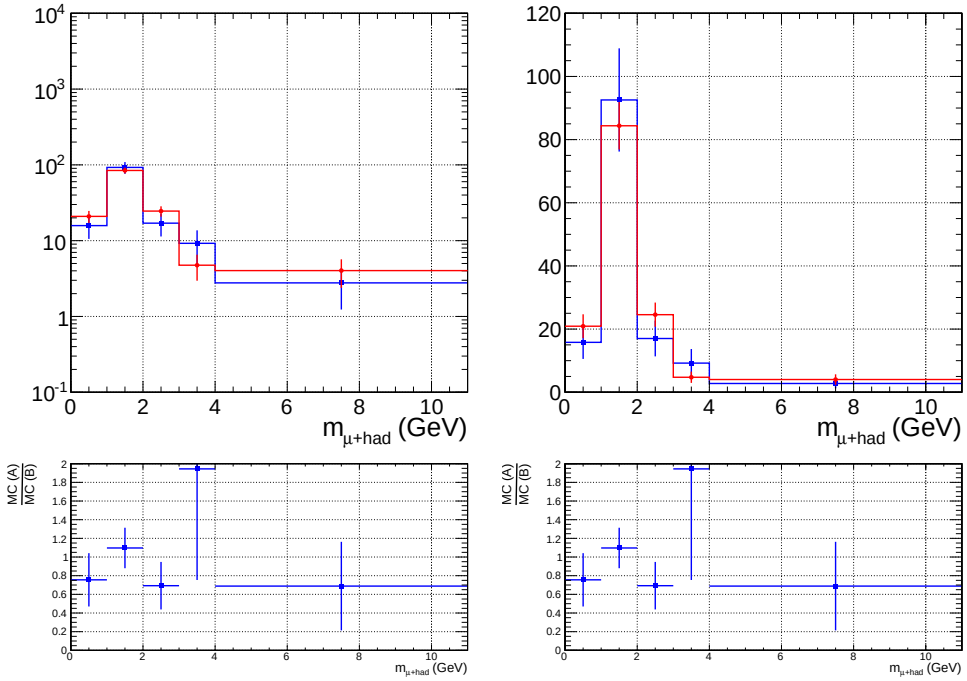

Figure 7.8. Comparison of the $m_{\mu+\mathrm{had}}$ distribution for the sum of the simulated backgrounds (all except QCD multi-jets) in the high-M_{T} bin in the signal region A (blue) with the same distribution in the sideband control region B (red), which is used to model the total background distribution in region A. The search region distribution is normalized to 19.7 fb^{-1}, while the control region distribution is normalized to the area of the search region distribution. The small plots beneath the main plots show the ratio of the search region distribution to the control region distribution. Errors are statistical only. (Left) Log scale for y axis. (Right) Linear scale for y axis.

7.2.2 QCD-enriched control regions in data

Two cross-checks are done to establish that the background from jets with double muon decays or isolated di-muon resonances ("resonances") is negligible. For these cross checks, two other control samples, defined from the trigger muon isolation sideband ("region C" and "region D", cf. Figure 7.2), are used. Since the trigger muon is required to be non-isolated, these regions are enriched in QCD di-jets, including heavy flavor jets with leptonic decays. In the $m_{\mu+\mathrm{had}}$ region around the J/ψ mass, $\tau_\mu\tau_{\mathrm{had}}$ resonances arise from $J/\psi \to \mu\mu$ decays where one muon fakes a tau. In the $m_{\mu+\mathrm{had}}$ region around the Υ mass, $\tau_\mu\tau_{\mathrm{had}}$ resonances arise from real $\Upsilon \to \tau_\mu\tau_{\mathrm{had}}$ decays. In addition, jets may have two semileptonic decays to muons. The first cross check, described in Section 7.3.1, utilizes a fit around the J/ψ mass to extract the J/ψ component, while the second, described in

Section 7.3.2, uses an independent control sample of three-muon events.

Data control regions C (isolated $\tau_\mu\tau_{\rm had}$, non-isolated trigger muon) and D (non-isolated $\tau_\mu\tau_{\rm had}$, non-isolated trigger muon, cf. Fig. 7.2) are also used to model the QCD di-jet backgrounds in regions A and B for the purpose of checking whether data with non-isolated taus (D) can accurately model data with isolated taus (C). The QCD shape in region A is taken from region C, defined by the trigger muon isolation sideband of the signal region (cf. Figure 7.2). The predicted QCD $\tau_\mu\tau_{\rm had}$ mass distribution in region A, denoted $f_{\rm A}^{\rm QCD}(m_{\mu+{\rm had}})$, is given by

$$f_{\rm A}^{\rm QCD}(m_{\mu+{\rm had}}) = R_{\rm A}^{\rm QCD} \cdot f_{\rm C}(m_{\mu+{\rm had}}) \tag{7.1}$$

where

$$R_{\rm A}^{\rm QCD} = [\int_0^{3GeV} (f_{\rm B}^{\rm data}(m_{\mu+{\rm had}}) - f_{\rm B}^{\rm MC}(m_{\mu+{\rm had}}))dm_{\mu+{\rm had}}]/[\int_0^{3GeV} f_{\rm D}(m_{\mu+{\rm had}})dm_{\mu+{\rm had}}] \tag{7.2}$$

and $f_{\rm C}(m_{\mu+{\rm had}})$ is the $\tau_\mu\tau_{\rm had}$ mass distribution in region C data. In other words, the region C data distribution is normalized by a factor similar to that in Eq. 7.4, except that the integral runs from 0 to 3 GeV instead of 0 to infinity.

This procedure assumes that the QCD $m_{\mu+{\rm had}}$ distributions in regions C and D have similar, compatible shapes. Analogously to the shape comparisons in regions A and B for the non-QCD background sources in simulation, a comparison of regions C and D is shown in Figure 7.9. The agreement is reasonable, especially in the low $M_{\rm T}$ bin for $m_{\mu+{\rm had}} \geq 2$ GeV.

To supplement the region A/B and C/D comparisons, Table 7.2 shows the number of events in the normalization sideband ($m_{\mu+{\rm had}} < 2$ GeV) and search window ($m_{\mu+{\rm had}} \geq 4$ GeV) for individual MC background sources and the data-driven QCD background in Regions A and B after the preselection, as well as the ratio of the number of normalization sideband events to the number of search window events in Regions A and B.

Table 7.2. Number of events after preselection in the normalization sideband ($m_{\mu+\text{had}} < 2$ GeV) and in the search window ($m_{\mu+\text{had}} \geq 4$ GeV) for individual MC backgrounds and the data-driven QCD background in Regions A and B, and the ratio of normalization sideband events to search window events. The MC backgrounds are normalized to 19.7 fb^{-1}. The data-driven QCD background shape for region A(B) comes from region C(D), while the normalization is discussed in Sections 7.2.1 and 7.2.3.

	Signal region (A)			Control region (B)		
$M_T \leq 50$ GeV	$m_{\mu+\text{had}} < 2$ GeV	$m_{\mu+\text{had}} \geq 4$ GeV	$\frac{m_{\mu+\text{had}} < 2\ GeV}{m_{\mu+\text{had}} \geq 4\ GeV}$	$m_{\mu+\text{had}} < 2$ GeV	$m_{\mu+\text{had}} \geq 4$ GeV	$\frac{m_{\mu+\text{had}} < 2\ GeV}{m_{\mu+\text{had}} \geq 4\ GeV}$
WW	0.315 ± 0.18	0	—	0.726 ± 0.3	0	—
ZZ	0.262 ± 0.096	0.108 ± 0.062	2.42 ± 1.7	0.998 ± 0.2	0.0421 ± 0.042	23.7 ± 24
WZ	0.344 ± 0.15	0	—	1.58 ± 0.35	0	—
W + jets	41.1 ± 10	0	—	180 ± 26	0.367 ± 0.37	491 ± 5e+02
Single top	1.56 ± 0.8	0.278 ± 0.28	5.61 ± 6.3	4.47 ± 1.3	0.373 ± 0.32	12 ± 11
$t\bar{t}$	9.76 ± 6.1	4.1 ± 4.1	2.38 ± 2.8	35.4 ± 11	0	—
Drell-Yan + jets	26.8 ± 11	5.25 ± 3.7	5.1 ± 4.1	110 ± 22	5.92 ± 4.3	18.5 ± 14
QCD (from data)	24.9 ± 7	0.89 ± 0.7	28 ± 23	117 ± 30	7.23 ± 2.8	16.2 ± 7.5
Tot. bkg. (MC + data-driven QCD)	105 ± 17	10.6 ± 5.6	9.88 ± 5.4	450 ± 46	13.9 ± 5.1	32.3 ± 12
Data	122	7	17.4	430	22	19.5
$M_T \geq 50$ GeV	$m_{\mu+\text{had}} < 2$ GeV	$m_{\mu+\text{had}} \geq 4$ GeV	$\frac{m_{\mu+\text{had}} < 2\ GeV}{m_{\mu+\text{had}} \geq 4\ GeV}$	$m_{\mu+\text{had}} < 2$ GeV	$m_{\mu+\text{had}} \geq 4$ GeV	$\frac{m_{\mu+\text{had}} < 2\ GeV}{m_{\mu+\text{had}} \geq 4\ GeV}$
WW	0.472 ± 0.24	0	—	3.66 ± 0.68	0	—
ZZ	0.263 ± 0.1	0.145 ± 0.085	1.82 ± 1.3	0.723 ± 0.17	0.235 ± 0.097	3.08 ± 1.5
WZ	1.08 ± 0.28	0.404 ± 0.17	2.66 ± 1.3	2.71 ± 0.46	0.276 ± 0.14	9.85 ± 5.2
W + jets	74.6 ± 14	1.76 ± 1.5	42.4 ± 36	409 ± 39	6.29 ± 2.8	65 ± 29
Single top	2.5 ± 0.86	0	—	15.5 ± 2.4	0.69 ± 0.43	22.5 ± 15
$t\bar{t}$	17.7 ± 8	0	—	59.9 ± 15	14.6 ± 8.6	4.11 ± 2.6
Drell-Yan + jets	11.7 ± 5.7	0.463 ± 0.46	25.4 ± 28	85.2 ± 18	0	—
QCD (from data)	16.7 ± 15	0.49 ± 0.74	34 ± 60	56.5 ± 36	2.09 ± 2	27 ± 31
Tot. bkg. (MC + data-driven QCD)	125 ± 23	3.26 ± 1.7	38.3 ± 21	633 ± 59	24.1 ± 9.2	26.2 ± 10
Data	166	14	11.9	577	20	28.9

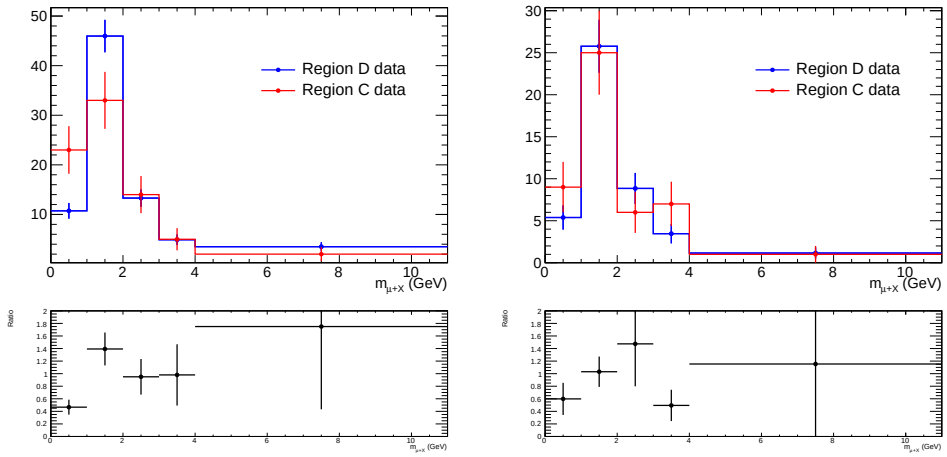

Figure 7.9. $m_{\mu+\mathrm{had}}$ distributions for data events in QCD-enriched regions C and D. The region D distribution is normalized such that $N_{\mathrm{C}}(m_{\mu+\mathrm{had}} < 3\ GeV) = N_{\mathrm{D}}(m_{\mu+\mathrm{had}} < 3\ GeV)$. The small plots beneath the main plots show the ratio of the region C distribution to the region D distribution. Errors are statistical only. (Left) Low M_{T}. (Right) High M_{T}.

7.2.3 Understanding background composition in jet fake control region (B)

In order to understand and have an independent validation of the background composition in the jet fake control region (B), comparisons of data and total non-QCD backgrounds from MC simulation are made for a number of different distributions characterizing region B. The individual MC background distributions are normalized to the data luminosity of 19.7 fb^{-1} using the cross sections given in Table 4.2, and event weights are applied to account for the different pileup distributions in data and MC according to the procedure described in Ref. [100]. The available QCD MC does not provide enough events passing the control region selection to have any statistical power, so the QCD contribution is taken instead from the QCD-enriched control region in data defined by the trigger muon isolation sideband ("region D", cf. Figure 7.2) of the jet fake control region. The predicted QCD $\tau_\mu\tau_{\mathrm{had}}$ mass distribution in region B, denoted $f_{\mathrm{B}}^{\mathrm{QCD}}(m_{\mu+\mathrm{had}})$, is given by

$$f_{\mathrm{B}}^{\mathrm{QCD}}(m_{\mu+\mathrm{had}}) = R_{\mathrm{B}}^{\mathrm{QCD}} \cdot f_{\mathrm{D}}(m_{\mu+\mathrm{had}}) \tag{7.3}$$

where

$$R_{\mathrm{B}}^{\mathrm{QCD}} = [\int_0^\infty (f_{\mathrm{B}}^{\mathrm{data}}(m_{\mu+\mathrm{had}}) - f_{\mathrm{B}}^{\mathrm{MC}}(m_{\mu+\mathrm{had}}))dm_{\mu+\mathrm{had}}]/[\int_0^\infty f_{\mathrm{D}}(m_{\mu+\mathrm{had}})dm_{\mu+\mathrm{had}}] \quad (7.4)$$

$f_{\mathrm{D}}(m_{\mu+\mathrm{had}})$ is the $\tau_\mu\tau_{\mathrm{had}}$ mass distribution in region D data, $f_{\mathrm{B}}^{\mathrm{data}}(m_{\mu+\mathrm{had}})$ is the $\tau_\mu\tau_{\mathrm{had}}$ mass distribution in region B data, and $f_{\mathrm{B}}^{\mathrm{MC}}(m_{\mu+\mathrm{had}})$ is the $\tau_\mu\tau_{\mathrm{had}}$ mass distribution in region B MC. In other words, the region D data distribution is normalized to the number of data events minus the sum of MC events in region B, where the MC sum is normalized to 19.7 fb^{-1}. Figure 7.10 shows the comparison of $m_{\mu+\mathrm{had}}$ shapes between region D data and region B data minus total non-QCD background from simulation, which are generally consistent with each other within statistical errors.

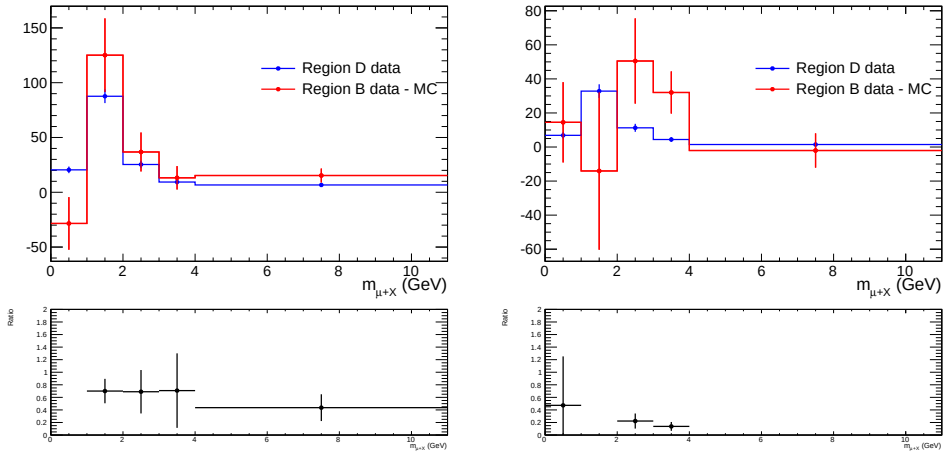

Figure 7.10. Comparison of $m_{\mu+\mathrm{had}}$ shapes between region D data and region B data minus total non-QCD backgrounds from simulation. The region D distribution is normalized such that $N_{\mathrm{D}}(m_{\mu+\mathrm{had}} < 3\ GeV) = N_{\mathrm{B}}^{\mathrm{data\ -\ MC}}(m_{\mu+\mathrm{had}} < 3\ GeV)$. (Left) Low M_{T}. (Right) High M_{T}.

Table 7.3 shows the breakdown of the expected SM contributions and observed data above $m_{\mu+\mathrm{had}} = 4$ GeV in the isolation sideband control region B, including statistical errors. Figures 7.11-7.26 show some comparisons between region B data, region B MC, and the region B QCD prediction (which is taken from region D).

Table 7.3. Expected SM events (from MC; QCD from region D data) and observed events above $m_{\mu+\text{had}} = 4$ GeV in Region B. The MC backgrounds are normalized to 19.7 fb^{-1}. The QCD normalization is given by Eqs. 7.3 and 7.4.

	$M_T \leq 50$ GeV	$M_T > 50$ GeV
WW	0	0
ZZ	0.0421 ± 0.042	0.235 ± 0.097
WZ	0	0.276 ± 0.14
W + jets	0.367 ± 0.37	6.29 ± 2.8
Single top	0.373 ± 0.32	0.69 ± 0.43
$t\bar{t}$	0	14.6 ± 8.6
Drell-Yan + jets	5.92 ± 4.3	0
QCD (from data)	7.23 ± 2.8	2.09 ± 2
Tot. expected SM	13.9 ± 5.1	24.1 ± 9.2
Data	22	20

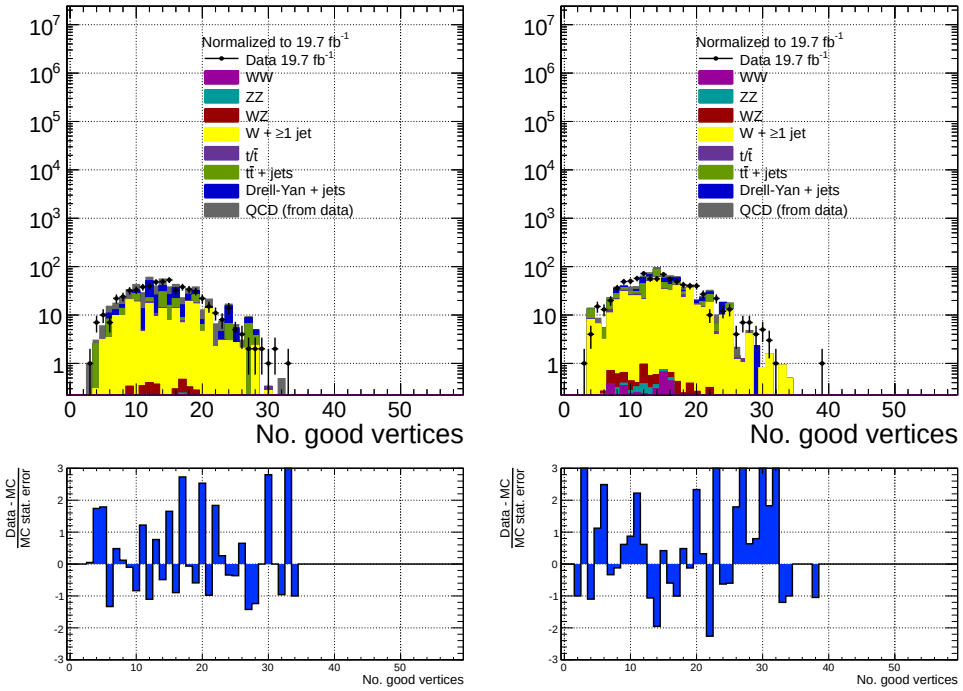

Figure 7.11. Distribution of the number of good reconstructed vertices for region B data (black points), region B total non-QCD backgrounds from MC (solid stacked histograms), and the QCD prediction from region D data (solid gray histogram). A good reconstructed vertex is required to not be fake, have >4 degrees of freedom, have z position $\leq$24 cm, and have radial position $\leq$2 cm. Errors are statistical only. (Left) Low-M_{T} bin. (Right) High-M_{T} bin.

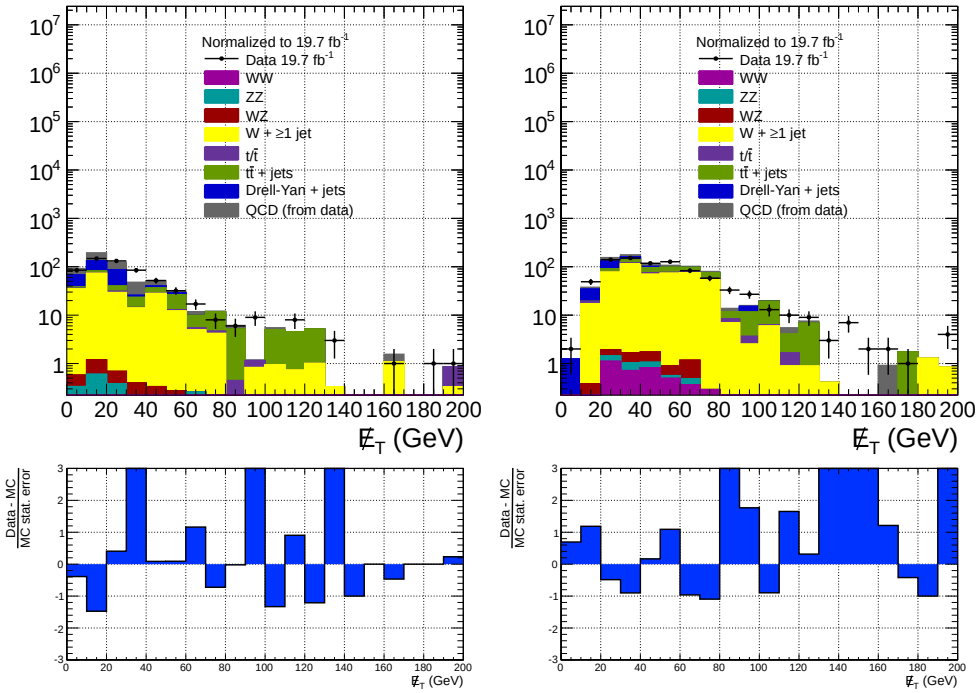

Figure 7.12. $\not\!\!E_{\mathrm{T}}$ distribution for region B data (black points), region B total non-QCD backgrounds from MC (solid stacked histograms), and the QCD prediction from region D data (solid gray histogram). Errors are statistical only. (Left) Low-M_{T} bin. (Right) High-M_{T} bin.

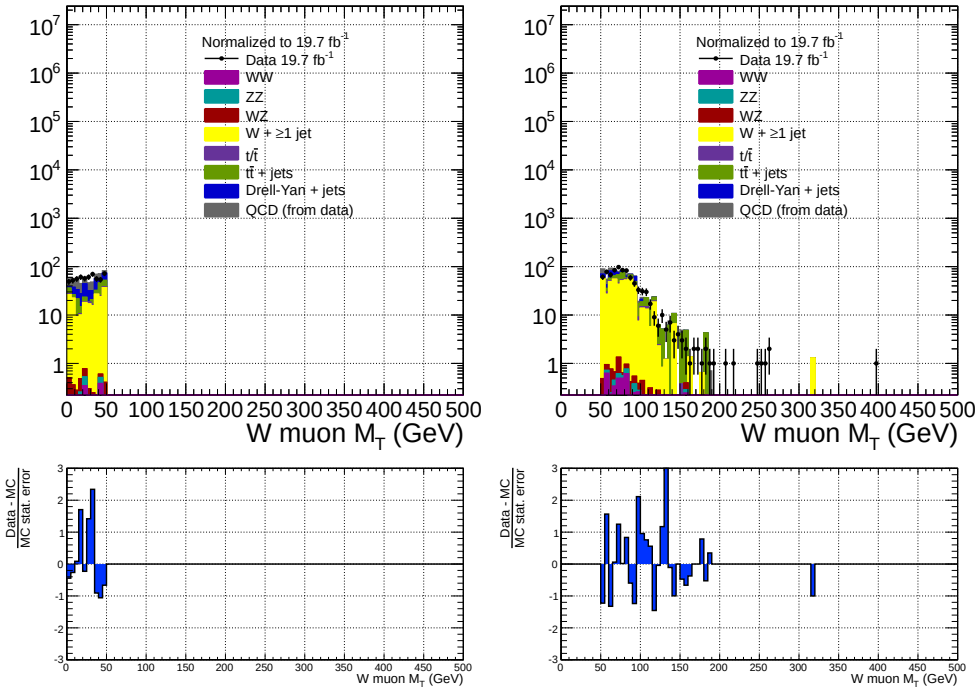

Figure 7.13. Trigger muon M_{T} distribution for region B data (black points), region B total non-QCD backgrounds from MC (solid stacked histograms), and the QCD prediction from region D data (solid gray histogram). Errors are statistical only. The term "W muon" in the label refers to the trigger muon, not necessarily a muon from a W decay (as in the case of the ggH signal, for instance). (Left) Low-M_{T} bin. (Right) High-M_{T} bin.

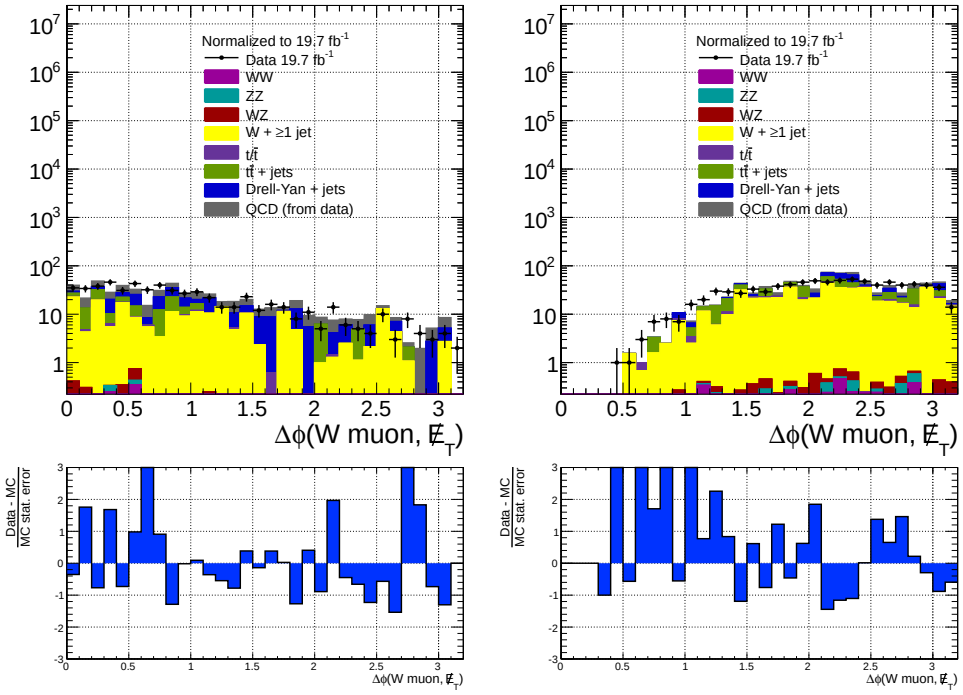

Figure 7.14. $\Delta\phi$(trigger muon, $\not\!\!E_{\mathrm{T}}$) distribution for region B data (black points), region B total non-QCD backgrounds from MC (solid stacked histograms), and the QCD prediction from region D data (solid gray histogram). Errors are statistical only. The term "W muon" in the label refers to the trigger muon, not necessarily a muon from a W decay (as in the case of the ggH signal, for instance). (Left) Low-M_{T} bin. (Right) High-M_{T} bin.

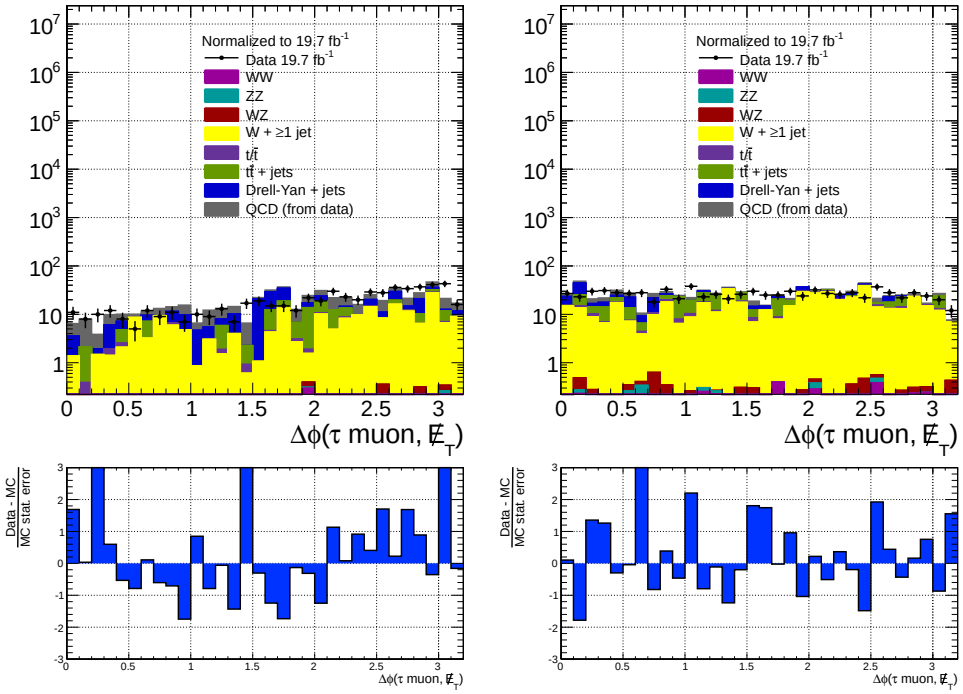

Figure 7.15. $\Delta\phi(\tau$ muon, $\not\!\!E_{\mathrm{T}})$ distribution for region B data (black points), region B total non-QCD backgrounds from MC (solid stacked histograms), and the QCD prediction from region D data (solid gray histogram). Errors are statistical only. (Left) Low-M_{T} bin. (Right) High-M_{T} bin.

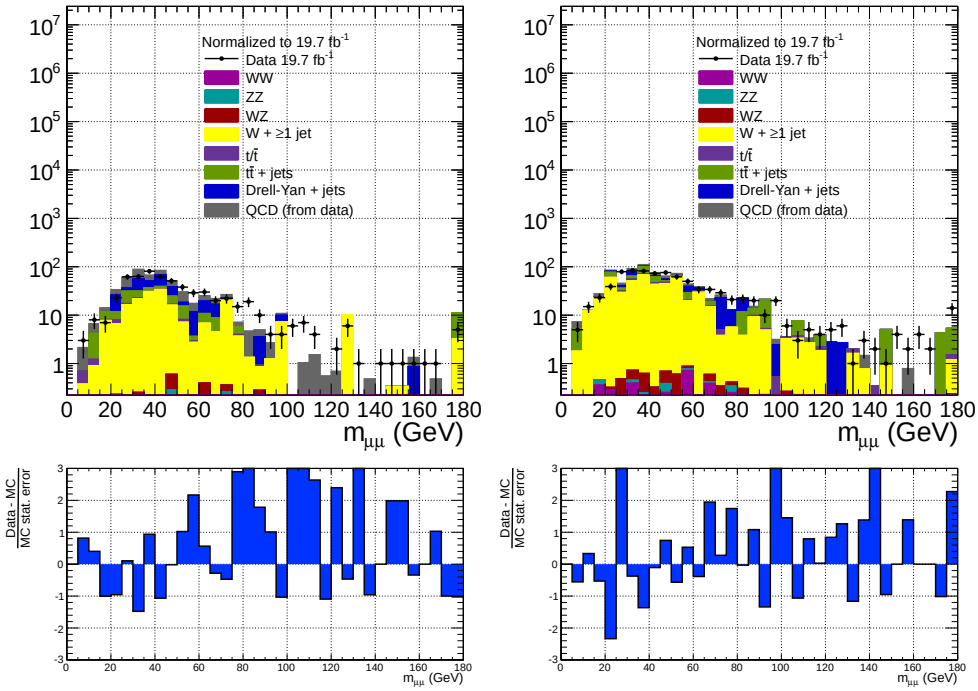

Figure 7.16. Distribution of the invariant mass of the trigger muon and τ muon for region B data (black points), region B total non-QCD backgrounds from MC (solid stacked histograms), and the QCD prediction from region D data (solid gray histogram). Errors are statistical only. (Left) Low-M_{T} bin. (Right) High-M_{T} bin.

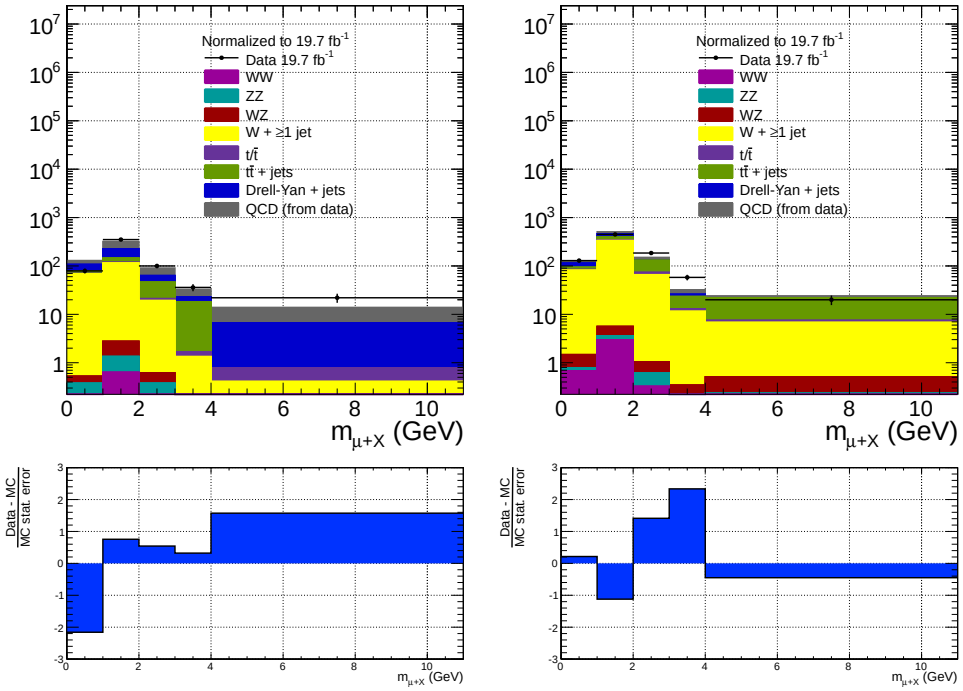

Figure 7.17. $m_{\mu+\text{had}}$ distribution for region B data (black points), region B total non-QCD backgrounds from MC (solid stacked histograms), and the QCD prediction from region D data (solid gray histogram). Errors are statistical only. (Left) Low-M_{T} bin. (Right) High-M_{T} bin.

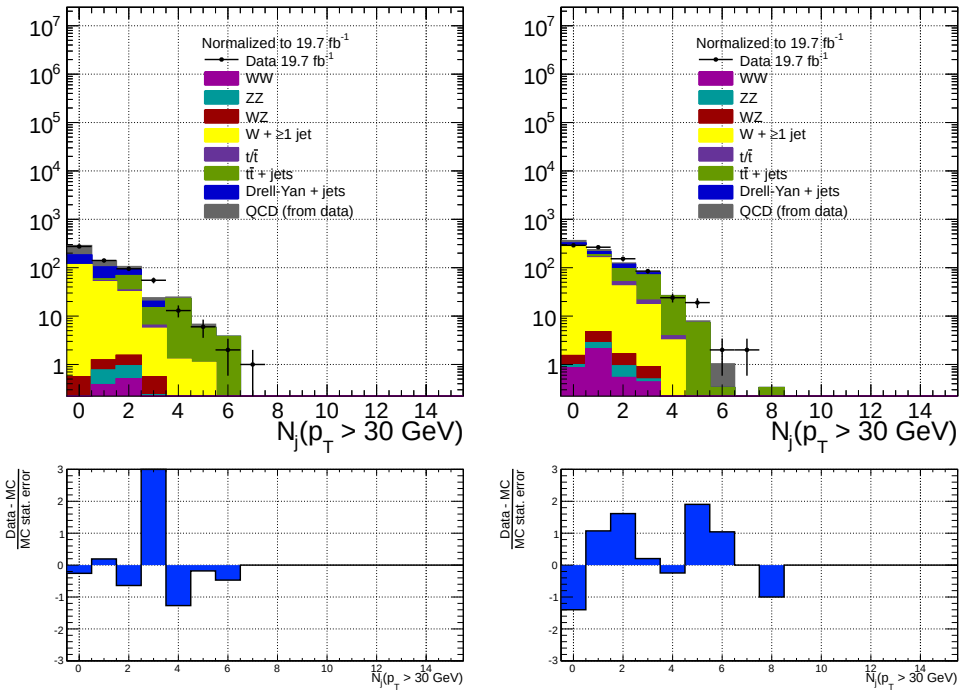

Figure 7.18. Number of anti-k_{T} R = 0.5 PF jets with L1FastL2L3 [93] corrected $p_T > 30$ GeVc. Distribution for region B data (black points), region B total non-QCD backgrounds from MC (solid stacked histograms), and the QCD prediction from region D data (solid gray histogram). Errors are statistical only. (Left) Low-M_{T} bin. (Right) High-M_{T} bin.

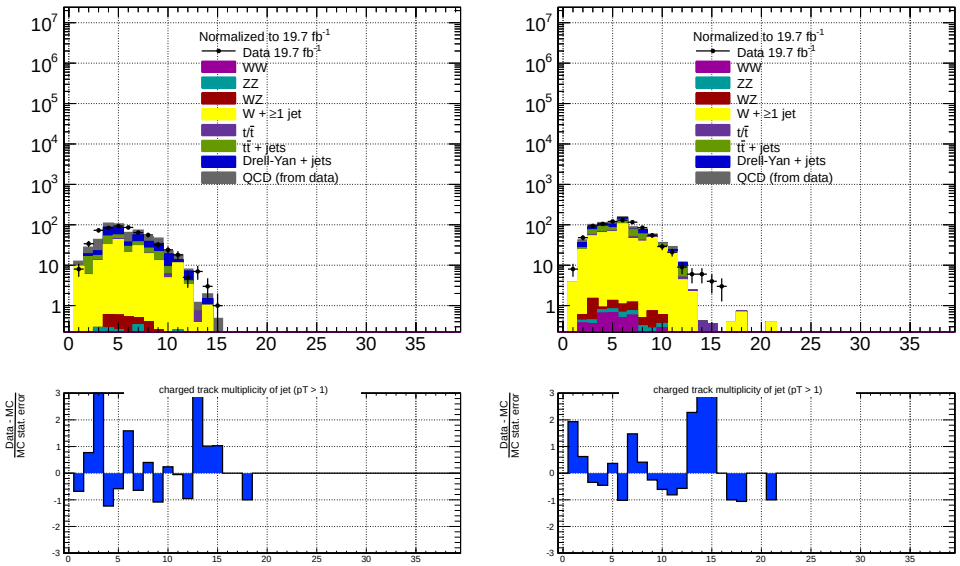

Figure 7.19. Number of charged tracks with $p_T > 1$ GeVc in the parent jet of the $\tau_\mu\tau_{\rm had}$ object. Distribution for region B data (black points), region B total non-QCD backgrounds from MC (solid stacked histograms), and the QCD prediction from region D data (solid gray histogram). Errors are statistical only. (Left) Low-$M_{\rm T}$ bin. (Right) High-$M_{\rm T}$ bin.

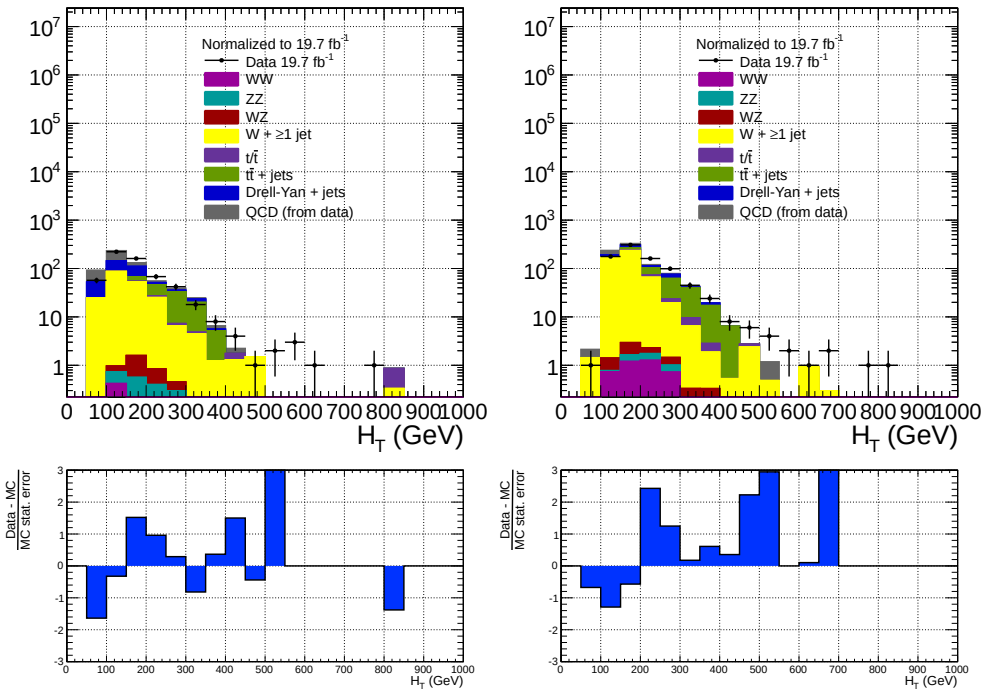

Figure 7.20. p_T sum of the tau muon, hadronic tau, highest p_T distinct jet, trigger muon, and $\not\!\!E_{\rm T}$. Distribution for region B data (black points), region B total non-QCD backgrounds from MC (solid stacked histograms), and the QCD prediction from region D data (solid gray histogram). Errors are statistical only. (Left) Low-$M_{\rm T}$ bin. (Right) High-$M_{\rm T}$ bin.

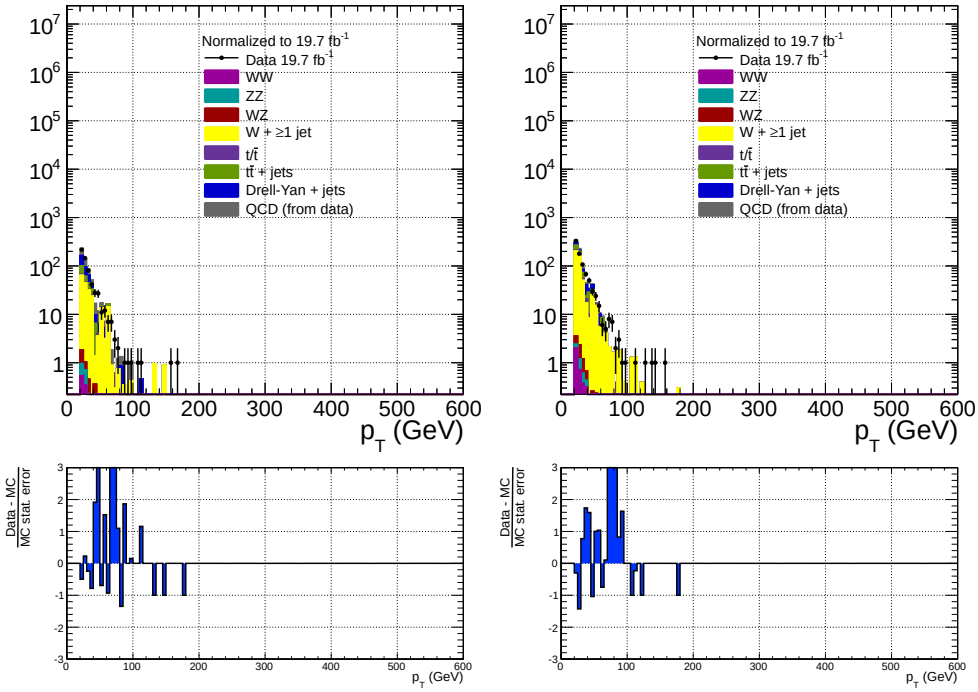

Figure 7.21. Hadronic tau p_T distribution for region B data (black points), region B total non-QCD backgrounds from MC (solid stacked histograms), and the QCD prediction from region D data (solid gray histogram). Errors are statistical only. (Left) Low-M_{T} bin. (Right) High-M_{T} bin.

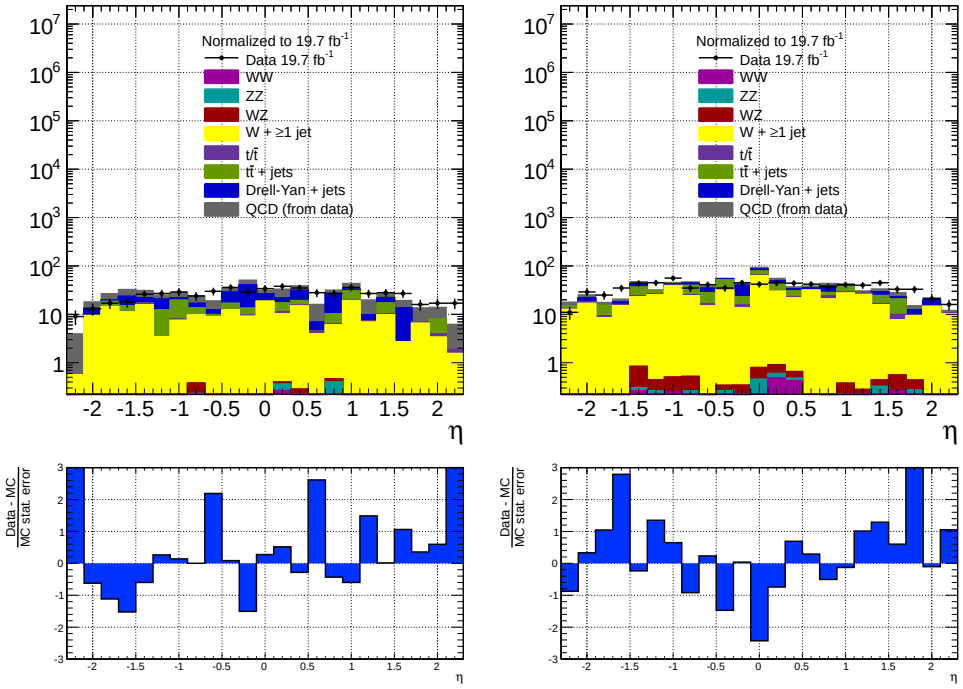

Figure 7.22. Hadronic tau η distribution for region B data (black points), region B total non-QCD backgrounds from MC (solid stacked histograms), and the QCD prediction from region D data (solid gray histogram). Errors are statistical only. (Left) Low-M_{T} bin. (Right) High-M_{T} bin.

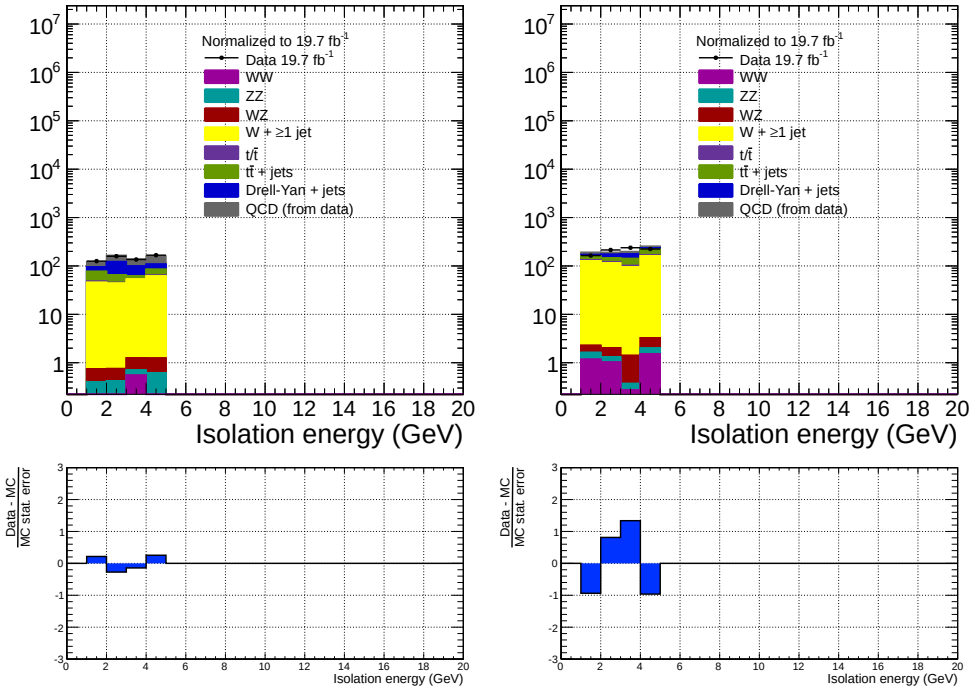

Figure 7.23. Hadronic tau isolation distribution for region B data (black points), region B total non-QCD backgrounds from MC (solid stacked histograms), and the QCD prediction from region D data (solid gray histogram). Errors are statistical only. (Left) Low-M_{T} bin. (Right) High-M_{T} bin.

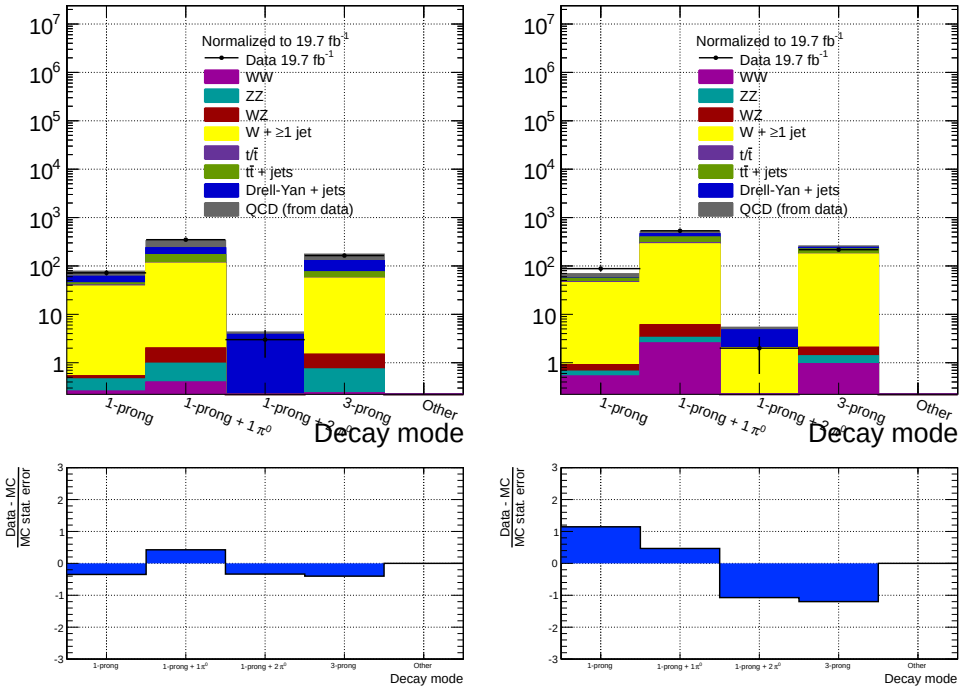

Figure 7.24. Hadronic tau decay mode distribution for region B data (black points), region B total non-QCD backgrounds from MC (solid stacked histograms), and the QCD prediction from region D data (solid gray histogram). Errors are statistical only. (Left) Low-M_{T} bin. (Right) High-M_{T} bin.

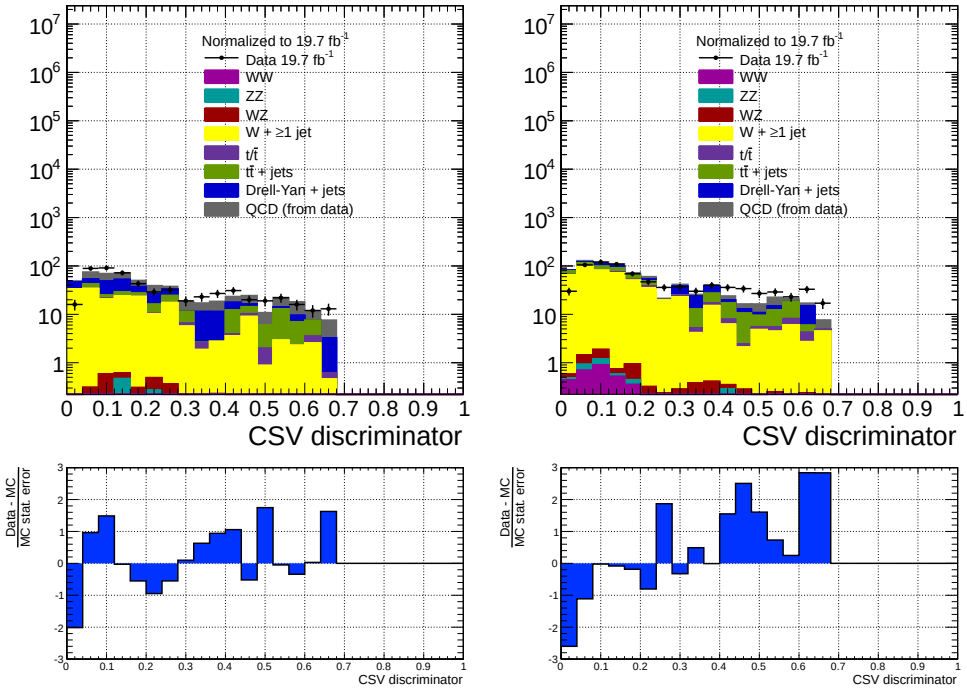

Figure 7.25. CSV discriminant distribution for region B data (black points), region B total non-QCD backgrounds from MC (solid stacked histograms), and the QCD prediction from region D data (solid gray histogram). Errors are statistical only. (Left) Low-M_{T} bin. (Right) High-M_{T} bin.

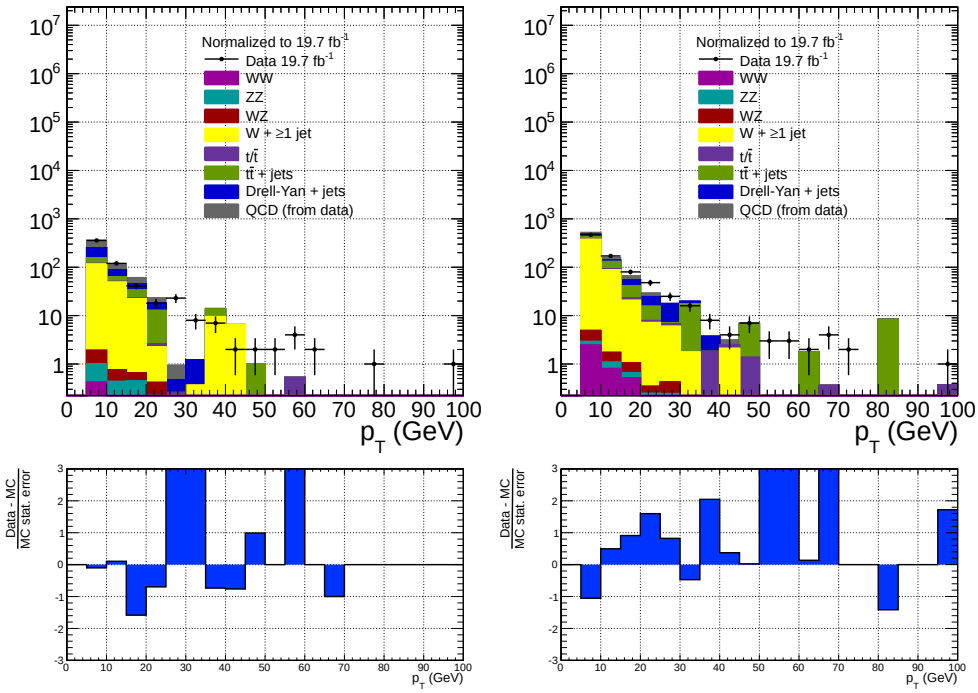

Figure 7.26. τ muon p_T distribution for region B data (black points), region B total non-QCD backgrounds from MC (solid stacked histograms), and the QCD prediction from region D data (solid gray histogram). Errors are statistical only. (Left) Low-M_{T} bin. (Right) High-M_{T} bin.

7.2.4 Final jet fake background prediction and systematic uncertainties

The region B data sample contains some mixture of Drell-Yan, W + jets, $t\bar{t}$ and single top, QCD di-jet, and diboson backgrounds. The low-M_{T} bin is dominated by QCD and Drell-Yan, while the high-M_{T} bin is dominated by W + jets and $t\bar{t}$. Within each bin, MC predicts the background composition in regions B and A to be reasonably similar, as shown in Table 7.4. Furthermore, the shape of the $m_{\mu+\mathrm{had}}$ distribution for each individual background component is predicted to be similar in regions A and B, as shown in Figs. 7.3-7.6. However, the lack of exact knowledge of the background composition of regions A and B (i.e., the relative ratios of each individual background source), coupled with small differences between the $m_{\mu+\mathrm{had}}$ distributions for each background, constitute the main systematic uncertainty of the method.

Table 7.4. Expected events below $m_{\mu+\text{had}} = 2$ GeV and above $m_{\mu+\text{had}} = 4$ GeV in Regions B and A. The MC backgrounds are normalized to 19.7 fb^{-1}. The QCD normalization is given by Eqs. 7.3 and 7.4 for region B and 7.1 and 7.2 for region A. The number in parentheses is the percentage of the total.

	$M_T \leq 50$ GeV				$M_T > 50$ GeV			
	Region A		Region B		Region A		Region B	
$m_{\mu+\text{had}}$	< 2 GeV	≥ 4 GeV	< 2 GeV	≥ 4 GeV	< 2 GeV	≥ 4 GeV	< 2 GeV	≥ 4 GeV
WW	0.315 ± 0.18 (0.3%)	0 ± 0 (0%)	0.726 ± 0.3 (0.16%)	0 ± 0 (0%)	0.472 ± 0.24 (0.38%)	0 ± 0 (0%)	3.66 ± 0.68 (0.58%)	0 ± 0 (0%)
ZZ	0.262 ± 0.096 (0.25%)	0.108 ± 0.062 (1%)	0.998 ± 0.2 (0.22%)	0.0421 ± 0.042 (0.3%)	0.263 ± 0.1 (0.21%)	0.145 ± 0.085 (4.4%)	0.723 ± 0.17 (0.11%)	0.235 ± 0.097 (0.97%)
WZ	0.344 ± 0.15 (0.33%)	0 ± 0 (0%)	1.58 ± 0.35 (0.35%)	0 ± 0 (0%)	1.08 ± 0.28 (0.86%)	0.404 ± 0.17 (12%)	2.71 ± 0.46 (0.43%)	0.276 ± 0.14 (1.1%)
W + jets	41.1 ± 10 (39%)	0 ± 0 (0%)	180 ± 26 (40%)	0.367 ± 0.37 (2.6%)	74.6 ± 14 (60%)	1.76 ± 1.5 (54%)	409 ± 39 (65%)	6.29 ± 2.8 (26%)
Single top	1.56 ± 0.8 (1.5%)	0.278 ± 0.28 (2.6%)	4.47 ± 1.3 (0.99%)	0.373 ± 0.32 (2.7%)	2.5 ± 0.86 (2%)	0 ± 0 (0%)	15.5 ± 2.4 (2.4%)	0.69 ± 0.43 (2.9%)
$t\bar{t}$	9.76 ± 6.1 (9.3%)	4.1 ± 4.1 (39%)	35.4 ± 11 (7.9%)	0 ± 0 (0%)	17.7 ± 8 (14%)	0 ± 0 (0%)	59.9 ± 15 (9.5%)	14.6 ± 8.6 (60%)
Drell-Yan + jets	26.8 ± 11 (25%)	5.25 ± 3.7 (49%)	110 ± 22 (24%)	5.92 ± 4.3 (42%)	11.7 ± 5.7 (9.4%)	0.463 ± 0.46 (14%)	85.2 ± 18 (13%)	0 ± 0 (0%)
QCD (from data)	24.9 ± 7 (24%)	0.89 ± 0.7 (8.4%)	117 ± 30 (26%)	7.23 ± 2.8 (52%)	16.7 ± 15 (13%)	0.49 ± 0.74 (15%)	56.5 ± 36 (8.9%)	2.09 ± 2 (8.7%)
Tot. expected	105 ± 17	10.6 ± 5.6	450 ± 46	13.9 ± 5.1	125 ± 23	3.26 ± 1.7	633 ± 59	24.1 ± 9.2

Table 7.5. Background prediction for $m_{\mu+\text{had}} \geq 4$ GeV after full selection, from Region B data and from the alternative background shapes (Region D data and Region B total EWK MC).

Background shape source	Low M_T	High M_T
Region B data	6.24 ± 1.48	5.75 ± 1.38
Region D data ("all-QCD")	7.53 ± 2.12	6.15 ± 3.65
Region B total MC ("all-EWK")	2.45 ± 1.61	6.35 ± 2.69

To assess this uncertainty, the background is estimated using two alternate $m_{\mu+\text{had}}$ shapes. One alternate shape is taken from region D data, and represents the "all-QCD" shape. The other is from the sum of the Drell-Yan, W + jets, $t\bar{t}$ and single top, and diboson MC in region B and represents the "all-electroweak (EW)" shape. The number of events in the $m_{\mu+\text{had}} \geq 4$ GeV bin in Region B data after the full selection is then compared to the yield predicted by these alternative shapes, which are normalized such that the integral of events in the normalization sideband is the same as the integral for the region B data shape; the comparison of these yields is shown in Table 7.5. The differences between the two alternates and the nominal are taken as an asymmetric error on the jet fake background prediction from region B. As the region D data and region B MC are themselves statistically limited, the statistical error on the background estimations from these two samples must also factor into the total error on the nominal background prediction.

The jet fake background prediction in the $m_{\mu+\text{had}} > 4$ GeV bin is taken to be the unweighted average of the Region B data yield and the yields from the other two alternative shapes in that bin. The asymmetric systematic uncertainties are then calculated by taking the difference between the average yield and the alternate shape yield whose central value plus $1 - \sigma$ is the farthest from the average yield (in the positive or negative direction, for the respective positive and negative systematic errors). The nominal, all-QCD, and all-EW $m_{\mu+\text{had}}$ distributions, properly normalized to the region A data, are shown in Figure 7.27.

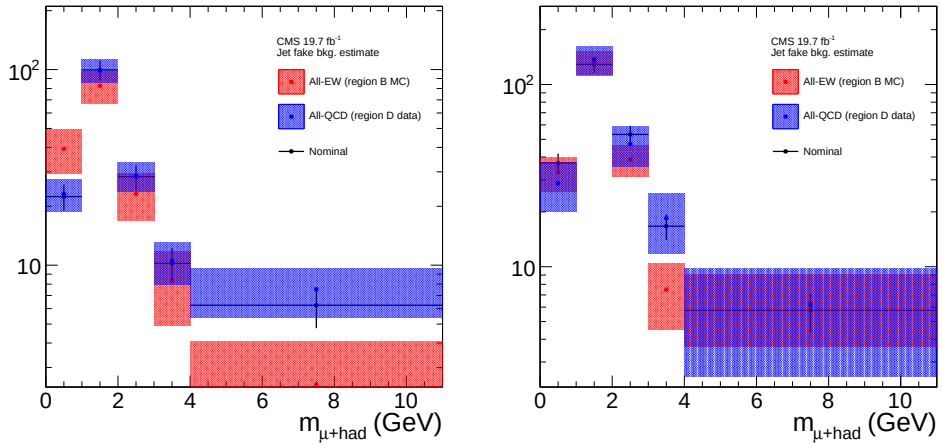

Figure 7.27. Alternative shapes for the jet fake background: nominal (black points), all-QCD (blue shaded band), and all-EW (red shaded band) $m_{\mu+\text{had}}$ distributions, properly normalized to the region A data. The statistical errors on each distribution are shown as shaded bands or y error bars. (Left) Low-M_{T} bin. (Right) High-M_{T} bin.

Including all statistical errors and the envelope formed by the alternative background shape templates, the final jet fake background predictions are

- Low-M_{T}: 5.41 $\pm$ 1 (stat.) +4.2/-4.6 (syst.)
- High-M_{T}: 6.08 $\pm$ 1.6 (stat.) +3.7/-3.6 (syst.)

7.3 Cross-checks for additional backgrounds

An isolated $\tau_\mu\tau_{\text{had}}$ pair can be faked by the collimated decay of a boosted di-muon resonance, where one muon is reconstructed as the HPS tau, or by two nearby hadrons decaying semileptonically. In the signal window $m_{\mu+\text{had}} \geq 4$ GeV, there is the additional possibility of Υ decays to $\tau_\mu\tau_{\text{had}}$ pairs. The cross section for double Υ [101] and $W+\Upsilon$ [102] for the kinematic cuts employed in this search is expected to be negligibly small at the 8 TeV LHC. However, it has to be confirmed that the background from double semileptonic decays is also small.

7.3.1 Cross check #1: fit for resonances in region C

In this check, data control region C is used to estimate this background. Region C (cf. Fig. 7.2) is composed of events passing all preselections except trigger muon isolation,

which they are required to fail. This trigger muon isolation inversion insures that region C is enriched in di-jet events. If boosted prompt Υ production is a background at all, it should be more prevalent in region C than in any other control region due to the tau isolation requirement. Non-prompt Υ from b decay should be somewhat suppressed due to the CSV veto for the jet that seeds the $\tau_\mu\tau_{\rm had}$ (cf. Sec. 5.5) and the tau isolation requirement.

Region C is composed of two types of events: di-jet events where one jet fakes the $\tau_\mu\tau_{\rm had}$ object, and boosted $\Upsilon \to \tau_\mu\tau_{\rm had}$ events. Jet-faking-tau events are like those of region B, with a smoothly falling $m_{\mu+\rm had}$ distribution due to the lack of a real massive resonance. Boosted $\Upsilon \to \tau_\mu\tau_{\rm had}$ events should have a broad peaking structure near 10 GeV, but due to neutrinos in the decay and the lack of simulation, there is considerable uncertainty as to what the actual shape would be. Therefore, we conservatively take all events in region C with $m_{\mu+\rm had} \geq 4$ GeV to be Υ events and normalize by the factor $R_{\rm A}^{\rm QCD}$ (cf. Eq. 7.2).

The choice of $m_{\mu+\rm had} < 3$ GeV insures adequate statistics for the calculation of the normalization factor, such that statistical error from region C, not normalization error, dominates the Υ background error. Furthermore, we assume that the rate of boosted Υ production (more pronounced at high $m_{\mu+\rm had}$) relative to all other backgrounds (more pronounced at low $m_{\mu+\rm had}$) is independent of the recoiling trigger muon isolation, implying that the normalization factor should be independent of the mass range in which it is calculated. (We do not make this assumption in cross check #2 in Section 7.3.2.)

The upsilon background predictions are

- Low-$M_{\rm T}$: 0.888 ± 0.664 (stat.)
- High-$M_{\rm T}$: 0.483 ± 0.737 (stat.)

Although it does not affect the signal region $m_{\mu+\rm had} \geq 4$ GeV, the background from $J/\psi \to \mu\mu$ where one muon fakes a tau is also estimated from region C data. This background affects the mass range 2 GeV $\leq m_{\mu+\rm had} < 4$ GeV that is used in the final unblinding as a cross check of the background prediction methods. The number of J/ψ

events is extracted from a one-dimensional binned likelihood fit to the $m_{\mu+\text{had}}$ distribution in region C. The fit function is a signal (i.e. J/ψ) Crystal Ball + background (i.e. non-peaking QCD from jets faking taus) exponential, and is restricted to the range $1.5 \leq m_{\mu+\text{had}} \leq 4$ GeV, since it is assumed that all region C data with $m_{\mu+\text{had}} \geq 4$ GeV are due to $\Upsilon \to \tau\tau$. The fitted Crystal Ball component is normalized with the same factor $R_{\text{A}}^{\text{QCD}}$ used for the Υ background determination.

Based on a comparison of $m_{\mu+\text{had}}$ distributions between regions C and D, the J/ψ background has been found to be relevant only for the high-M_{T} bin. If the background from boosted $J/\psi \to \mu\mu$ where one muon fakes a tau is significant, it should be more prevalent in region C than region D due to the $\tau_\mu\tau_{\text{had}}$ isolation requirement in region C. As shown in Figure 7.28 (top), there is a vague peaking structure around the J/ψ mass in the region C distribution with respect to the region D distribution for the high-M_{T} bin only. A cross-check with more statistics is shown in Figure 7.28 (bottom) for data triggered by `HLT_Mu40_eta2p1` and with trigger muon p_T threshold increased to 41 GeV. This trigger has no muon isolation requirement, so the trigger bias present in the nominal `HLT_IsoMu24_eta2p1`-triggered regions C and D (which are required to fail the trigger muon offline PF isolation requirement, cf. Fig. 7.2) is absent in the `HLT_Mu40_eta2p1`-triggered regions C and D. In this cross-check region with no trigger muon isolation, the region C peaking structure in the high-M_{T} bin is more pronounced, but no such pronouncement is observed in the low-M_{T} bin.

Two different methods are used to extract the exponential decay constant. The first is a fit of the Crystal Ball + exponential shape to the region C data with all fit parameters floating, where the decay constant is taken as the fitted exponential decay constant parameter. The second is a fit of the exponential shape only to the data in region D within the mass range $1.5 \leq m_{\mu+\text{had}} \leq 4$ GeV. The fit in region D is more precise due to the larger number of events than in region C, but the fit in region C utilizes the true sidebands (from jets faking taus) in region C and so is more accurate. The exponential background shape with decay constant taken as the weighted average of that found from the two fits is taken as nominal, and the J/ψ component is extracted from a Crystal Ball

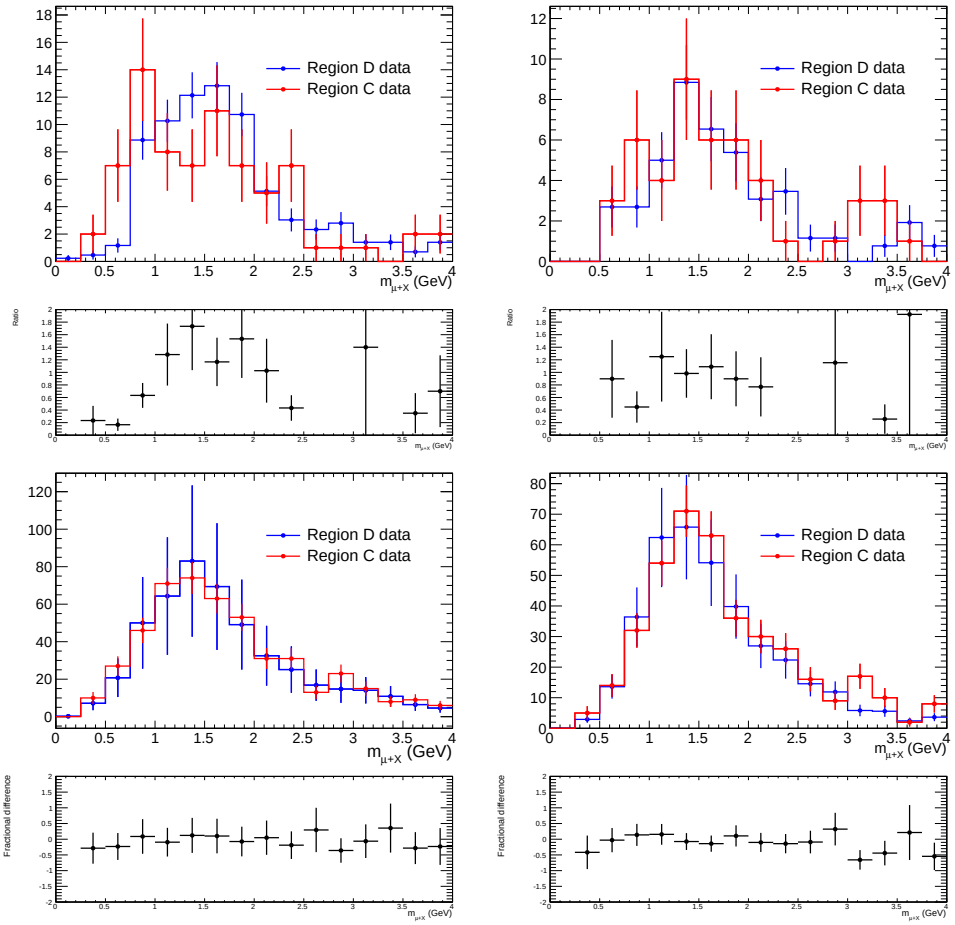

Figure 7.28. $m_{\mu+\text{had}}$ distributions for regions C and D. The region D distribution is normalized such that $N_{\text{C}}(m_{\mu+\text{had}} < 3\ GeV) = N_{\text{D}}(m_{\mu+\text{had}} < 3\ GeV)$. (Top Left) Low M_{T}, `HLT_IsoMu24_eta2p1`. (Top Right) High M_{T}, `HLT_IsoMu24_eta2p1`. (Bottom Left) Low M_{T}, `HLT_Mu40_eta2p1` (tau $p_T > 10$ GeV). (Bottom Right) High M_{T}, `HLT_Mu40_eta2p1` (tau $p_T > 10$ GeV).

signal fit with the background shape fixed to nominal.

Region C data with the three different fit functions overlaid, as well as region D data with the background fit function overlaid, are shown in Figure 7.29. The fit quality is affected by the poor statistics in region C, and the signal shape and normalization is correlated to the choice of background shape (compare Fig. 7.29 top Rightand bottom Left). These are indications that the $J/\psi \to \mu\mu$ background is likely very small. The same fit study was performed for the `HLT_Mu40_eta2p1` cross-check region to make sure that the J/ψ background is not being underestimated due to trigger bias. The `HLT_Mu40_eta2p1` fit results are shown in Figure 7.30. The fit results are a little more precise than for the `HLT_IsoMu24_eta2p1`-triggered data, but the fitted signal is similarly small.

Table 7.6 compares the predicted J/ψ, Υ, and jet fake background yields from the nominal
(`HLT_IsoMu24_eta2p1`) and cross-check (`HLT_Mu40_eta2p1`) data samples. The ratios of Υ to jet fake and J/ψ to jet fake backgrounds are stable across the two triggers, indicating no serious underestimation of the resonance backgrounds in the nominal `HLT_IsoMu24_eta2p1` data due to trigger bias. If anything, the J/ψ background is probably overestimated in the `HLT_IsoMu24_eta2p1` data due to a poor fit with low statistics.

The difference in signal yield as a function of fixed background shape (decay constant from region D fit or region C sideband fit) is taken as a systematic error of the fit method. A plot comparing the nominal J/ψ background shape to the band formed by the varied shapes is shown in Figure 7.31. Note that the J/ψ yield is zero in the signal region $m_{\mu+\mathrm{had}}$ $\geq$ 4 GeV—the J/ψ background estimation is done for $m_{\mu+\mathrm{had}} < 4$ GeV to get a handle on the sensitivity of this search to boosted di-lepton resonances.

From Table 7.6, we see that the predicted J/ψ yield in the high M_{T} bin of region A is negligible compared to the jet fake background prediction in the same mass window. The conservative Υ prediction is at most 20% of the jet fake background prediction for the signal window $m_{\mu+\mathrm{had}} \geq 4$ GeV, and the low significance of the J/ψ peak indicates that the search is probably insensitive to $\Upsilon \to \mu\mu$ or $\Upsilon \to \tau\tau$. In addition, no J/ψ peak is found in the low M_{T} bin of region C, where it might be expected to be most prominent.

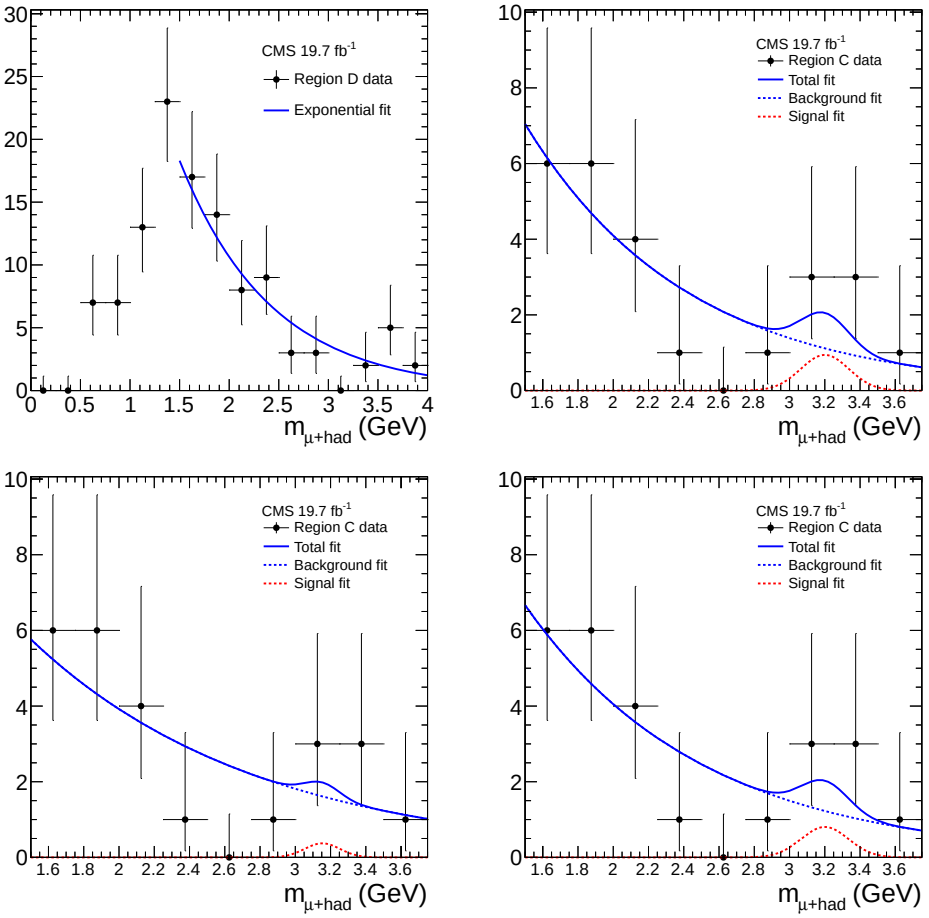

Figure 7.29. Fits of the $m_{\mu+\mathrm{had}}$ distribution in regions C and D to extract the J/ψ component. Data are shown for the high M_{T} bin and `HLT_IsoMu24_eta2p1` trigger. (Top Left) Exponential-only fit to region D. (Top Right) Crystal Ball + exponential fit to region C with exponential decay constant fixed to value fitted in region D. (Bottom Left) Crystal Ball + exponential fit to region C with all parameters floating. (Bottom Right) Crystal Ball + exponential fit to region C with exponential decay constant fixed to weighted average of values found in region D exponential-only and region C all-parameters-floating fits.

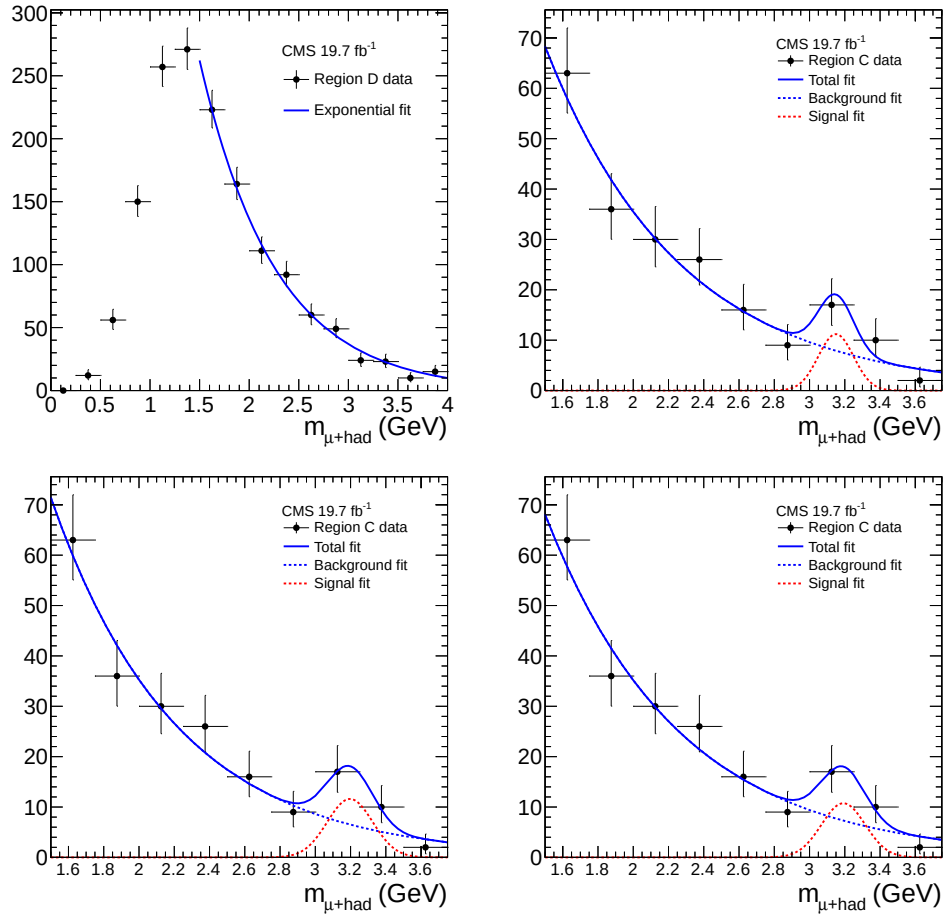

Figure 7.30. Fits of the $m_{\mu+\mathrm{had}}$ distribution in regions C and D to extract the J/ψ component. Data are shown for the high M_{T} bin and `HLT_Mu40_eta2p1` trigger. Tau $p_T > 10$ GeV. (Top Left) Exponential-only fit to region D. (Top Right) Crystal Ball + exponential fit to region C with exponential decay constant fixed to value fitted in region D. (Bottom Left) Crystal Ball + exponential fit to region C with all parameters floating. (Bottom Right) Crystal Ball + exponential fit to region C with exponential decay constant fixed to weighted average of values found in region D exponential-only and region C all-parameters-floating fits.

Table 7.6. Predicted J/ψ, Υ, and jet fake background yields from the nominal (HLT_IsoMu24_eta2p1) and cross-check (HLT_Mu40_eta2p1) data samples. For the HLT_Mu40_eta2p1 data, tau $p_T > 10$ GeV. Only statistical errors are quoted.

	HLT_IsoMu24_eta2p1 (nominal)		HLT_Mu40_eta2p1 (cross-check)	
	$M_T \leq 50$ GeV	$M_T > 50$ GeV	$M_T \leq 50$ GeV	$M_T > 50$ GeV
Bkg. from reg. B (j $\to \tau$) ($m_{\mu+\text{had}} \geq 4$ GeV)	6.2 ± 1.4	6.0 ± 1.4	6.5 ± 1.3	7.4 ± 1.4
Bkg. from reg. B (j $\to \tau$) (2 GeV $\leq m_{\mu+\text{had}} < 4$ GeV)	—	38 ± 14	—	72 ± 6
Bkg. from reg. C (Υ)	0.9 ± 0.7	0.48 ± 0.74	0.7 ± 0.3	1.1 ± 0.5
Bkg. from reg. C (J/ψ)	—	0.9 ± 1.0	—	2.1 ± 0.8
Ratio Υ : jet fake ($m_{\mu+\text{had}} \geq 4$ GeV)	0.14 ± 0.11	0.08 ± 0.13	0.10 ± 0.04	0.15 ± 0.07
Ratio J/ψ : jet fake (2 GeV $\leq m_{\mu+\text{had}} < 4$ GeV)	—	0.012 ± 0.014	—	0.005 ± 0.002

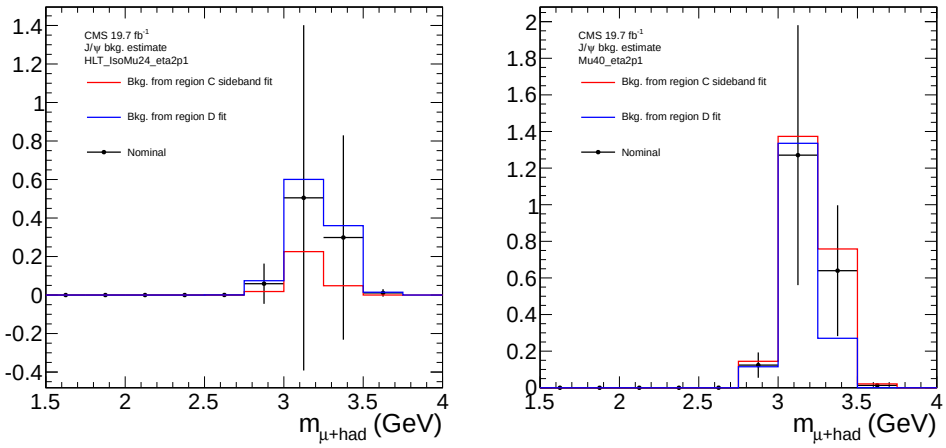

Figure 7.31. Nominal and varied J/ψ background shapes for high-M_T data. Tau $p_T > 10$ GeV. (Left) HLT_IsoMu24_eta2p1. (Right) HLT_Mu40_eta2p1.

To address these issues, another cross check of the background from double muon decays is presented in the following section.

7.3.2 Cross check #2: three-muon events

Di-muon resonances or double semileptonic decays faking the $\tau_\mu \tau_{\text{had}}$ object are expected to show up in a subsample of the preselected data with three muons: one trigger muon,

one tau decay muon, and a muon reconstructed from the same track as a signal candidate of the HPS tau. As in Sec. 7.3.1, three-muon events are expected to be visible in region C more than in region D due to the isolated $\tau_\mu\tau_{\rm had}$ object, but they may not scale to region A via the same factor $R_{\rm A}^{\rm QCD}$.

In this cross check, the data sample is split into exclusive 3-muon and non-3-muon categories. The estimated $\tau_\mu\tau_{\rm had}$ mass distribution for the background from jets with double muon decays, denoted $f_{\rm A}^{3\mu}(m_{\mu+\rm had})$, is given by

$$f_{\rm A}^{3\mu}(m_{\mu+\rm had}) = R_{\rm A}^{3\mu} \cdot f_{\rm C}^{3\mu}(m_{\mu+\rm had}) \tag{7.5}$$

where

$$R_{\rm A}^{3\mu} = \int_0^\infty f_{\rm B}^{3\mu}(m_{\mu+\rm had})dm_{\mu+\rm had} / \int_0^\infty f_{\rm D}^{3\mu}(m_{\mu+\rm had})dm_{\mu+\rm had} \tag{7.6}$$

and $f_{\rm i}^{3\mu}(m_{\mu+\rm had})$ is the $\tau_\mu\tau_{\rm had}$ mass distribution for three-muon events in region i, i = B,C,D. The estimated $\tau_\mu\tau_{\rm had}$ mass distribution for the background from jets with single muon decays, denoted $f_{\rm A}^{\rm fake}(m_{\mu+\rm had})$, is given by

$$f_{\rm A}^{\rm fake}(m_{\mu+\rm had}) = R_{\rm A}^{\text{non-}3\mu} \cdot f_{\rm C}^{\text{non-}3\mu}(m_{\mu+\rm had}) \tag{7.7}$$

where

$$R_{\rm A}^{\text{non-}3\mu} = \int_0^{2GeV} f_{\rm A}^{\text{non-}3\mu}(m_{\mu+\rm had})dm_{\mu+\rm had} / \int_0^{2GeV} f_{\rm B}^{\text{non-}3\mu}(m_{\mu+\rm had})dm_{\mu+\rm had} \tag{7.8}$$

and $f_{\rm i}^{\text{non-}3\mu}(m_{\mu+\rm had})$ is the $\tau_\mu\tau_{\rm had}$ mass distribution for non-three-muon events in region i, i = B,C,D. In this scheme, the single muon jet fake background estimated from region B is allowed to have a different scale factor to region A than the double muon background estimated from region C.

An event is classified as a 3-muon event if any of the HPS tau signal candidates from the 20 GeV reconstructed HPS tau that's part of the $\tau_\mu\tau_{\rm had}$ object shares a `TrackRef` with the best `TrackRef` (i.e. `reco::Muon::muonBestTrack()`) of any PF `reco::Muon` in the event (i.e. `reco::Muon::isPFMuon()`) that is not already identified as the trigger muon

(e.g. the highest p_T muon with $p_T > 25$ GeV satisfying the trigger muon criteria) or the partner tau decay muon (e.g. the highest p_T soft muon with $p_T > 5$ GeV from among the muons "removed" from the jet that feeds the HPS reconstruction).

As shown in Figure 7.32, the 3-muon control sample in region C is quite small, even nonexistent in the low M_{T} bin. Therefore, the results are similar to what is obtained without treating the 3-muon and non-3-muon samples separately. There are no 3-muon events in the $m_{\mu+\mathrm{had}} \geq 4$ GeV signal window of region C, indicating a negligible background from jets with double muon decays.

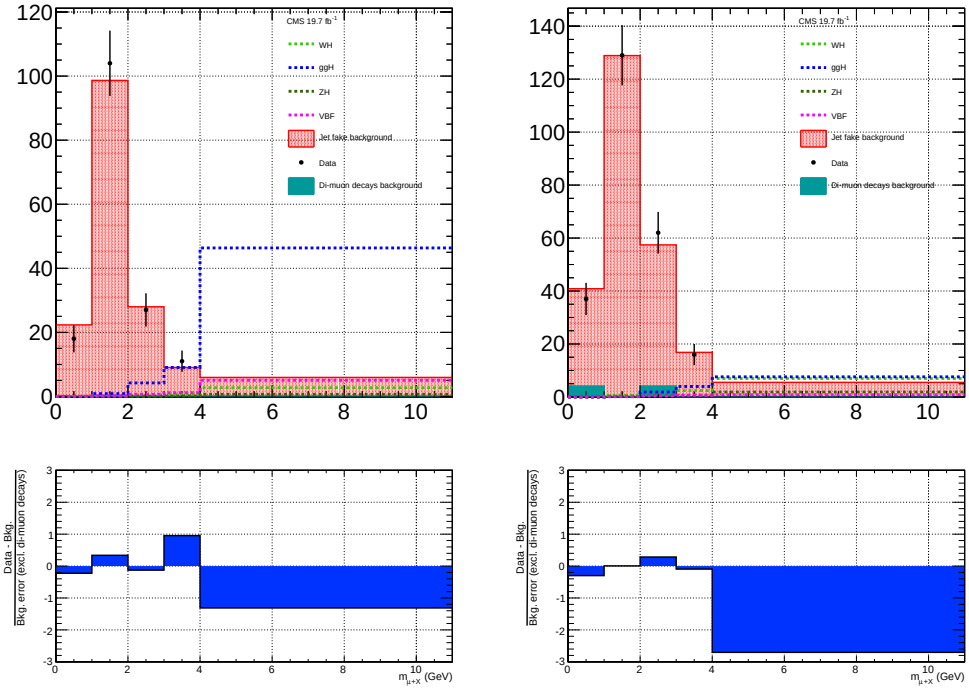

Figure 7.32. $m_{\mu+\mathrm{had}}$ distribution for non-three-muon region B data (red), three-muon region C data (teal), and all region A data (black). The non-three-muon and three-muon background predictions are normalized as described in the text. (Left) Low-M_{T} bin. (Right) High-M_{T} bin.

Looser definitions of the three-muon sample also do not predict any events in the $m_{\mu+\mathrm{had}} \geq 4$ GeV signal window of region C. These studies are documented in Ref. [103]. We conclude that jets from double muon decays and boosted di-muon resonances are a negligible background to this search.

7.4 Total background

Figure 7.33 shows the jet fake background estimate, the data, and the four signal models for the nominal `HLT_IsoMu24_eta2p1`-triggered sample. Table 7.7 shows the breakdown of the expected signal and background contributions above $m_{\mu+\text{had}} = 4$ GeV in Region A, including statistical errors.

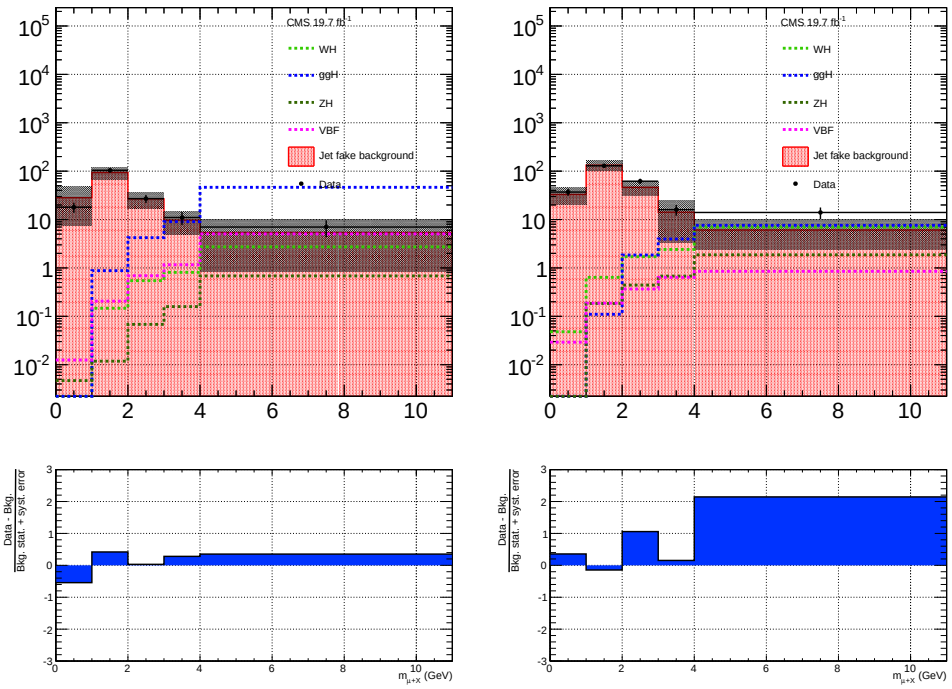

Figure 7.33. Jet fake background estimate, data, and four signal models. The different-colored pull distributions beneath the plots are evaluated for different choices of background and background statistical error according to the systematic jet fake background shape variations described in Sec. 7.2.4. The pulls are meaningless above 4 GeV because the data is blinded there. (Left) Low M_{T}. (Right) High M_{T}.

Table 7.7. Expected signal and background events and observed events above $m_{\mu+\text{had}}$ = 4 GeV in Region A. All errors are statistical only except the region B data error, which includes all systematic contributions discussed in Sec. 7.2.4.

	$M_T \leq 50$ GeV	$M_T > 50$ GeV
WW	0	0
ZZ	0.108 ± 0.062	0.145 ± 0.085
WZ	0	0.404 ± 0.17
W + jets	0	1.76 ± 1.5
Single top	0.278 ± 0.28	0
$t\bar{t}$	4.1 ± 4.1	0
Drell-Yan + jets	5.25 ± 3.7	0.463 ± 0.46
QCD (from data)	0.89 ± 0.7	0.49 ± 0.74
Pred. bkg. (region B data)	5.41 ± 1 (stat.) +4.2/-4.6 (syst.)	6.08 ± 1.6 (stat.) +3.7/-3.6 (syst.)
WH	2.72 ± 0.22	6.97 ± 0.35
ggH	46.4 ± 2.8	7.6 ± 1.1
ZH	0.683 ± 0.046	1.87 ± 0.077
VBF	5.05 ± 0.28	0.851 ± 0.11

Chapter 8

Results and interpretation

This chapter presents the final results of the search after the signal region is unblinded, followed by the interpretation of these results in terms of a model-independent limit on the branching ratio Br($H \rightarrow aa \rightarrow 4\tau$).

8.1 Observed and expected results

Table 8.1 shows the expected numbers of signal events from each generated pseudoscalar mass point in each of the M_{T} bins, followed by the background prediction (obtained from Region B data) and the actual number of events observed in Region A data for $m_{\mu+\mathrm{had}} > 4$ GeV.

As can be seen from Table 8.1, after the full selection, the observed number of events in the search region $m_{\mu+\mathrm{had}} > 4$ GeV matches the predicted $m_{\mu+\mathrm{had}}$ background (obtained as described in Sec. 7.2.4, and normalized according to Sec. 5.8) within the allowed statistical and systematic errors. 14 events were observed in the high-M_T bin, which is in excess of the background prediction by about 8 events, but this excess is still within 2σ of the background prediction, where σ is the combined statistical and systematic uncertainty for that bin. Thus, for both M_T regions, no significant excess is observed above the Standard Model prediction.

Table 8.1. Observed data, estimated background, and expected signal from each generated pseudoscalar mass point in each of the M_{T} bins assuming SM cross sections and 100% Br($H \rightarrow aa \rightarrow 4\tau$) Only statistical error is shown for the signal, while the full error is shown for the background.

		$M_{\mathrm{T}} \leq 50$ GeV	$M_{\mathrm{T}} > 50$ GeV
WH	$m_a = 5$ GeV	0.11 ± 0.05	0.10 ± 0.04
	$m_a = 7$ GeV	1.5 ± 0.2	3.8 ± 0.3
	$m_a = 9$ GeV	2.7 ± 0.2	7.0 ± 0.3
	$m_a = 11$ GeV	4.2 ± 0.3	8.8 ± 0.4
	$m_a = 13$ GeV	3.5 ± 0.2	9.9 ± 0.4
	$m_a = 15$ GeV	3.6 ± 0.2	8.4 ± 0.4
ggH	$m_a = 5$ GeV	0.31 ± 0.22	0
	$m_a = 7$ GeV	21 ± 2	1.9 ± 0.6
	$m_a = 9$ GeV	46 ± 3	7.6 ± 1.1
	$m_a = 11$ GeV	64 ± 3	11 ± 1
	$m_a = 13$ GeV	63 ± 3	18 ± 2
	$m_a = 15$ GeV	41 ± 3	11 ± 1
ZH	$m_a = 5$ GeV	0.03 ± 0.01	0.03 ± 0.01
	$m_a = 7$ GeV	0.38 ± 0.04	1.0 ± 0.1
	$m_a = 9$ GeV	0.68 ± 0.05	1.9 ± 0.1
	$m_a = 11$ GeV	1.1 ± 0.05	2.3 ± 0.1
	$m_a = 13$ GeV	0.88 ± 0.06	2.7 ± 0.1
	$m_a = 15$ GeV	0.91 ± 0.06	2.3 ± 0.1
VBF	$m_a = 5$ GeV	0.03 ± 0.02	0
	$m_a = 7$ GeV	2.3 ± 0.2	0.22 ± 0.06
	$m_a = 9$ GeV	5.1 ± 0.3	0.9 ± 0.1
	$m_a = 11$ GeV	7.0 ± 0.4	1.2 ± 0.1
	$m_a = 13$ GeV	6.9 ± 0.4	2.0 ± 0.2
	$m_a = 15$ GeV	4.5 ± 0.3	1.3 ± 0.2
SM Background		5.41 ± 1 (stat.) $^{+4.2}_{-4.6}$ (syst.)	6.08 ± 1.6 (stat.) $^{+3.7}_{-3.6}$ (syst.)
Data (observed)		7	14

8.2 Limit calculation with the CL_s method

Various statistical techniques exist for assessing the compatibility of observed data with the background-only (or null) hypothesis and the signal + background hypothesis. With these techniques, one can set an upper limit on the parameter under study, such as the signal process cross section, beyond which one can exclude the signal hypothesis with a

desired level of confidence (the conventional confidence level 95% is used in this search).

In this search, the modified frequentist method – also known as CL_s – is used to set conservative upper limits on the branching ratio $\mathrm{Br}(H \rightarrow aa \rightarrow 4\tau)$. A brief explanation of the CL_s method follows here, based on the overview given in [104]; a more in-depth treatment can be found in [105].

The expected background and signal yields b and s come respectively from the Standard Model and from the theory predicting the signal process (in this case, the NMSSM). The values of b and s are affected by systematic uncertainties due to various sources, which need to be accounted for in the design of the experiment. These uncertainties are represented by a set of nuisance parameters θ, and thus the predicted signal and background yields can be considered functions of these nuisance parameters: $s(\theta)$ and $b(\theta)$.

Given the observed data, the set of nuisance parameters θ, $s(\theta)$, and $b(\theta)$, a likelihood function $\mathcal{L}$ can be constructed:

$$\mathcal{L}(n|\mu,\theta) = Poisson(n|\mu,\theta) \cdot \rho(\theta) \tag{8.1}$$

In the case of this simple counting experiment with a single bin for the $m_{\mu+\mathrm{had}} > 4$ GeV search region (for a given M_T region), the Poisson function is simply the Poisson probability for observing n events in the search region bin. The variable μ is the signal strength modifier, which scales all predicted Higgs production cross sections by a factor of μ. The posterior function $\rho(\theta)$ is the probability distribution function (p.d.f.) for the systematic uncertainties. In this study, for sources of uncertainty not related to stochastic effects or sample size, the nuisance parameter p.d.f.'s are parametrized by log-normal functions, while nuisance parameters errors related to statistical uncertainties in sample sizes are parametrized by gamma functions [106]. For a full list and description of the sources of systematic uncertainty accounted for in this search, see Section 8.3.

From the likelihood function, the test statistic q_μ is defined for a given μ as:

$$q_\mu = -2 \ln \frac{\mathcal{L}(n|\mu,\hat{\theta}_\mu)}{\mathcal{L}(n|\hat{\mu},\hat{\theta})} \tag{8.2}$$

Here, $\hat{\mu}$ and $\hat{\theta}$ are the value of μ and the values of the nuisance parameters θ that give

the global maximum of $\mathcal{L}$, and $\hat{\theta}_\mu$ is the set of values of θ that maximize $\mathcal{L}$ for the given value of μ. One can then find the observed value of q_μ given the observed number of events n and the value of μ used in the signal+background hypothesis, as well as the values θ_0^{obs} and θ_μ^{obs} of the nuisance parameters that maximize the likelihood for the background-only hypothesis and the signal+background hypothesis respectively.

Monte Carlo simulation is used to generate p.d.f.'s $f(q_0|\mu, \theta_0^{obs})$ and $f(q_\mu|\mu, \theta_\mu^{obs})$ for the background-only hypothesis and signal+background hypothesis respectively. One can then define the CL_s parameter as follows:

$$CL_s(\mu) = \frac{\int_{q_\mu^{obs}}^{\infty} f(q_\mu|\mu, \theta_\mu^{obs}) dq_\mu}{\int_{q_0^{obs}}^{\infty} f(q_\mu|0, \theta_0^{obs}) dq_\mu} \tag{8.3}$$

If (1 - $CL_s(\mu)$) is less than the desired confidence level of 95%, then the signal+background hypothesis is said to be excluded at that level, and an upper limit can be set on the parameter of interest (i.e., the branching ratio Br($H \rightarrow aa \rightarrow 4\tau$)) based on the value of μ for which (1 - $CL_s(\mu)$) $< 95\%$.

In calculating observed limits, the value of n in Equation 8.2 is the experimentally observed number of events in the search region. Expected limits are calculated by taking n to be equal to the predicted number of background events in the search region; these are the limits that would be set if the experimental observation were to coincide exactly with the background-only hypothesis prediction. The signal model can be excluded where the observed limits are lower than the expected limits (meaning that fewer events were observed than were expected from the background model).

8.3 Systematic uncertainties

The following is a list of the sources of systematic uncertainty used in the calculation of the total uncertainty in this search, some of which have been mentioned in previous chapters. For the limit calculation, these systematics are all treated as nuisance parameters affecting only the scale of the expected signal or background yields, and they are modelled with log-normal distributions.

- **Luminosity:** As assessed in summer 2013 [35], the uncertainty on the integrated

luminosity is taken to be 2.6%.

- **Muon trigger efficiency:** As measured in dedicated single-muon HLT efficiency studies [96] (see Section 6.1.2), the systematic uncertainty from the single-muon trigger `HLT_IsoMu24_eta2p1` is 0.2% for the WH and ZH signals. For the ggH and VBF signals, because of the effect of the nearby lepton filter applied to the trigger muon, a larger systematic uncertainty of 4.2% is applied (see Sec. 6.1.2 for details).

- **Tight muon ID efficiency:** Using the results from muon ID efficiency studies [96] (see Section 6.1.1), the systematic uncertainty on the trigger muon tight ID is 0.5%.

- **Muon isolation efficiency:** The systematic uncertainty on the trigger muon isolation is taken to be 0.2% [96], as measured in the muon ID efficiency studies (see Section 6.1.4). This is applied to signal events in the WH and ZH channels. For the ggH and VBF channels, an uncertainty of 10% is used instead, to account for the fact that the muon which fires the trigger comes from a boosted $\tau_\mu\tau_{\rm had}$ topology, and that the isolation efficiency for the trigger muon is largely recovered if the nearby reconstructed tau is subtracted from its isolation cone; the 10% figure is taken from the CMS recommendation for the HPS tau ID efficiency for this boosted configuration (see Section 6.2.2 for details).

- **Soft muon ID efficiency:** According to muon ID efficiency measurements in $J/\psi \to \mu\mu$ events at CMS [96, 98], the systematic uncertainty on the τ_μ ID is 1.5%.

- **HPS ID efficiency:** The accepted value of 6% is used [97], as measured from τ_{had} ID efficiency studies [92] (see Section 6.2.2).

- **Tau charge misidentification rate:** The accepted value of -1%/+2% from CMS τ ID efficiency studies is used [97].

- **b-veto efficiency:** Two systematic uncertainties are considered for the b-veto efficiency. The first uncertainty stems from the fact that b-veto data/MC scale factors for light jets are applied to the tau jets on which the b-veto is applied; since

the actual data/MC scale factors are expected to be somewhere between light jets and b-jets, the percent difference in signal yields when using light jet scale factors and when using b-jet scale factors is taken as a systematic uncertainty, and the magnitude ranges between 1.8-8.5% depending on the M_T bin and signal process. The second source of systematic uncertainty comes from the uncertainty on the light-jet scale factors used; following the BTV recommendations, the scale factors are shifted coherently by $\pm 1\sigma$, and the difference between the nominal and shifted expected signal yields is taken as the systematic uncertainty; errors range up to 5.2% depending on signal sample. Because the VBF signal is expected to have a similar selection efficiency as the ggH channel, the errors calculated for each mass point in the ggH channel are applied to the VBF prediction for each analogous mass point; for similar reasons, the errors calculated for the WH channel are applied to the ZH channel.

- **ID efficiencies for nearby lepton filter around trigger muon:** Systematic uncertainties are assigned to the ID efficiency data/MC scale factors of the PF electrons, muons, and taus used for the neighbouring lepton veto around the trigger muon. For the PF electrons, since no ID is applied beyond the requirement that they pass PF reconstruction, have $p_T > 7$ GeV, and $|\eta| < 2.5$, we apply a conservative error of 1.1%, based on the highest uncertainty for the low-p_T electrons passing Loose ID requirements [95]. For the PF muons, since the same soft ID is used as for the reconstructed τ_μ, the same systematics uncertainty of 1.5% is applied [96]. For the PF taus, a conservative uncertainty of 10% is used. This came from our studies of the HPS tau ID efficiency for taus reconstructed from jets via the jet-cleaning method (cf. Section 5.2.2) with $p_T > 10$ GeV; this value was estimated by taking the standard uncertainty of 6% [97], adding in quadrature the discrepancy of at least 1% observed between the HPS tau ID efficiencies for our signal and Drell-Yan MC events in the studies described in Chapter 6, and rounding upwards.

- **Background:** To obtain the final jet-faking-tau background prediction in the $m_{\mu+\text{had}} > 4$ GeV bin in Region A, we take the unweighted average of the nomi-

nal background prediction from Region B and the predictions from the alternative non-QCD (from MC) or all-QCD (from region D data) background shapes. For the systematic uncertainty on this background prediction, we look at the central value $\pm\sigma$ of the alternative background shapes and compare it to the final background prediction; the greatest positive (negative) difference between the final background prediction and one of these values is then taken to be the positive (negative) systematic error on the final background prediction. This results in asymmetric systematic errors of +77.6%/-85.0% for the low-M_T bin and +60.9%/-59.2% for the high-M_T bin.

- $M_{\mathbf{T}}$**:** Using calculations done on our signal samples (calculation procedures described at [91]), errors range up to 12.2% depending on the pseudoscalar mass and production channel. Just as for the b-veto errors, the MET errors calculated for each mass point in the ggH channel are applied to the VBF prediction for each analogous mass point, while the errors calculated for the WH channel are applied to the ZH channel.

- **VBF and ZH predictions:** The expected signal yields from the VBF and ZH channels are calculated by scaling the ggH and WH expected yields (from MC) respectively to the appropriate SM cross-sections. To account for the extra element of uncertainty introduced by this indirect method of estimation, the percent difference between the number of VBF (or ZH) events from MC and from the indirect estimation method after the full selection at the 9 GeV pseudoscalar mass point (the only mass point for which MC samples were generated for the VBF and ZH channels) is taken as an estimate of the error on the VBF and ZH predictions for all pseudoscalar mass points. For the low-M_T bin, the errors were 23.2% for VBF and 19.1% for ZH; for the high-M_T bin, the errors were 25.3% for VBF and 24.3% for ZH.

8.4 Interpretation

8.4.1 Limit calculation

Algorithms from the `HiggsAnalysis/CombinedLimit` package (full code found at [107], documented in [106]) are used to calculate the limits in this search. The signal strengths calculated by this package are defined as:

$$\mu = \frac{\sigma_{\text{prod}} \cdot Br(H \to aa \to 4\tau)}{(\sigma_{\text{prod}})_{\text{expected}}} \tag{8.4}$$

where σ_{prod} is the production cross-section for the channel of interest and the denominators are calculated using the SM Higgs production cross sections given in [80] (e.g., $\sigma(\text{ggH}) = 19.27$ fb and $\sigma(\text{WH}) = 0.7046$ fb for $m_H = 125.0$ GeV). For the cases of WH and ZH, the production cross-sections are considered to be multiplied by the appropriate SM branching ratio for the decay of the vector boson to leptons. Since Standard Model Higgs production cross-sections are assumed, the cross-sections at the numerator and denominator cancel out, leaving μ equal to the branching ratio $\text{Br}(H \to aa \to 4\tau)$; thus, the limits calculated on μ by the `combine` [108] package are limits on this branching ratio.

The CL_s method is used to obtain observed and expected upper limits for each M_T bin and pseudoscalar mass point. For each M_T bin, limits are calculated for the signal strength parameter corresponding to the combination of all four signal channels (ggH, WH, VBF, and ZH).

8.4.2 Model-independent limits

The total expected yield from all four signal channels was used to calculate model-independent expected and observed CL_s limits on the signal strength. Plots of the observed and expected limits (median, $\pm 1\sigma$, and $\pm 2\sigma$) for the low-M_T bin, the high-M_T bin, and the combination of the two M_T bin at different m_a points are shown in Figure 8.1. The limits are reported in terms of the total branching ratio $\text{Br}(H \to aa \to 4\tau)$, assuming SM Higgs production cross-sections.

These model-independent observed and expected CL_s limits, expressed in terms of limits on $\text{Br}(H \to aa \to 4\tau)$ assuming SM Higgs production cross-sections, are also tabulated

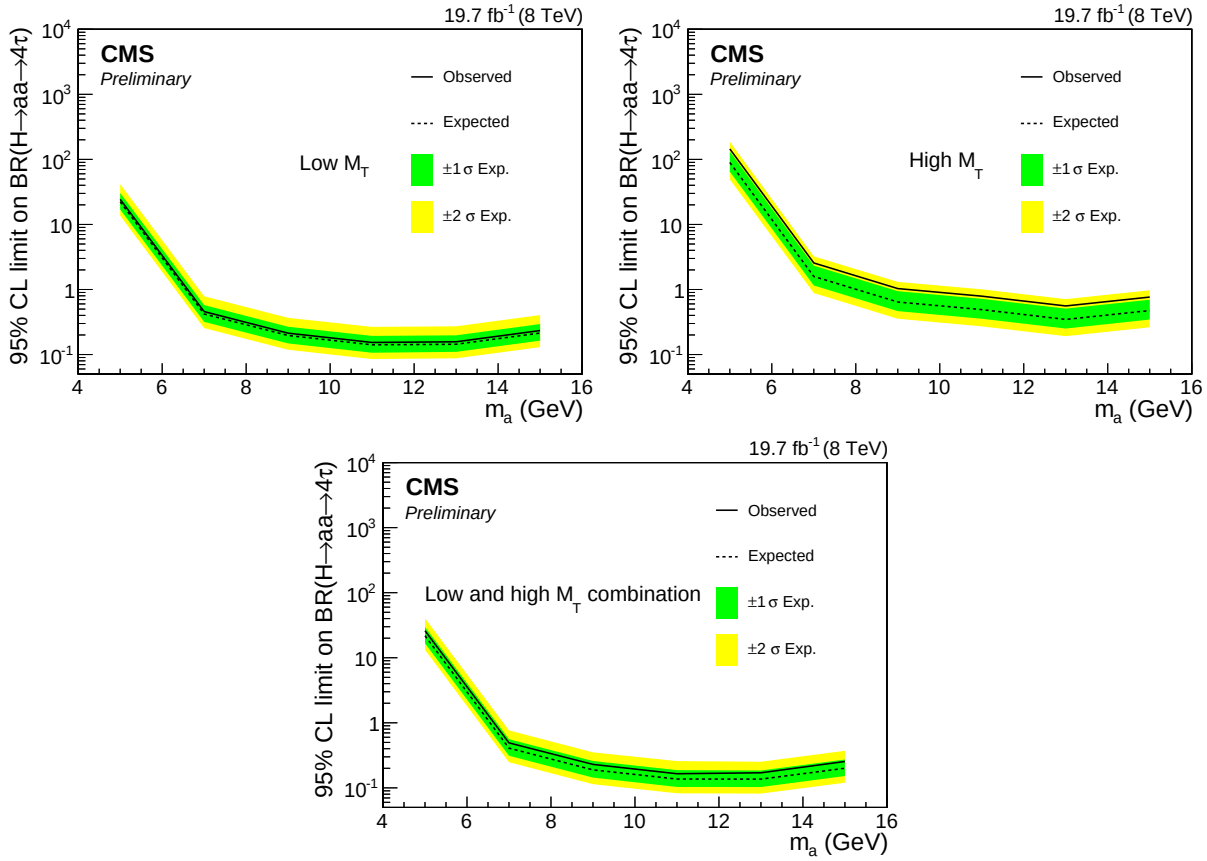

Figure 8.1. Observed 95% C.L. limits (solid black curve) on the branching ratio Br($H \rightarrow aa \rightarrow 4\tau$), compared to expected limits (dotted black curve, with $\pm 1\sigma$ bands in green and $\pm 2\sigma$ bands in yellow) at pseudoscalar mass points $m_a = 5$ through 15 GeVcc. (Top Left) $M_{\mathrm{T}} < 50$ GeV. (Top Right) $M_{\mathrm{T}} > 50$ GeV. (Bottom) Combined result between the low- and high-M_{T} bins.

in Tables 8.2-8.4.

Table 8.2. Observed and expected CL_s limits on $Br(H \rightarrow aa \rightarrow 4\tau)$ assuming SM Higgs production cross-sections in the low-M_T bin.

m_a (GeV)	-2σ	-1σ	Median expected	+1σ	+2σ	Observed
5	13.6	17.0	22.3	30.4	41.7	24.3
7	0.255	0.318	0.420	0.576	0.787	0.457
9	0.119	0.148	0.196	0.268	0.367	0.213
11	0.0852	0.107	0.141	0.193	0.266	0.153
13	0.0875	0.110	0.144	0.197	0.272	0.157
15	0.130	0.163	0.214	0.294	0.404	0.234

Table 8.3. Observed and expected CL_s limits on $Br(H \rightarrow aa \rightarrow 4\tau)$ assuming SM Higgs production cross-sections in the high-M_T bin.

m_a (GeV)	-2σ	-1σ	Median expected	+1σ	+2σ	Observed
5	49.2	64.7	89.6	132	186	145
7	0.883	1.15	1.59	2.33	3.23	2.56
9	0.355	0.465	0.643	0.942	1.31	1.03
11	0.271	0.355	0.490	0.719	1.005	0.789
13	0.192	0.251	0.347	0.508	0.711	0.559
15	0.262	0.345	0.475	0.696	0.973	0.765

Table 8.4. Observed and expected CL_s limits on $Br(H \rightarrow aa \rightarrow 4\tau)$ assuming SM Higgs production cross-sections for the combination of the low- and high-M_T bins.

m_a (GeV)	-2σ	-1σ	Median expected	+1σ	+2σ	Observed
5	13.2	16.7	21.8	29.7	40.4	26.0
7	0.248	0.312	0.408	0.560	0.765	0.491
9	0.114	0.144	0.189	0.259	0.352	0.230
11	0.0825	0.103	0.137	0.187	0.258	0.165
13	0.0817	0.103	0.136	0.185	0.252	0.171
15	0.120	0.152	0.200	0.273	0.371	0.254

Chapter 9

Conclusions

We have performed search for the decay $H \rightarrow aa/hh \rightarrow 4\tau$ in the gluon fusion, W associated, Z associated, and vector boson fusion production modes. The observed data was consistent with Standard Model expectations, and no evidence of this exotic decay was found. Thus, we have set model-independent upper limits on the branching ratio to new physics, assuming SM production of the 125 GeV Higgs. For a 9 GeV pseudoscalar, an upper limit of 18.9% was set on $BR(\mathrm{H} \rightarrow 4\tau)$. The most stringent limits, 13.7% and 13.6%, were set at the 11 and 13 GeV pseudoscalar mass points respectively. These branching ratios can be interpreted in the context of any 2HDM models that allow the decay of H to light scalars or pseudoscalars. This result is the first of its kind at the LHC.

The boosted tau identification techniques developed in this search can find promising use in future searches during the LHC's Run II at 13 TeV. Decays such as $H \rightarrow aa \rightarrow 2\mu 2\tau$ and $H \rightarrow aa \rightarrow 2\tau b\bar{b}$ remain to be explored. Searches need not be limited to Higgs studies either, as any other event topology involving boosted tau pairs could benefit from these techniques.

Appendices

Appendix A

N-subjettiness

Highly boosted particles can result from the decay of high-mass resonances, whose production in the LHC is made possible by the high collision energies achievable. The decay products of a boosted object appear as a collimated spray of tracks in the detector; with a sufficiently high boost factor and thus sufficiently high collimation, these decay products can be reconstructed as a single jet and are thus not identified as distinct objects. Techniques for probing jet substructure are important for identifying and analyzing boosted jets. One method is the use of N-subjettiness (τ_N), a parameter that measures the degree to which the energy within a jet is aligned along N candidate subjet axes
The formula for N-subjettiness is as follows:

$$\tau_N = \frac{1}{d_0} \sum_k p_{T,k} min(\Delta R_{1,k}, \Delta R_{2,k}, \ldots, \Delta R_{N,k}) \tag{A.1}$$

The index k goes over all the constituent particles in the jet, $p_{T,k}$ is the transverse momentum of the k^{th} particle, and $\Delta \mathrm{R}_{n,k}$ is the distance in $\eta - \phi$ space between the k^{th} particle and the axis of the n^{th} candidate subjet. The term d_0 is given by

$$d_0 = \sum_k p_{T,k} R_0 \tag{A.2}$$

where R_0 is the radius used in the original jet clustering algorithm [109].
In a jet whose particles are closely aligned with N or fewer subjets, the terms $p_{T,k}\mathrm{min}(\Delta \mathrm{R}_{1,k},\ \Delta \mathrm{R}_{2,k},\ \ldots,\ \Delta \mathrm{R}_{N,k})$ in the sum will be very small, and thus τ_N will be

closer to zero, while in a jet whose energy is distributed away from the N subjet axes will have a larger value of τ_N and must have at least N + 1 subjets.

N-subjettiness has been used successfully in the identification of boosted objects such as top quarks and W bosons. A new use of N-subjettiness for identifying jets seeded by boosted tau pairs (referred to as boosted ditau jets) was probed in a theoretical study by Englert et al. [110], which suggested that the ratio τ_3/τ_1 could provide discrimination between boosted ditau jets and QCD jets. This study used 14-TeV Monte Carlo signal and background samples under conditions of zero pileup, where the signal process was $h_1 \to 2a_1 \to 4\tau$, with all inclusive tau decay modes considered.

In this study, I explored N-subjettiness ratios τ_3/τ_1, τ_2/τ_1, τ_1/τ_2, τ_2/τ_3, and τ_3/τ_4 for their possible discriminatory power in identifying boosted tau jets. An important issue that arose was the influence of pileup on the N-subjettiness distribution, which tends to impair the discriminatory power of N-subjettiness ratios; for instance, the mean of the τ_3/τ_1 distribution was observed to increase with increasing pileup for signal Monte Carlo, causing it to become increasingly indistinguishable from the τ_3/τ_1 distribution for jets from W+NJets events. Jet pruning was used to remove pileup from jets, and although this recovered some of the discriminatory power when comparing unit-normalized signal and W+NJets Monte Carlo τ_3/τ_1 distribution shapes, further analysis – comparing signal Monte Carlo and all other background Monte Carlo samples except QCD, after applying pileup reweighting and all the preselection cuts to these events – has not shown significant discrimination between signal and background.

Also, as a result of the jet pruning, a significant number of jets were left with only 3 or fewer constituents, resulting in τ_3/τ_1 values of exactly 0. This suggests the need for using less aggressive methods of pileup removal from jets. Figures A.1 and A.2 show the distributions of two N-subjettiness ratios, τ_3/τ_1 and τ_1/τ_2, for the low and high M_T bins, illustrating both the currently insufficient discriminatory power of these variables and the problematic peak at zero (for τ_3/τ_1) caused by jet pruning. Further investigation will eventually be required to find an optimal method of pileup removal and potentially an improvement in discriminatory power for N-subjettiness ratios.

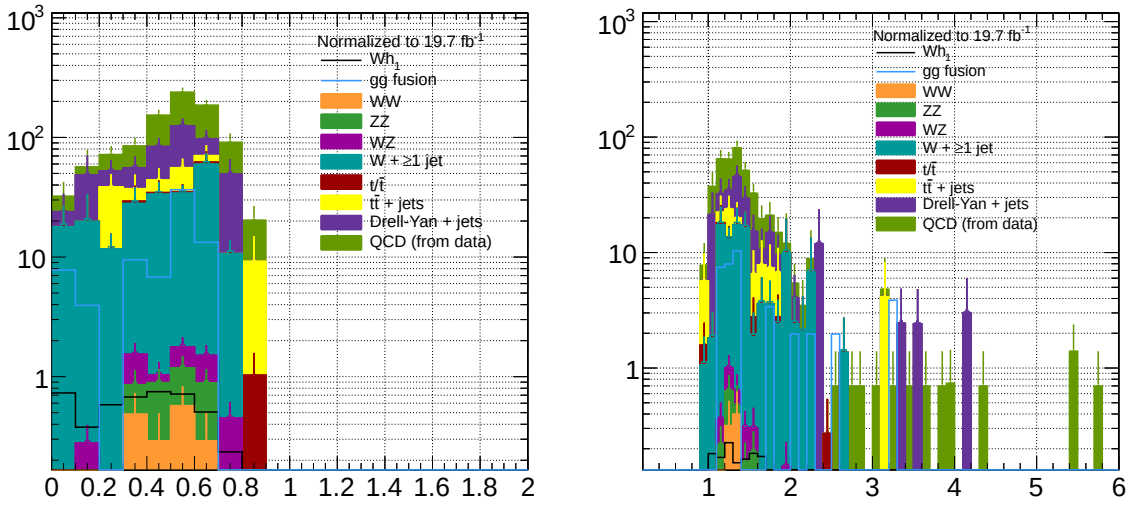

Figure A.1. Examples of N-subjettiness ratio distributions for the low-M_T bin, comparing distributions for two signal models and all backgrounds discussed in Sec. 5 including data-driven QCD, after all the preselection cuts have been applied. (Left) τ_3/τ_1. (Right) τ_1/τ_2.

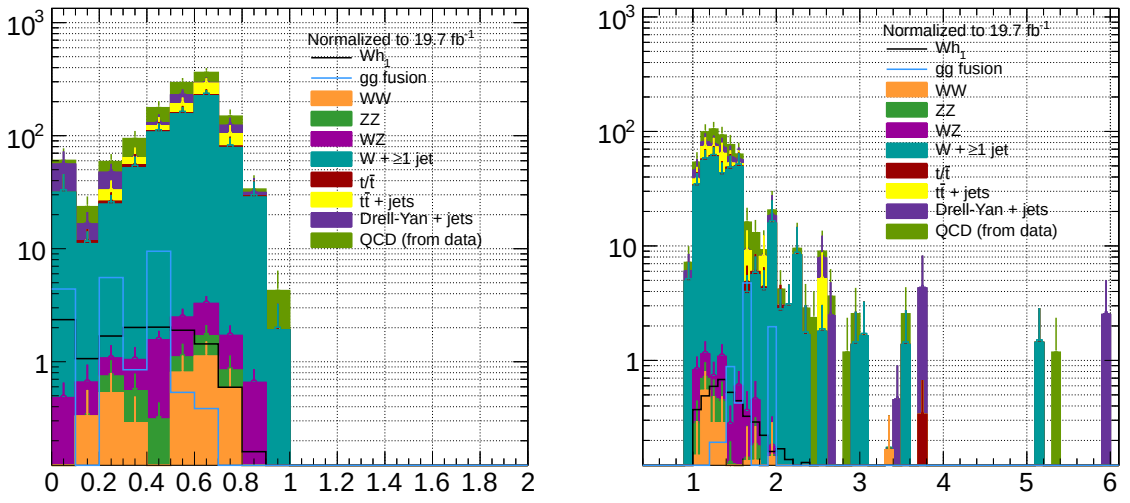

Figure A.2. Examples of N-subjettiness ratio distributions for the high-M_T bin, comparing distributions for two signal models and all backgrounds discussed in Sec. 5 including data-driven QCD, after all the preselection cuts have been applied. (Left) τ_3/τ_1. (Right) τ_1/τ_2.

Appendix B

Pixel detector geometry description

The CMS detector simulation uses an XML schema called the Detector Description Language (DDL) to encode the description of the detector geometry and material composition [111]. Together with two other auxiliary packages, the Algorithm Description Language (ADL) and Configuration Description Language (CDL), DDL interfaces with GEANT4 to provide the volumes, positions, and material compositions of the simulated detector elements.

All the various components and subcomponents of the CMS detector form a geometrical hierarchy, with each individual object being a subcomponent of some larger whole. In a system of XML files, DDL defines the basic data structures for describing the dimensions and materials of these parts as well as their geometrical hierarchies.

Two different classes of material definitions exist in DDL. Elementary materials correspond to elements of the periodic table and are identified by name, periodic table symbol, density, atomic number, and atomic weight. Definitions of composite materials are built by specifying their fractional composition in terms of elementary materials, or even other composite materials; the data structure allows one to customize the density associated with a particular composite material. The radiation length and hadronic interaction length of composite materials are calculated from the material definitions in the XML files; these parameters are needed ifor the modelling of particle interactions with detector components.

Detector parts are defined by their material, their shape, and their position in the de-

tector. The various types of 3-dimensional shapes (such as rectangular boxes, trapezoids, and cylinders) allowed in this package are based on the GEANT syntax. Various parameters such as angles and Cartesian coordinates encode a detector component's spatial position, often with respect to a larger structure of which it is a component. Algorithms from ADL, referred to as DDAlgorithms, are used to position multiple copies of a detector component in a specific pattern, to represent symmetrical or repeating structures.

The rest of this chapter treats the projects involving the CMS pixel detector geometry description, in which I have participated.

B.1 Pilot system simulation

From February 2013 to February 2015, the LHC was turned off. During this period of planned off-time, while the LHC was being prepared for proton-proton collisions at the design centre-of-mass energy $\sqrt{s} = 14$ TeV, the CMS collaboration took the opportunity to perform repairs and maintenance on the CMS detector and install a new beampipe with a smaller diameter. To make room for the beampipe replacement, the pixel detector was extracted from the experimental cavern. While it was being stored in a lab aboveground, longstanding problems with its panels and electronics were diagnosed and fixed, and the barrel and endcap systems were calibrated in preparation for reinstallation in the cavern.

In addition to repairs and calibration, one extra endcap disk was installed on the -z side of the forward pixel detector (FPIX). 8 prototype modules for the planned Phase I upgrade [25] were mounted on the blades of this extra disk; Figure B.1 shows the third half-disk in one of the -z FPIX half-cylinders, with the Phase I prototype modules mounted. DC-DC converters and portcards were also installed in the service cylinder, to power the third disk and read out the modules respectively. Since this extra disk and its associated electronics serve as a pilot run for the Phase I modules, the ensemble is referred to as the pilot system.

The tracker geometry and material description at the time did not include a description of the pilot system. In order to have an accurate representation of the material distribution (often referred to as the material budget) in the tracker for generating MC events with this

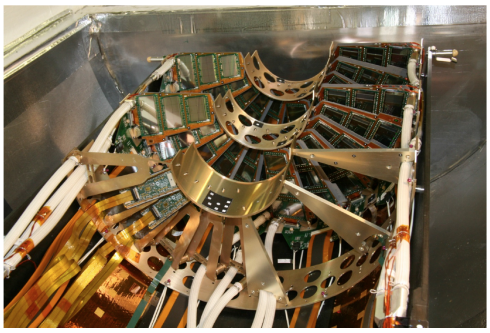

Figure B.1. -z FPIX half-cylinder, containing one pilot half-disk (foreground) in addition to its two standard half-disks.

new pixel detector configuration, the pilot system needed to be added to the description.

A disk on the -z endcap was cloned into the position where the new disk was installed; new objects were declared in the geometry description to represent the shape and material of the new modules. The large uniform blocks roughly representing the FPIX portcard electronics in the service cylinder also were updated, since the material composition had changed due to the addition of copper-containing DC-DC converters and other new electronics for the pilot system's readout. Thus, a new composite material representing the new average material composition was defined. Figure B.2 shows a visualization of the simulated pilot disk with its modules. Figure B.3 shows the configuration of FPIX service cylinder electronics prior to the pilot installation, as well as an illustration of the pilot system in the final simulation.

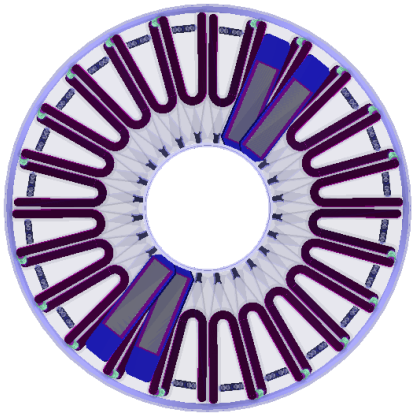

Figure B.2. Fireworks visualization of the added pilot disk with its modules in the FPIX geometry description.

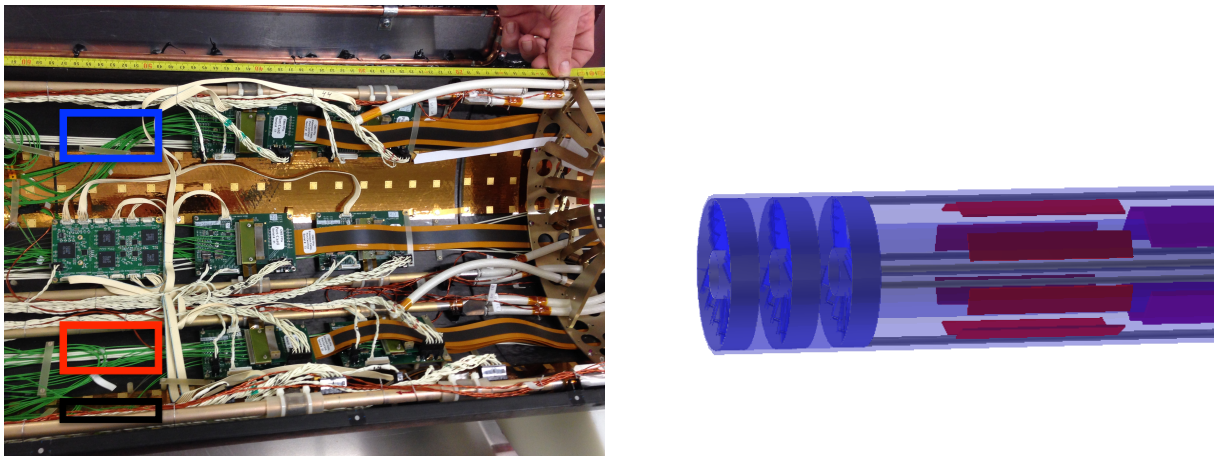

Figure B.3. (Left) The FPIX service cylinder before the pilot system installation. The rightmost two columns of circuitboards are the portcards for the two standard FPIX half-disks. The leftmost column contains only the digital communication and control unit (CCU) for the two FPIX half-disks. The coloured rectangles indicate the intended positions of the pilot electronics to be installed. Red: CCU board. Blue: Portcard. Black: DC-DC converter board. (Right) Fireworks visualization of the pilot system in the FPIX geometry description. The pilot disk is the third from the left, and the red rectangles represent the portcard electronics in the service cylinder.

B.2 Phase I pixel geometry simulation

During the next long shutdown of the LHC scheduled for 2018, the CMS detector will undergo another round of detector upgrades referred to as the Phase I upgrade [25], to fix weaknesses in the current systems and improve detector performance at the higher luminosities expected in the future. The pixel detector will acquire one new barrel layer and one new endcap on each side of the interaction point, to provide redundancy in track hit pattern recognition, reduce fake rates at high pileup, and still allow decent tracking performance even if the inner layer undergoes more radiation damage than expected. Faster electronics will be installed, for efficient operation at high event rates. The cooling system, which currently uses C_6F_{14} as a coolant, will be replaced with a more lightweight cooling system that uses cold carbon dioxide instead. In general, the design of the upgraded pixel detector's support structures is aimed at decreasing the material budget in the tracking volume.

The CMS geometry description currently has XML files that describe the Phase I pixel geometry and material. The barrel pixel part is mostly accurate aside from a few

minor changes involving the dimensions of some support structures and the addition of aluminium cabling in some regions. However, the forward pixel part is very inaccurate and needs significant updating in order to represent the actual upgraded components that will be installed. The main challenges are the following:

- The geometrical arrangement and symmetry of the FPIX blades have been found to be fundamentally incorrect. This is arguably the most critical issue with the Phase 1 FPIX description, since the positions of the pixel sensors in simulation are determined by the position of the blades. Comparisons with the blueprints of the actual Phase I FPIX blades in their disks to determine the correct arrangement are currently underway, together with efforts to correct the positioning of the blades in the detector description code.

- The FPIX support rings are currently modelled by flat, uniform rings in simulation, whereas the actual support rings have a zigzagging shape (see Figure B.4). To achieve a more accurate description of the material budget, the shape of the rings needs to be corrected. These shapes are, however, extremely difficult to render using the shapes allowed in DDL. A tentative solution is being developed and tested, in which a series of "infinitesimally" thin blocks are positioned in the form of a zigzag-shaped ring using a DDAlgorithm that gives each block an appropriate displacement along z as a function of ϕ. Figure B.5 illustrates this proposed solution.

- In simulation, the endcap disks on the +z side are obtained by rotating the -z disks about the y axis, whereas the actual Phase I disks have a mirror symmetry about the xy plane instead.

- The modules and the blades currently have very rough and basic shapes in the simulation (see Figure B.6); the dimensions and positions of the modules need to be checked for accuracy.

- The composite material used to describe the Phase I portcard objects is the same material used to describe the portcard objects in the Run I geometry before the

installation of the pilot system. This is clearly incorrect, and a new composite material needs to be declared that matches the average composition of the actual Phase I FPIX portcards and DC-DC converters.

Since January 2015 and continuing into my postdoctoral position, I have been the convener of the CMS tracker material budget group, overseeing any tasks that involve changes to the geometry/material description or studies of the tracker structure in data. I have been directly involved in updating the Phase I pixel geometry description, as well as coordinating the other individuals involved in this effort, and much of this work is still ongoing. In particular, I designed the solution shown in Figure B.5 for representing the zigzag-shaped FPIX support rings, which is still being tested and refined, and I am developing a new simulated material to represent the Phase I FPIX supply tube electronics. I am communicating with the engineers who are building the Phase I FPIX and BPIX components, to gather the necessary information about their dimensions and material composition in order to represent them accurately in the simulated description.

Also, I am overseeing a group who are updating a software package for reconstructing and analyzing nuclear interaction vertices, in order to use reconstructed nuclear interaction vertices in the tracker material to image the structure of the tracker, measure the positions of tracker structures and the beam-pipe (see Figure B.7), and validate the geometry description in simulation. My contribution to this effort consists of technical support with software and providing a point of contact with and the tracking algorithm experts and pixel detector operations group.

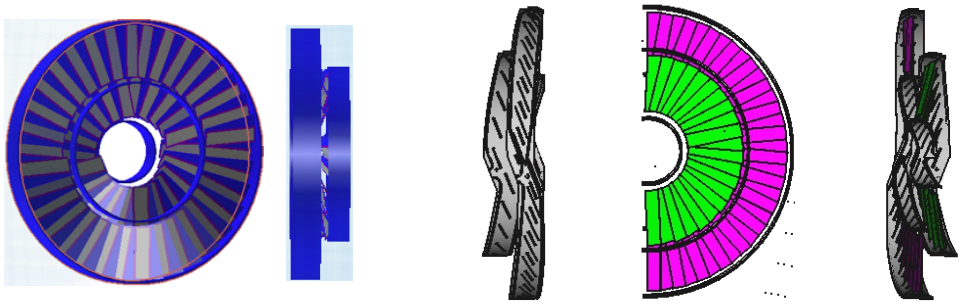

Figure B.4. Comparison of the shape of the Phase I FPIX half-disks in simulation (Left) and in the actual detector being constructed (Right).

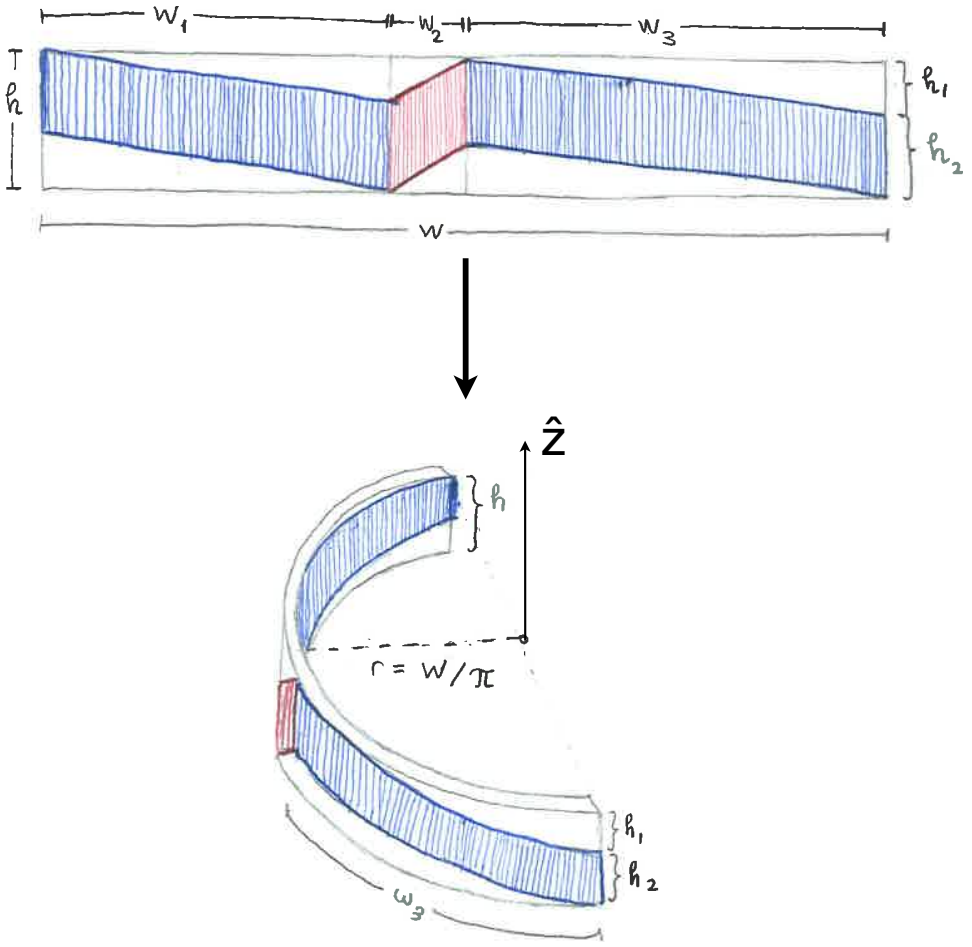

Figure B.5. Illustration of the proposed solution for modelling the zigzag-shaped support rings of the Phase I FPIX half-disks. The blue and red colours denote the three basic legs that make up the zigzag. A DDAlgorithm is used to position thin rectangular blocks into an approximation of the zigzagging shape by arranging them in a ring and giving each block an incremental displacement in the z-direction from the plane of the ring, where the displacement depends on the azimuthal angle ϕ around the ring.

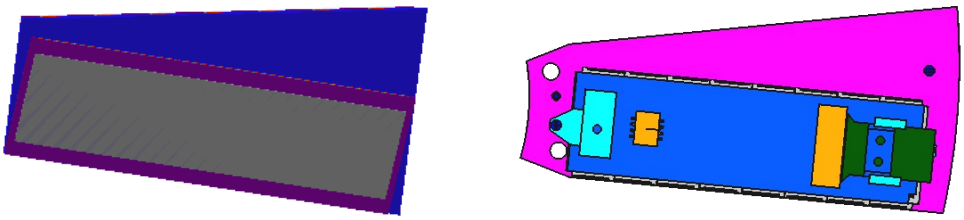

Figure B.6. Comparison of the shape of the Phase I FPIX blades and modules in simulation (Left) and in the actual detector being constructed (Right).

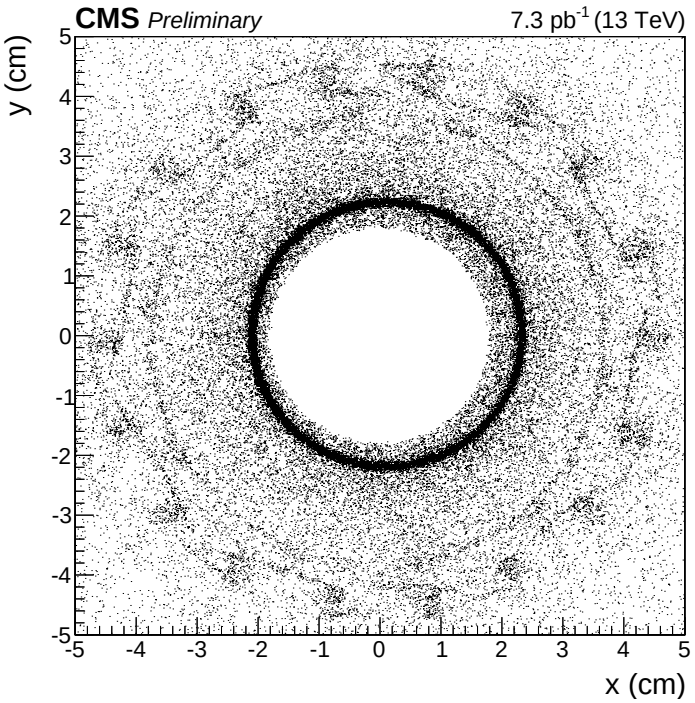

Figure B.7. Cross-sectional view of the beam pipe (dark circle at radius 2.25 cm) surrounded by support structures of the barrel pixel detector, imaged with reconstructed nuclear interaction vertices in data from 13-TeV proton-proton collisions in 2015. [112]

Appendix C

Forward pixel maintenance during LS1

As mentioned in Section B.1, the CMS pixel detector was extracted from the experimental cavern and housed in a special facility aboveground for repairs, maintenance, and calibration during the first long shutdown (LS1) of the LHC from February 2013 to February 2015.

Starting in the summer of 2012 until the extraction of the pixel detector in May 2013, I worked with the UC Davis forward pixel team to set up the lab that would be used to house the extracted pixel detector. Each of the four half-cylinders of the forward pixel (FPIX) detector and two half-barrels of the barrel pixel (BPIX) detector was stored in a protective insulated box called a "cold box", with temperature and humidity controlled by the following systems:

- **Water-glycol cooling system:** A system of pipes cycles a chilled 50%/50% mixture by volume of water and glycol continuously through the cooling tubes of the pixel half-cylinder supply tube. The water-glycol chiller set point is 0° C for the BPIX cold boxes and -2° C for the FPIX cold boxes; these set points have been empirically optimized so as to achieve a temperature below 10° C within the cold boxes, as monitored by temperature sensors at various positions in the boxes. The purpose of keeping the general environment of the pixel detector cold is to minimize the spread of radiation-induced defects in the crystal structure of the silicon sensors

due to thermal agitation.

- **C_6F_{14} cooling system:** A system of pipes cycles chilled liquid C_6F_{14} through the cooling tubes that serve the sensitive regions of the pixel-half cylinder (panels in FPIX and modules in BPIX). The C_6F_{14} is kept at a controlled temperature by a dedicated C_6F_{14} chiller. When the panels or modules of the half-cylinder are turned on for testing and maintenance, extra heating occurs due to the currents in the electronics; thus, the C_6F_{14} cooling system serves to counter this extra heating and keep the temperature at the panels or modules below 20° C. The C_6F_{14} cooling system, with the chiller set point set to -15° C, is only used when the half-cylinder panels or modules are powered on.

- **Dry air system:** To prevent humidity from damaging the sensitive electronics of the pixel detector, the interior of the cold box is supplied with room-temperature dry air from an air dryer via pipe lines leading into the boxes, while the humidity inside the boxes is monitored by dew point sensors at various positions.

I helped to prepare the cold boxes for transport to the lab and to set up all three of the above-listed systems to serve the cold boxes, as well as the power supply crates and DAQ systems to power and read-out respectively the forward pixel detector half-cylinders. Figure C.1 shows a photo of two of the cold boxes, as well as the power supply and DAQ crates and the cooling system.

Once the pixel detector was extracted and the half-cylinders and half-barrels were successfully installed in their respective cold boxes, the forward pixel team proceeded to tackle known longstanding problems with certain panels in the FPIX. Here is a brief list of the issues:

- Panels needing replacement, due to showing no analog signal.

- "Slow" channels: the analog signal from certain panels has a slow rise time, requiring investigation (see Figure C.2 for an illustration).

- Difficulty in programming several panels, requiring investigation.

Figure C.1. Photo of the pixel lab at the CMS experimental site, showing two of the FPIX cold boxes, with the system of cooling pipes in the back bringing water-glycol and C_6F_{14} to cool the interior of the boxes, and the power supply and DAQ crates (blue) to the left for reading out one half-cylinder at a time.

- Inability to communicate with one analogy optical hybrid (AOH).
- The sense wires carrying digital and analog power voltages from the portcard to the panel had become disconnected and need to be reconnected.
- The connectors of the some of the optical fibres had got contaminated with dust particles and needed cleaning.
- A panel with one dead ROC. (This problem is insignificant and not worth the delicate operation of replacing the panel, so it was left alone.)

Several of the above problems were traced to misaligned or dislodged flex cables that provide communication from the portcards to the panels, for various functions such as programming the ROCs and reading out analog signals. In particular: the slow-channel issues, the difficulty with programming certain panels, and faulty communication with one of the AOH's were all fixed by correcting the alignment of these flex cables in their connectors. I participated in the diagnosis of these problems by assisting in reading out the analog signal coming from problematic panels at different points along the readout circuit,

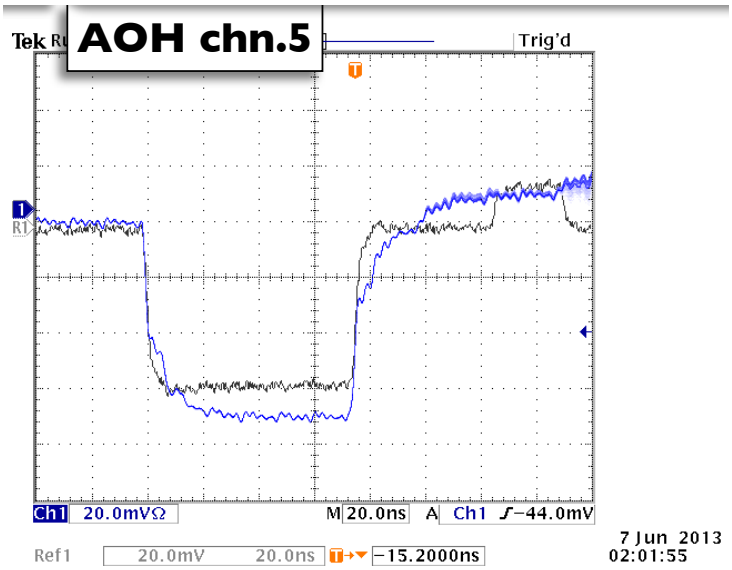

Figure C.2. Oscilloscope reading of the shape of the analog header signal from a normally-functioning channel (black curve) compared to a "slow" channel (blue curve). Note the lack of a clean drop at the beginning of the blue curve's header signal, and the slow rise time at the end.

and by replacing portcards when necessary. I also helped to clean out the optical fibre connectors and took week-long shifts to monitor the condition of the pixel half-cylinders in their cold-boxes.

Figure C.3 illustrates the fraction of live channels in the four FPIX disks at the end of Run 1, before the pixel detector was extracted. The problematic panels and ROCs can be seen as blank regions in the live channel display; roughly 16% of the FPIX channels were not functional. After the repairs during the long shutdown, including the replacement of certain panels, the inactive channels in the FPIX were largely recovered, and the total live channel fraction increased from 84% to more than 99.9% after the repairs. The live channel displays for FPIX after the repairs are shown in Figure C.4

After all the planned repairs were completed, the FPIX calibration sequence was performed on each half-cylinder individually. This set of calibration procedures programs and optimizes the various parameters of the FPIX readout system (illustrated schematically in Figure C.5); a treatment of these calibration procedures is beyond the scope of this dissertation, but a detailed description can be found in this source [113].

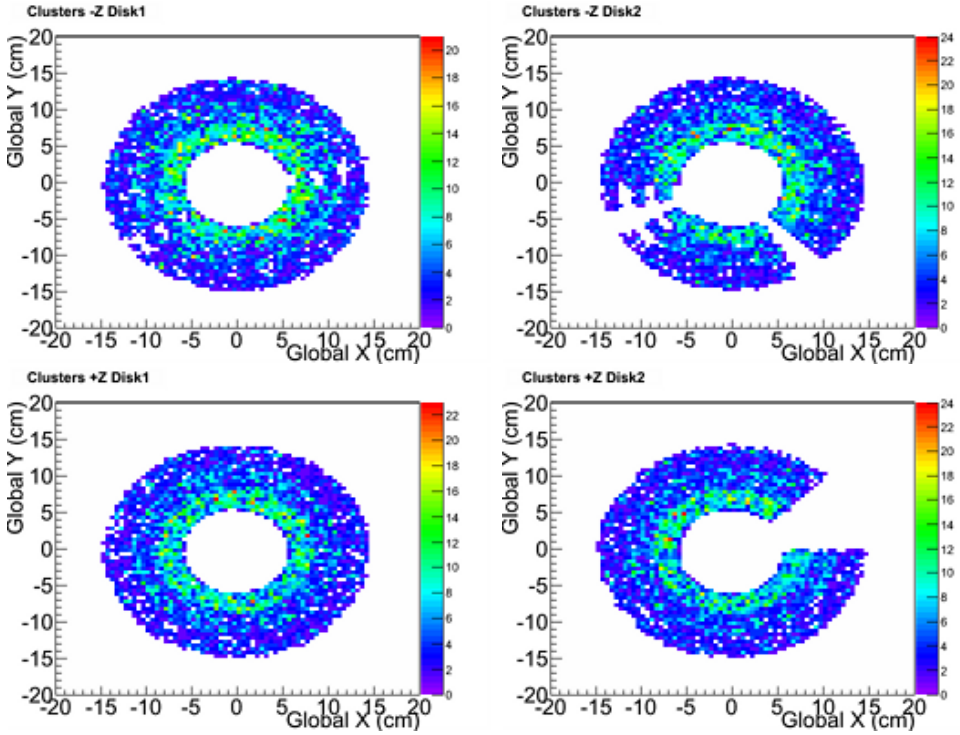

Figure C.3. Fraction of live channels in each of the four FPIX disks at the end of Run 1. Blank holes in the display indicate problems with panels and electronics, with the fraction of malfunctioning channels amounting to a total of 16%. (Top Left) -z Disk 1. (Top Right) -z Disk 2. (Bottom Left) +z Disk 1. (Bottom Right) +z Disk 2.

All four FPIX half-cylinders were successfully calibrated in the lab; I participated in the calibration of three out of the four. After the pixel detector was reinstalled in the experimental cavern at the end of 2014, I assisted the forward pixel team in calibrating the FPIX once more (this time calibrating all half-cylinders simultaneously instead of separately) to prepare it for data-taking at the new proton-proton collision energy of 13 TeV.

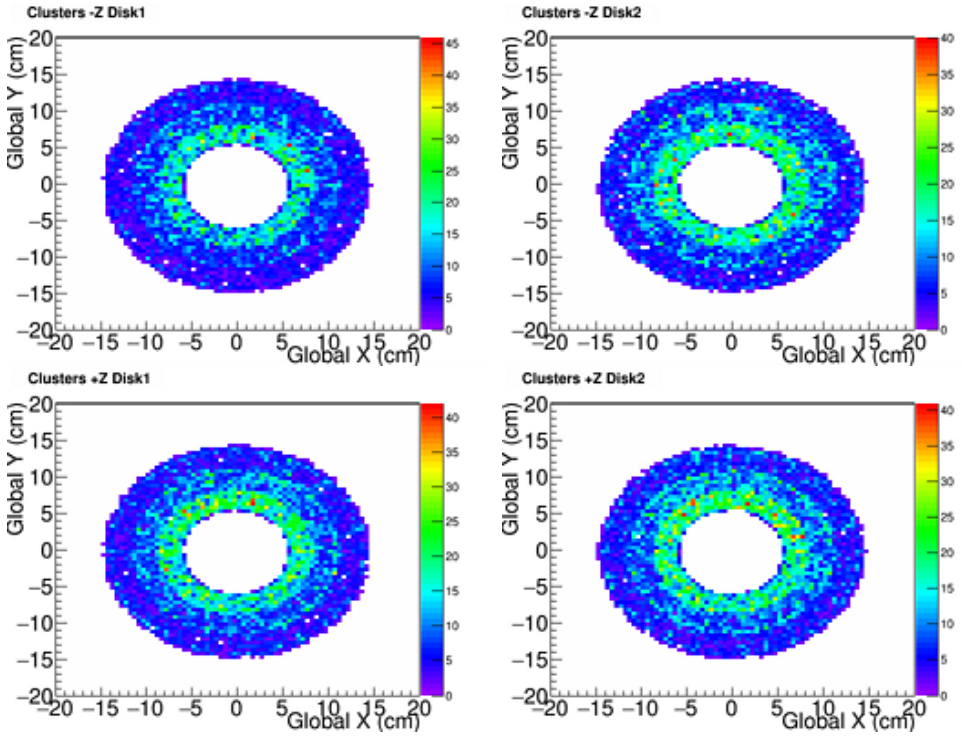

Figure C.4. Fraction of live channels in each of the four FPIX disks after the repairs during the long shutdown, revealing the almost complete recovery of the nonfunctioning channels; the fraction of live channels now exceeds 99.9%. (Top Left) -z Disk 1. (Top Right) -z Disk 2. (Bottom Left) +z Disk 1. (Bottom Right) +z Disk 2.

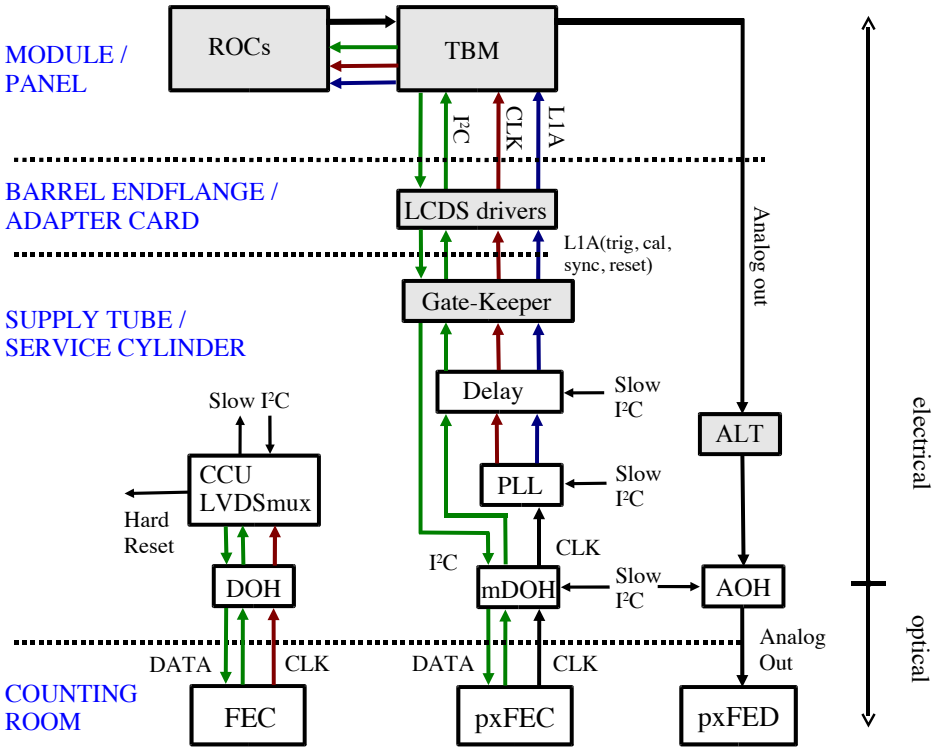

Figure C.5. Diagram of the CMS pixel control and readout system. [24]

List of Figures

List of Tables

References

[1] Albert Messiah and G. M. Temmer (translator). *Quantum Mechanics*, volume 1. John Wiley & Sons, Inc., 1958. 1

[2] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to Quantum Field Theory*. Westview Press, 1995. 2

[3] Mark Thomson. *Modern Particle Physics*. Cambridge University Press, 1st edition, 2013. 2, 6

[4] Alessandro Bettini. *Introduction to Elementary Particle Physics*. Cambridge University Press, 2nd edition, 2014. 2, 5, 8

[5] F. H. Combley. (g-2) factors for muon and electron and the consequences for QED. *Reports on Progress in Physics*, 42(12):1889, 1979. URL: `http://stacks.iop.org/0034-4885/42/i=12/a=001`. 6

[6] Donald H. Perkins. *Introduction to High Energy Physics*. Cambridge University Press, 4th edition, 2000. 6

[7] V. A. Bednyakov, N. D. Giokaris, and A. V. Bednyakov. On Higgs mass generation mechanism in the Standard Model. *Phys. Part. Nucl.*, 39:13–36, 2008. `arXiv:hep-ph/0703280`, `doi:10.1007/s11496-008-1002-9`. 6

[8] G. Arnison et al. Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540$ GeV. *Phys. Lett.*, B122:103–116, 1983. [,611(1983)]. `doi:10.1016/0370-2693(83)91177-2`. 6

[9] G. Arnison et al. Experimental Observation of Lepton Pairs of Invariant Mass Around 95-GeV/c^2 at the CERN SPS Collider. *Phys. Lett.*, B126:398–410, 1983. `doi:10.1016/0370-2693(83)90188-0`. 6

[10] Abdelhak Djouadi. The Anatomy of electro-weak symmetry breaking. I: The Higgs boson in the standard model. *Phys. Rept.*, 457:1–216, 2008. `arXiv:hep-ph/0503172`, `doi:10.1016/j.physrep.2007.10.004`. 7, 173

[11] Serguei Chatrchyan et al. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Phys.Lett.*, B716:30–61, 2012. `arXiv:1207.7235`, `doi:10.1016/j.physletb.2012.08.021`. 8, 58, 173

[12] Georges 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.*, B716:1–29, 2012. `arXiv:1207.7214`, `doi:10.1016/j.physletb.2012.08.020`. 8, 58

[13] Stephen P. Martin. A Supersymmetry primer. 1997. [Adv. Ser. Direct. High Energy Phys.18,1(1998)]. `arXiv:hep-ph/9709356`, `doi:10.1142/9789812839657_0001,10.1142/9789814307505_0001`. 9, 10, 11, 173

[14] Ian J. R. Aitchison. Supersymmetry and the MSSM: An Elementary introduction. 2005. arXiv:hep-ph/0505105. 11

[15] Serguei Chatrchyan et al. Search for a light pseudoscalar Higgs boson in the dimuon decay channel in pp collisions at $\sqrt{s} = 7$ TeV. *Phys.Rev.Lett.*, 109:121801, 2012. arXiv:1206.6326, doi:10.1103/PhysRevLett.109.121801. 12, 60

[16] Brando Bellazzini, Csaba Csaki, Antonio Delgado, and Andreas Weiler. SUSY without the Little Hierarchy. *Phys. Rev.*, D79:095003, 2009. arXiv:0902.0015, doi:10.1103/PhysRevD.79.095003. 12

[17] M. Maniatis. The Next-to-Minimal Supersymmetric extension of the Standard Model reviewed. *Int. J. Mod. Phys.*, A25:3505–3602, 2010. arXiv:0906.0777, doi:10.1142/S0217751X10049827. 12

[18] Ulrich Ellwanger and Michel Rausch Traubenberg. Natural range of higgs masses in supersymmetry. *Zeitschrift für Physik C Particles and Fields*, 53(3):521–527. URL: http://dx.doi.org/10.1007/BF01625914, doi:10.1007/BF01625914. 12

[19] J. Ellis, J. F. Gunion, H. E. Haber, L. Roszkowski, and F. Zwirner. Higgs bosons in a nonminimal supersymmetric model. *Phys. Rev. D*, 39:844–869, Feb 1989. URL: http://link.aps.org/doi/10.1103/PhysRevD.39.844, doi:10.1103/PhysRevD.39.844. 12

[20] Lyndon Evans and Philip Bryant. LHC Machine. *Journal of Instrumentation*, 3(08):S08001, 2008. URL: http://stacks.iop.org/1748-0221/3/i=08/a=S08001. 13

[21] APS News. https://www.aps.org/publications/apsnews/201409/backpage.cfm. Accessed: 2016-01-08. 14, 173

[22] Torbjörn Sjöstrand, Stephen Mrenna, and Peter Skands. PYTHIA 6.4 physics and manual. *JHEP*, 05:026, 2006. arXiv:hep-ph/0603175, doi:10.1088/1126-6708/2006/05/026. 16, 52

[23] CMS Collaboration. CMS Detector Drawings. URL: https://cmsinfo.web.cern.ch/cmsinfo/Media/Images/Detector/Detector%20Drawings/index.html. 18, 173

[24] The CMS Collaboration. The CMS experiment at the CERN LHC. *Journal of Instrumentation*, 3(08):S08004, 2008. URL: http://stacks.iop.org/1748-0221/3/i=08/a=S08004. 17, 23, 25, 27, 28, 29, 32, 34, 35, 37, 39, 40, 41, 42, 44, 45, 172, 174, 175, 176, 190

[25] A Dominguez, D Abbaneo, K Arndt, N Bacchetta, A Ball, E Bartz, W Bertl, G M Bilei, G Bolla, H W K Cheung, M Chertok, S Costa, N Demaria, A Dominguez, K Ecklund, W Erdmann, K Gill, G Hall, K Harder, F Hartmann, R Horisberger,

W Johns, H C Kaestli, K Klein, D Kotlinski, S Kwan, M Pesaresi, H Postema, T Rohe, C Schfer, A Starodumov, S Streuli, A Tricomi, P Tropea, J Troska, F Vasey, and W Zeuner. CMS Technical Design Report for the Pixel Detector Upgrade. Technical Report CERN-LHCC-2012-016. CMS-TDR-11, CERN, Geneva, Sep 2012. Additional contacts: Jeffrey Spalding, Fermilab, Jeffrey.Spalding@cern.ch Didier Contardo, Universite Claude Bernard-Lyon I, didier.claude.contardo@cern.ch. URL: https://cds.cern.ch/record/1481838. 19, 20, 21, 159, 161, 174

[26] V Karimäki, M Mannelli, P Siegrist, H Breuker, A Caner, R Castaldi, K Freudenreich, G Hall, R Horisberger, M Huhtinen, and A Cattai. *The CMS tracker system project: Technical Design Report*. Technical Design Report CMS. CERN, Geneva, 1997. URL: http://cds.cern.ch/record/368412. 22, 174

[27] Serguei Chatrchyan et al. Description and performance of track and primary-vertex reconstruction with the CMS tracker. *JINST*, 9(10):P10009, 2014. arXiv:1405.6569, doi:10.1088/1748-0221/9/10/P10009. 24, 26, 174

[28] Serguei Chatrchyan et al. Energy Calibration and Resolution of the CMS Electromagnetic Calorimeter in *pp* Collisions at $\sqrt{s} = 7$ TeV. *JINST*, 8:P09009, 2013. [JINST8,9009(2013)]. arXiv:1306.2016, doi:10.1088/1748-0221/8/09/P09009. 30, 175

[29] The CMS Collaboration. Design, Performance and Calibration of the CMS Forward Calorimeter Wedges. Technical Report CMS-NOTE-2006-044, CERN, Geneva, Feb 2006. URL: https://cds.cern.ch/record/951395. 35, 175

[30] N Amapane, V Andreev, V Drollinger, V Karimki, V Klyukhin, and T Todorov. Volume-based Representation of the Magnetic Field. (CMS-CR-2005-011):5 p, 2005. URL: http://cds.cern.ch/record/883293. 37, 38, 175

[31] G. Acquistapace et al. CMS, the magnet project: Technical design report. 1997. 37, 175

[32] The CMS Collaboration. *The CMS muon project: Technical Design Report*. Technical Design Report CMS. CERN, Geneva, 1997. URL: https://cds.cern.ch/record/343814. 38

[33] Serguei Chatrchyan et al. Performance of CMS muon reconstruction in *pp* collision events at $\sqrt{s} = 7$ TeV. *JINST*, 7:P10002, 2012. arXiv:1206.4071, doi:10.1088/1748-0221/7/10/P10002. 43, 44, 93, 94, 176

[34] Sergio Cittolin, Attila Rácz, and Paris Sphicas. *CMS The TriDAS Project: Technical Design Report, Volume 2: Data Acquisition and High-Level Trigger. CMS trigger and data-acquisition project*. Technical Design Report CMS. CERN, Geneva, 2002. URL: http://cds.cern.ch/record/578006. 46

[35] CMS Collaboration. CMS Luminosity Based on Pixel Cluster Counting - Summer 2013 Update. Technical Report CMS-PAS-LUM-13-001, CERN, Geneva, 2013. 46, 146

[36] CMS Collaboration. CMS Luminosity Public Results. URL: https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults. 47, 48, 176

[37] CMS Collaboration. Particle-Flow Event Reconstruction in CMS and Performance for Jets, Taus, and MET. Technical Report CMS-PAS-PFT-09-001, CERN, 2009. Geneva, Apr 2009. URL: https://cds.cern.ch/record/1194487. 49, 51, 177

[38] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. The anti-k$_T$ jet clustering algorithm. *JHEP*, 2008(04):063, 2008. doi:10.1088/1126-6708/2008/04/063. 52, 78

[39] CMS Collaboration. Tau identification in CMS. 2011. 52, 61

[40] GEANT4 Collaboration. *Geant4 user's documents*. CERN, Geneva, 1998. URL: https://cds.cern.ch/record/998155. 52

[41] Johan Alwall, Michel Herquet, Fabio Maltoni, Olivier Mattelaer, and Tim Stelzer. MadGraph 5: going beyond. *JHEP*, 2011:1–40, 2011. URL: http://dx.doi.org/10.1007/JHEP06(2011)128, doi:10.1007/JHEP06(2011)128. 54, 62

[42] C F Berger, Z Bern, L J Dixon, F Febres Cordero, D Forde, T Gleisberg, H Ita, D A Kosower, and D Maitre. Multi-jet cross sections at NLO with BlackHat and Sherpa. Technical Report arXiv:0905.2735. MIT-CTP 4038. SACLAY-IPHT-T09-059. SLAC-PUB-13630. UCLA-09-TEP-48, May 2009. Comments: 4 pages, contribution to the proceedings of the XLIIIth Rencontres de Moriond (QCD). URL: https://cds.cern.ch/record/1177972. 54

[43] Simone Alioli. NLO and Parton Showers: The POWHEG-BOX. In *Physics at the LHC2010. Proceedings, 5th Conference, PLHC2010, Hamburg, Germany, June 7-12, 2010*, pages 204–208, 2010. URL: http://inspirehep.net/record/867509/files/arXiv:1009.2348.pdf, arXiv:1009.2348. 54

[44] Stefaan Tavernier. *Experimental techniques in nuclear and particle physics*. Springer, Berlin, 2009. URL: https://cds.cern.ch/record/1172614. 55

[45] CMS Collaboration. Measurement of Tracking Efficiency. 2010. 55

[46] Serguei Chatrchyan et al. Search for $B_s^0 \to \mu^+\mu^-$ and $B^0 \to \mu^+\mu^-$ decays. *JHEP*, 04:033, 2012. arXiv:1203.3976, doi:10.1007/JHEP04(2012)033. 55

[47] Georges Aad et al. Updated coupling measurements of the Higgs boson with the ATLAS detector using up to 25/fb of proton-proton collision data, 2014. ATLAS-CONF-2014-009. 58

[48] Combination of standard model Higgs boson searches and measurements of the properties of the new boson with a mass near 125 GeV, 2013. CMS PAS HIG-13-005. 58

[49] Georges Aad et al. Measurements of Higgs boson production and couplings in diboson final states with the ATLAS detector at the LHC. *Phys.Lett.*, B726:88–119, 2013. arXiv:1307.1427, doi:10.1016/j.physletb.2013.08.010. 58

[50] Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV. Technical Report ATLAS-CONF-2015-044, CERN, Geneva, Sep 2015. URL: http://cds.cern.ch/record/2052552. 59, 177

[51] Radovan Dermisek and John F. Gunion. Escaping the large fine tuning and little hierarchy problems in the next to minimal supersymmetric model and $h \to aa$ decays. *Phys.Rev.Lett.*, 95:041801, 2005. arXiv:hep-ph/0502105, doi:10.1103/PhysRevLett.95.041801. 58

[52] Radovan Dermisek and John F. Gunion. The NMSSM Close to the R-symmetry Limit and Naturalness in $h \to aa$ Decays for $m(a) < 2m(b)$. *Phys.Rev.*, D75:075019, 2007. arXiv:hep-ph/0611142, doi:10.1103/PhysRevD.75.075019. 58

[53] Spencer Chang, Radovan Dermisek, John F. Gunion, and Neal Weiner. Nonstandard Higgs Boson Decays. *Ann.Rev.Nucl.Part.Sci.*, 58:75–98, 2008. arXiv:0801.4554, doi:10.1146/annurev.nucl.58.110707.171200. 58

[54] David Curtin, Rouven Essig, Stefania Gori, Prerit Jaiswal, Andrey Katz, et al. Exotic Decays of the 125 GeV Higgs Boson. 2013. arXiv:1312.4992. 58, 59

[55] G. Belanger, B. Dumont, U. Ellwanger, J. F. Gunion, and S. Kraml. Status of invisible Higgs decays. *Phys. Lett.*, B723:340–347, 2013. arXiv:1302.5694, doi:10.1016/j.physletb.2013.05.024. 58

[56] Jeremy Bernon, John F. Gunion, Yun Jiang, and Sabine Kraml. Light Higgs bosons in Two-Higgs-Doublet Models. *Phys. Rev.*, D91(7):075019, 2015. arXiv:1412.3385, doi:10.1103/PhysRevD.91.075019. 58

[57] Alejandro Celis, Victor Ilisie, and Antonio Pich. LHC constraints on two-Higgs doublet models. *JHEP*, 1307:053, 2013. arXiv:1302.4022, doi:10.1007/JHEP07(2013)053. 58

[58] Benjamin Grinstein and Patipan Uttayarat. Carving Out Parameter Space in Type-II Two Higgs Doublets Model. *JHEP*, 1306:094, 2013. arXiv:1304.0028, doi:10.1007/JHEP06(2013)094. 58

[59] Baradhwaj Coleppa, Felix Kling, and Shufang Su. Constraining Type II 2HDM in Light of LHC Higgs Searches. 2013. arXiv:1305.0002. 58

[60] Chien-Yi Chen, S. Dawson, and Marc Sher. Heavy Higgs Searches and Constraints on Two Higgs Doublet Models. *Phys.Rev.*, D88:015018, 2013. arXiv:1305.1624, doi:10.1103/PhysRevD.88.015018. 58

[61] Nathaniel Craig, Jamison Galloway, and Scott Thomas. Searching for Signs of the Second Higgs Doublet. 2013. arXiv:1305.2424. 58

[62] Lei Wang and Xiao-Fang Han. Status of the aligned two-Higgs-doublet model confronted with the Higgs data. *JHEP*, 1404:128, 2014. arXiv:1312.4759, doi:10.1007/JHEP04(2014)128. 58

[63] Julien Baglio, Otto Eberhardt, Ulrich Nierste, and Martin Wiebusch. Benchmarks for Higgs Pair Production and Heavy Higgs Searches in the Two-Higgs-Doublet Model of Type II. 2014. arXiv:1403.1264. 58

[64] S.F. King, M. Muehlleitner, R. Nevzorov, and K. Walz. Natural NMSSM Higgs Bosons. *Nucl.Phys.*, B870:323–352, 2013. arXiv:1211.5074, doi:10.1016/j.nuclphysb.2013.01.020. 59

[65] Junjie Cao, Fangfang Ding, Chengcheng Han, Jin Min Yang, and Jingya Zhu. A light Higgs scalar in the NMSSM confronted with the latest LHC Higgs data. *JHEP*, 1311:018, 2013. arXiv:1309.4939, doi:10.1007/JHEP11(2013)018. 59

[66] Neil D. Christensen, Tao Han, Zhen Liu, and Shufang Su. Low-Mass Higgs Bosons in the NMSSM and Their LHC Implications. *JHEP*, 1308:019, 2013. arXiv:1303.2113, doi:10.1007/JHEP08(2013)019. 59

[67] David G. Cerdeno, Pradipta Ghosh, and Chan Beom Park. Probing the two light Higgs scenario in the NMSSM with a low-mass pseudoscalar. *JHEP*, 1306:031, 2013. arXiv:1301.1325, doi:10.1007/JHEP06(2013)031. 59

[68] G. Chalons and F. Domingo. Analysis of the Higgs potentials for two doublets and a singlet. *Phys.Rev.*, D86:115024, 2012. arXiv:1209.6235, doi:10.1103/PhysRevD.86.115024. 59

[69] Amine Ahriche, Abdesslam Arhrib, and Salah Nasri. Higgs Phenomenology in the Two-Singlet Model. *JHEP*, 1402:042, 2014. arXiv:1309.5615, doi:10.1007/JHEP02(2014)042. 59

[70] G. Belanger, B. Dumont, U. Ellwanger, J.F. Gunion, and S. Kraml. Global fit to Higgs signal strengths and couplings and implications for extended Higgs sectors. 2013. arXiv:1306.2941. 59

[71] Jeremy Bernon, Beranger Dumont, and Sabine Kraml. Status of Higgs couplings after run 1 of the LHC. *Phys. Rev.*, D90:071301, 2014. arXiv:1409.1588, doi:10.1103/PhysRevD.90.071301. 59

[72] CMS Collaboration. Constraints on the Higgs boson width from off-shell production and decay to $ZZ \to \ell\ell\ell'\ell'$ and $\ell\ell\nu\nu$. 60

[73] Radovan Dermisek and John F. Gunion. Direct production of a light CP-odd Higgs boson at the Tevatron and LHC. *Phys.Rev.*, D81:055001, 2010. arXiv:0911.2460, doi:10.1103/PhysRevD.81.055001. 60

[74] Beranger Dumont, John F. Gunion, Yun Jiang, and Sabine Kraml. Constraints on and future prospects for Two-Higgs-Doublet Models in light of the LHC Higgs signal. 2014. arXiv:1405.3584. 60

[75] K. A. Olive et al. Review of Particle Physics. *Chin. Phys.*, C38:090001, 2014. doi:10.1088/1674-1137/38/9/090001. 61

[76] CMS Collaboration. CMS HLT configurations. URL: http://j2eeps.cern.ch/cms-project-confdb-hltdev/browser/. 62

[77] CMS Collaboration. PREP: Production and Reprocessing management tool. URL: http://cms.cern.ch/iCMS/jsp/mcprod/admin/requestmanagement.jsp?pwg=JME&campid=Summer12. 63, 191

[78] CMS Collaboration. Standard Model CrossSections at 8TeV. URL: https://twiki.cern.ch/twiki/bin/viewauth/CMS/StandardModelCrossSectionsat8TeV. 63, 191

[79] Torbjörn Sjöstrand, Stephen Mrenna, and Peter Skands. Pythia 6.4 physics and manual. *JHEP*, 2006(05):026, 2006. doi:doi:10.1088/1126-6708/2006/05/026. 62

[80] LHC Higgs Cross Sections Working Group. SM Higgs production cross sections at $\sqrt{s} = 8$ TeV. URL: https://twiki.cern.ch/twiki/bin/view/LHCPhysics/CERNYellowReportPageAt8TeV. 64, 150

[81] G. Bozzi, S. Catani, D. de Florian, and M. Grazzini. The q(T) spectrum of the Higgs boson at the LHC in QCD perturbation theory. *Phys.Lett.*, B564:65–72, 2003. arXiv:hep-ph/0302104, doi:10.1016/S0370-2693(03)00656-7. 65

[82] Grazzini, M. HqT version 2.0. URL: http://theory.fi.infn.it/grazzini/codes/note20.pdf. 65

[83] https://twiki.cern.ch/twiki/bin/viewauth/CMS/MuonHLT. URL: https://twiki.cern.ch/twiki/bin/viewauth/CMS/MuonHLT#SingleMu_Efficiency. 68, 93

[84] Serguei Chatrchyan et al. Measurement of the properties of a Higgs boson in the four-lepton final state. *Phys. Rev.*, D89(9):092007, 2014. arXiv:1312.5353, doi:10.1103/PhysRevD.89.092007. 69

[85] CMS Collaboration. J/Psi prompt and non-prompt cross sections in pp collisions at sqrt(s) = 7 TeV. 2010. 69, 77, 78

[86] CMS. Performance of CMS muon reconstruction in pp collision events at $\sqrt{s}$ = 7 TeV. *Journal of Instrumentation*, 7(10):P10002, 2012. doi:doi:10.1088/1748-0221/7/10/P10002. 72, 75, 95

[87] M. Bachtis, S. Dasu, and A. Savin. Prospects for measurement of $\sigma(pp \to Z) \cdot B(Z \to \tau^+\tau^-)$ with CMS in pp collisions at $\sqrt{s} = 7$ TeV. CMS Note 2010/082, CERN, 2010. URL: http://cms.cern.ch/iCMS/jsp/openfile.jsp?tp=draft&files=AN2010_082_v1.pdf. 79

[88] https://twiki.cern.ch/twiki/bin/view/CMSPublic/SWGuidePFTauID. URL: https://twiki.cern.ch/twiki/bin/view/CMSPublic/SWGuidePFTauID#Tau_ID_2014_preparation_for_AN1. 79

[89] Christian Weiser. A Combined Secondary Vertex Based B-Tagging Algorithm in CMS. Technical Report CMS-NOTE-2006-014, CERN, Geneva, Jan 2006. 82

[90] https://twiki.cern.ch/twiki/bin/viewauth/CMS/BTagPerformanceOP. URL: https://twiki.cern.ch/twiki/bin/viewauth/CMS/BTagPerformanceOP#B_tagging_Operating_Points_for_5. 82

[91] https://twiki.cern.ch/twiki/bin/viewauth/CMS/BTagSFMethods. URL: https://twiki.cern.ch/twiki/bin/viewauth/CMS/BTagSFMethods. 82, 149

[92] Vardan Khachatryan et al. Reconstruction and identification of ? lepton decays to hadrons and ?$_?$ at CMS. *JINST*, 11(01):P01019, 2016. arXiv:1510.07488, doi:10.1088/1748-0221/11/01/P01019. 84, 95, 147

[93] S. Chatrchyan et al. Determination of jet energy calibration and transverse momentum resolution in CMS. *JINST*, 6(11):P11002, 2011. doi:10.1088/1748-0221/6/11/P11002. 85, 120, 185

[94] CMS Collaboration. MET Systematics Tools Twiki. URL: https://twiki.cern.ch/twiki/bin/view/CMSPublic/SWGuidePATTools#MET_Systematics_Tools. 85

[95] https://twiki.cern.ch/twiki/bin/view/Main/EGammaScaleFactors2012. URL: https://twiki.cern.ch/twiki/bin/view/Main/EGammaScaleFactors2012#2012_8_TeV_Jan22_Re_recoed_data. 94, 148

[96] https://twiki.cern.ch/twiki/bin/viewauth/CMS/MuonReferenceEffs. URL: https://twiki.cern.ch/twiki/bin/viewauth/CMS/MuonReferenceEffs#Note_on_uncertainties. 94, 147, 148

[97] https://twiki.cern.ch/twiki/bin/viewauth/CMS/TauIDRecommendation. URL: https://twiki.cern.ch/twiki/bin/viewauth/CMS/TauIDRecommendation#Systematic_Uncertainties. 94, 147, 148

[98] https://twiki.cern.ch/twiki/bin/view/CMS/MuonReferenceEffs. URL: https://twiki.cern.ch/twiki/bin/view/CMS/MuonReferenceEffs. 95, 147

[99] CMS Particle Flow and Tau Identification Results. URL: https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsPFT. 97

[100] CMS Collaboration. Pileup Reweighting Utilities Twiki. URL: https://twiki.cern.ch/twiki/bin/viewauth/CMS/PileupMCReweightingUtilities. 110

[101] P. Ko, Jungil Lee, and Chaehyun Yu. Inclusive double-quarkonium production at the Large Hadron Collider. *JHEP*, 2011(1):70, 2011. doi:doi:10.1007/JHEP01(2011)070. 131

[102] Li Gang, Ma Wen-Gan, Song Mao, Zhang Ren-You, and Guo Jian-You. Associated production of Υ(1S)W at LHC in next-to-leading order QCD. *Journal of High Energy Physics*, 2013(1):34, 2013. doi:doi:10.1007/JHEP01(2013)034. 131

[103] Q&A for HIG-14-022. URL: https://twiki.cern.ch/twiki/bin/viewauth/CMS/HIG14022QA. 140

[104] Procedure for the LHC Higgs boson search combination in Summer 2011. Technical Report CMS-NOTE-2011-005. ATL-PHYS-PUB-2011-11, CERN, Geneva, Aug 2011. URL: https://cds.cern.ch/record/1379837. 145

[105] A. L. Read. Modified frequentist analysis of search results (the CL_s method). In L. Lyons, Y. Perrin, and F.E. James, editors, *Proceedings of the First Workshop on Confidence Limits, Geneva, 2000*, page 81, Geneva, 2000. CERN. URL: https://cdsweb.cern.ch/record/411537. 145

[106] CMS Collaboration. Higgs Analysis Combined Limit Twiki. URL: https://twiki.cern.ch/twiki/bin/viewauth/CMS/SWGuideHiggsAnalysisCombinedLimit. 145, 150

[107] CMS Collaboration. Higgs Analysis Combined Limit GitHub Page. URL: https://github.com/cms-analysis/HiggsAnalysis-CombinedLimit. 150

[108] Glen Cowan, Kyle Cranmer, Eilam Gross, and Ofer Vitells. Asymptotic formulae for likelihood-based tests of new physics. *Eur. Phys. J. C*, 71:1–19, 2011. doi:10.1140/epjc/s10052-011-1554-0. 150

[109] Jesse Thaler and Ken Van Tilburg. Identifying Boosted Objects with N-subjettiness. *JHEP*, 1103:015, 2011. arXiv:1011.2268, doi:10.1007/JHEP03(2011)015. 155

[110] Christoph Englert, Tuhin S. Roy, and Michael Spannowsky. Ditau jets in Higgs searches. *Phys.Rev.*, D84:075026, 2011. arXiv:1106.4545, doi:10.1103/PhysRevD.84.075026. 156

[111] M. Case, M. Liendl, and M. Lingen. XML based Detector Description Language. CMS Note 2005/000, CERN, 2005. URL: http://cmsdoc.cern.ch/cms/software/ddd/www/draft1.3_cms_ddl.pdf. 158

[112] https://twiki.cern.ch/twiki/bin/viewauth/CMS/TrackerMaterial2015. URL: https://twiki.cern.ch/twiki/bin/viewauth/CMS/TrackerMaterial2015. 165, 189

[113] Souvik Das, Karl Ecklund, Ben Kreis, Anders Ryd, Steve Stroiney, and Josh Thompson. CMS Pixel Online Software and Calibrations. URL: https://cms-docdb.cern.ch/cgi-bin/DocDB/RetrieveFile?docid=2226&filename=UsersGuide.pdf&version=1. 169