

Search for associated production of Z and Higgs bosons in
proton-antiproton collisions at 1.96 TeV.

John Alexander BackusMayes

A dissertation submitted in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

University of Washington

2010

Program Authorized to Offer Degree: Physics

University of Washington
Graduate School

This is to certify that I have examined this copy of a doctoral dissertation by

John Alexander BackusMayes

and have found that it is complete and satisfactory in all respects,
and that any and all revisions required by the final
examining committee have been made.

Chair of the Supervisory Committee:

Gordon Watts

Reading Committee:

Gordon Watts

Henry Lubatti

Stephen Ellis

Date: _____

In presenting this dissertation in partial fulfillment of the requirements for the doctoral degree at the University of Washington, I agree that the Library shall make its copies freely available for inspection. I further agree that extensive copying of the dissertation is allowable only for scholarly purposes, consistent with "fair use" as prescribed in the U.S. Copyright Law. Requests for copying or reproduction of this dissertation may be referred to ProQuest Information and Learning, 300 North Zeeb Road, Ann Arbor, MI 48106-1346, 1-800-521-0600, to whom the author has granted "the right to reproduce and sell (a) copies of the manuscript in microform and/or (b) printed copies of the manuscript made from microform."

Signature_____

Date_____

University of Washington

Abstract

Search for associated production of Z and Higgs bosons in proton-antiproton collisions at 1.96 TeV.

John Alexander BackusMayes

Chair of the Supervisory Committee:
Associate Professor Gordon Watts
Physics

We present a search for associated production of Z and Higgs bosons in 4.2 fb^{-1} of $\bar{p}p$ collisions at $\sqrt{s} = 1.96 \text{ TeV}$, produced in RunII of the Tevatron and recorded by the $D\bar{O}$ detector. The search is performed in events containing at least two muons and at least two jets. The ZH signal is distinguished from the expected backgrounds by means of multivariate classifiers known as random forests. Binned random forest output distributions are used in comparing the data to background-only and signal+background hypotheses. No excess is observed in the data, so we set upper limits on ZH production with a 95% confidence level.

TABLE OF CONTENTS

	Page
List of Figures	iv
List of Tables	vi
Chapter 1: Introduction	1
1.1 Fundamental Particles	3
1.1.1 Fermions	3
1.1.2 Bosons	6
1.2 The Standard Model	7
1.2.1 Quantum Chromodynamics	8
1.2.2 The Electroweak Theory	10
1.3 Higgs in Experiment	16
Chapter 2: Experimental Apparatus	20
2.1 Acceleration	21
2.1.1 Proton Beam	21
2.1.2 Antiproton Beam	26
2.1.3 The Tevatron	28
2.2 Collision	30
2.2.1 Cross Section and Luminosity	30
2.2.2 Coordinate System	32
2.3 Detection	34
2.3.1 Inner Tracking	36
2.3.2 Calorimetry	41
2.3.3 Muon Tracking	46
2.3.4 Luminosity Measurement	48
2.3.5 Triggering	49

Chapter 3:	Reconstruction	53
3.1	Tracking	54
3.2	The Primary Vertex	56
3.3	Muons	58
3.4	Jets	59
3.4.1	Identification	60
3.4.2	Selection	61
3.4.3	Calibration	62
3.5	b Tagging	66
Chapter 4:	Simulation	69
4.1	Monte Carlo Event Generators	71
4.1.1	PYTHIA	71
4.1.2	ALPGEN	71
4.2	Detector Simulation	73
4.2.1	GEANT	73
4.2.2	Minimum Bias Overlay	73
4.2.3	Triggering	74
4.3	Corrections to Reconstructed Objects	75
4.3.1	Primary Vertex	75
4.3.2	Muons	75
4.3.3	Jets	76
4.3.4	b Tagging	77
Chapter 5:	Analysis	78
5.1	Data	81
5.2	Pre-Tag Selection	82
5.2.1	$Z \rightarrow \mu^+ \mu^-$	82
5.2.2	Jets	84
5.3	Monte Carlo	85
5.4	Multijet Sample	90
5.5	Background Normalization	91
5.6	b Jet Selection	94
5.6.1	Taggability	94
5.6.2	Event Weights from TRF's	95
5.6.3	Tagging in the Multijet Sample	101

5.7	Event Yields	104
5.8	Kinematic Fit	107
5.9	Multivariate Classifier	109
5.9.1	Decision Trees	110
5.9.2	The Random Forest	111
5.9.3	Input Variables	112
5.9.4	Performance	115
5.10	Systematic Uncertainties	118
5.10.1	Flat Systematics	119
5.10.2	Shape Systematics	120
5.11	Limit Setting	126
Chapter 6:	Conclusion	133
Bibliography	135

LIST OF FIGURES

Figure Number	Page
1.1 Indirect constraints on the Higgs mass	17
1.2 Higgs production cross sections	18
1.3 Branching ratios for Higgs decays	19
2.1 The accelerator complex at Fermilab [8].	21
2.2 The magnetron H^- ion source [9].	22
2.3 Fermilab's Drift Tube Linac	23
2.4 Fermilab's Side-Coupled Linac	24
2.5 The DØ detector	35
2.6 DØ's innermost subdetectors [12].	36
2.7 DØ's silicon microstrip tracker	38
2.8 Cellular structure of the sampling calorimeter	44
2.9 Pseudo-projective calorimeter cell towers	45
2.10 Flow of data through the DØ trigger system [12].	50
3.1 Jet energy scale correction	64
3.2 Output of the neural-net b -tagger	67
3.3 b -tagging efficiency and misidentification rate	68
4.1 Distribution of primary vertices in the z direction	76
5.1 Dimuon invariant mass spectrum	83
5.2 Dijet invariant mass spectrum	84
5.3 Distribution of Z candidate p_T	86
5.4 Angular distributions of jets	87
5.5 Corrections from single-muon to inclusive triggering	88
5.6 Muon pseudorapidity distributions	89
5.7 Taggability in RunIIb	96
5.8 Anti-taggability in RunIIb	97
5.9 p_T dependence of taggability in RunIIa	98
5.10 Taggability in RunIIa	99

5.11	Anti-tagability in RunIIa	100
5.12	Dijet invariant mass distributions before and after the kinematic fit	108
5.13	Distributions of random forest input variables	114
5.14	Pre-tag RF output distributions	116
5.15	Post-tag RF output distributions	117
5.16	Observed and expected LLR values, $\mu\mu$ leptonic channel only	128
5.17	Upper limits on ZH production, $\mu\mu$ leptonic channel only	130
5.18	Observed and expected LLR values, combining all $ZH \rightarrow \ell\ell b\bar{b}$ channels	131
5.19	Upper limits on ZH production, combining all $ZH \rightarrow \ell\ell b\bar{b}$ channels	131

LIST OF TABLES

Table Number	Page
1.1 Properties of the fundamental fermions	4
1.2 Properties of the fundamental bosons	6
5.1 NLO correction factors to the simulated backgrounds	85
5.2 Combined normalization fit results and statistical uncertainties.	93
5.3 2L b -tagging probabilities	102
5.4 1T b -tagging probabilities	102
5.5 Scale factors for indirect b tagging in the multijet sample	103
5.6 Event yields before and after b tagging	105
5.7 Expected S/B and S/\sqrt{B} , before and after b tagging	106
5.8 Theoretical cross-section uncertainties.	120
5.9 Shape systematics in 1T-tagged events	124
5.10 Shape systematics in 2L-tagged events	125
5.11 Upper limits on ZH production, $\mu\mu$ leptonic channel only	130
5.12 Upper limits on ZH production, combining all $ZH \rightarrow \ell\ell b\bar{b}$ channels	132

Chapter 1

INTRODUCTION

Particle physics is concerned with discovering and studying the most fundamental pieces of the universe. We use the word *fundamental* in the reductionist sense: we hypothesize that everything in the universe can be understood as an interacting collection of these fundamental pieces, called particles. Theoretically, it should be possible to model even macroscopic objects in terms of their constituent particles, but usually such models are prohibitively complicated or mathematically intractable. Thus, particle physics is used to predict the behavior of physical systems containing a relatively small number of particles interacting over distances smaller than the size of an atomic nucleus. These particles are described in section 1.1.

The quantum field theory that governs the dynamics of the fundamental particles is known as the *standard model* (SM), described in section 1.2. The SM is arguably one of the most successful theories in modern physics, or even in modern science. Among many other experimentally verified predictions, the SM predicted the existence of the Z boson and the top quark, both required by the local symmetries that are at the heart of our understanding of particle physics. In the many years since the development of the standard model, it is quite remarkable that every result from collider experiments has been in agreement with the SM, though some astrophysical observations and the discovery of neutrino oscillations provide evidence that the SM is not complete.

The only SM particle that has yet to be discovered is the Higgs boson, which is central to the standard model's explanation of nonzero weak boson and fermion masses. Excluding the Higgs boson would be tantamount to excluding the SM, and could well lead to an indication of physics beyond the SM. Already, there is some tension between indirect and direct constraints on the mass of the Higgs, as described in section 1.3. On the other hand, discovering the Higgs boson would be yet another feather in the standard model's cap,

and further study of the properties of the Higgs could also point the way towards a more complete particle theory.

If Higgs bosons exist, then rather extreme experimental conditions are required to produce them. The largest particle accelerators in the world are required to achieve energy densities high enough to have a non-negligible chance of directly producing a Higgs boson. The research described in this dissertation is an analysis of proton-antiproton collisions produced in the Tevatron collider at Fermilab, with a center-of-mass energy of 1.96 TeV. Large, sophisticated detectors are required for precise measurement of the products of these collisions, and equally sophisticated data processing is needed to record and analyze the signals produced by such detectors. The data analyzed for this dissertation were collected with the DØ detector, one of two detectors surrounding collision points at the Tevatron. Chapter 2 describes Fermilab's particle accelerators and the DØ detector in detail.

In chapter 3, I describe the algorithms used to transform raw detector signals into measurements of particle kinematics, along with the methods by which different types of particles are identified. Chapter 4 describes how simulated data is obtained from Monte Carlo techniques. Finally, all of these ingredients are brought together in a search for associated production of Higgs and Z bosons, presented in chapter 5. Sophisticated multivariate techniques are used to maximize our sensitivity to the Higgs signal while accepting as many potential signal events as possible. As no significant excess is observed in the data, we extract upper limits on ZH production by comparing the observed data to the simulated signals and backgrounds.

1.1 Fundamental Particles

In this section, I present the particles currently hypothesized to be the fundamental building blocks of the universe. All of these particles except the Higgs boson have been directly detected in one or more experiments, and experimental results thus far are consistent with each of these particles being truly fundamental, i.e. not composite. It should be noted that many of these particles are not in fact “particles” in the strictest sense, as they are unstable (e.g. the μ or τ leptons) or cannot be observed alone (e.g. the quarks and gluons). We use the word *particle* to refer to the quanta from which we build variations in the fundamental dynamic fields of the universe.

There are actually a few phenomena which cannot be described in terms of the standard-model particles and fields. Most notably, gravity is entirely absent from the theory, but fortunately the effect of gravity is generally negligible in the small-scale physical systems we study. Also apparently missing are explanations for astrophysical dark matter and neutrino oscillations. These phenomena strongly suggest that the standard model is incomplete, and there are ongoing experimental and theoretical efforts to extend the theory. The search for the last undiscovered SM particle, the Higgs boson, may help to rule out some hypothetical SM extensions and to indicate better the path toward a more complete theory of fundamental particles.

1.1.1 Fermions

Fermions are particles with intrinsic angular momentum, or *spin*, in half-integer multiples of \hbar . Equivalently, they are the quanta of fields that change sign under spatial rotations of 2π radians. As a consequence of this, fermions obey the Pauli exclusion principle, and systems of many identical fermions are described by Fermi-Dirac statistics.

All fundamental fermions in the standard model have spin $1/2$, so they transform according to the fundamental representation of the spatial rotation group $SU(2)$. The standard model is a chiral theory, so left- and right-handed fermions must often be treated separately, transforming as Weyl spinors under the restricted Lorentz group omitting discrete parity (P) and time-reversal (T) symmetries. For convenience in calculations where P or T in-

	Y	T	T_3	Q	Mass		
ν_L	-1	1/2	1/2	0	0	0	0
ℓ_L^-			-1/2				
ℓ_R^-	-2	0	0	-1	511 keV	106 MeV	1.78 GeV
u_R	4/3	0	0				
u_L			1/2	2/3	[1.7, 3.3] MeV	[1.18, 1.34] GeV	172 GeV
d_L	1/3	1/2	-1/2				
d_R	-2/3	0	0	-1/3	[4.1, 5.8] MeV	[80, 130] MeV	[4.13, 4.85] GeV

Table 1.1: Properties of the fundamental fermions. Note that each quark listed is actually a triplet of quarks, distinguished only by “color.” Additionally, corresponding to every fermion is an anti-fermion with opposite charge. Definitions of Y , T , T_3 , and Q are given in the text.

variance should be manifest, the left- and right-handed components may be combined into a Dirac spinor.

Table 1.1 shows the quantum numbers and masses that distinguish the 21 different fundamental fermions (ignoring the three quark “colors” and the anti-fermions) in the SM. Y is the weak hypercharge, and T_3 is the third component of the weak isospin T , which distinguishes between left- and right-handed fermions. The electric charge Q is equal to $T_3 + Y/2$ and distinguishes between up- and down-type fermions as well as between leptons and quarks. The three columns of masses correspond to the three generations of SM fermions, in order of increasing mass.

Leptons

Each of the three generations of leptons contains left-handed (ℓ_L) and right-handed (ℓ_R) leptons with charge -1 and an electrically neutral left-handed neutrino ν_L . The charged lepton is massive, while the neutrino is massless, as nonzero neutrino mass would require the existence of right-handed neutrinos. Actually, recent observations of neutrino flavor oscillations imply nonzero neutrino masses, indicating that right-handed neutrinos may

exist despite the lack of direct experimental evidence for them. For the purposes of my research, we can safely ignore any small but nonzero neutrino mass, as we do not expect this to significantly affect Higgs production at the Tevatron.

The first generation of leptons contains the electron, an essential component of all atomic matter in the universe. The charged leptons in the other two generations are identical to the electron in every aspect except their mass. The muon is more than 200 times more massive than the electron, and the tau lepton is almost 17 times as massive as the muon. The muon and tau lepton are unstable and decay to particles of smaller mass via the weak nuclear interactions. With a mean lifetime of approximately $2.2 \mu\text{s}$, muons decay almost entirely in the process $\mu \rightarrow e\bar{\nu}_e\nu_\mu$. The larger mass of the tau lepton makes many other decay modes kinematically available, which reduces its mean lifetime to roughly 0.3 ps.

Quarks

The other type of fundamental SM fermion, quarks are the constituents of hadrons, such as the protons and neutrons of which every atomic nucleus is composed. In fact, the strong nuclear interaction actually *confines* quarks inside hadrons. Heuristically, as a quark is removed from a hadron, the force coupling the quark to that hadron increases, similar to the stretching of a strip of rubber. Eventually, it is energetically favorable for the rubber strip to break, which results in two unstretched strips, representing two hadrons. Therefore, any quark we attempt to remove from a hadron becomes confined in a new hadron. Because of confinement, the mass of an individual quark is very difficult to measure. In table 1.1, the 68% confidence intervals for quark masses are shown.

Ordered by mass, the up-type quarks with charge $2/3$ are known as up (u), charm (c), and top (t). The down-type quarks with charge $-1/3$ are known as down (d), strange (s), and bottom (b). Altogether, the quark masses range over five orders of magnitude. The top quark is the most massive, being famously as heavy as a gold nucleus. Such a large mass makes the top quark unique as the only quark that is not a constituent of hadrons: its mean lifetime is significantly less than the time required for hadronization. Because of this, the mass of the top quark can be precisely measured at high-energy collider experiments such

	Spin	Y	T	T_3	Q	Mass (GeV)	Couples to:
G	1	0	0	0	0	0	self and quarks
W^\pm	1	0	1	± 1	± 1	80.4	all particles with $T > 0$
Z	1	0	$0 \oplus 1$	0	0	91.2	W^\pm and all fermions
A	1	0	$0 \oplus 1$	0	0	0	all particles with $Q > 0$
H	0	1	1/2	-1/2	0	$[114.4, 158] \cup [175, \infty)$	all particles with mass

Table 1.2: Properties of the fundamental bosons. Note that G is actually an octet of gluons, distinguished only by “color.” Definitions of Y , T , T_3 , and Q are given in the text.

as $D\emptyset$.

1.1.2 Bosons

Bosons are particles with spin in integer multiples of \hbar . Equivalently, they are the quanta of fields invariant under spatial rotations of 2π radians. As a consequence of this, systems of many identical bosons are described by Bose-Einstein statistics.

Fundamental bosons in the standard model have spin 0 or 1. The spin-1 bosons transform according to the fundamental representation of the spatial rotation group $SO(3)$ and as four-vectors under Lorentz transformations. The spin-0 field, known as the Higgs field, transforms as a scalar under rotations and Lorentz boosts.

Table 1.2 shows the properties that distinguish the six fundamental bosons (ignoring the eight gluon “colors”) in the SM. The gluon G is the massless boson which mediates the strong nuclear interaction that binds quarks into hadrons and hadrons into nuclei. The W and Z bosons mediate the weak nuclear interaction that facilitates radioactive β -decay and many other phenomena. The photon A is the massless boson which mediates electromagnetic interactions. Finally, the Higgs boson H is the massive, spin-0 particle predicted by the Higgs mechanism of electroweak symmetry breaking, which provides mass to the W and Z bosons.

1.2 The Standard Model

Starting with the fundamental spin-1/2 fields corresponding to the quarks and leptons, we derive the standard model by requiring that the Lagrangian (from which we derive equations of motion) be invariant under global Lorentz transformations and local $U(1)$, $SU(2)$, and $SU(3)$ gauge transformations. To begin, I present a simplified model containing one spin-1/2 fermion ψ that transforms according to the fundamental representation of the gauge group $SU(N)$. The simplest invariant Lagrangian that describes this theory is shown in equation 1.1.

$$\mathcal{L} = -\frac{1}{4}(F_{\mu\nu}^a)^2 + \bar{\psi}(i\not{D} - m)\psi \quad (1.1)$$

The fermion ψ is a Dirac spinor, and $\bar{\psi}$ is the hermitian conjugate of ψ multiplied by the zeroth Dirac matrix γ^0 . An $SU(N)$ transformation affects the field ψ as shown in equation 1.2.

$$\psi(x) \rightarrow U(x)\psi(x) = e^{i\alpha^a(x)t^a}\psi(x) \quad (1.2)$$

Since ψ transforms according to the fundamental representation, U is a unitary $N \times N$ matrix of unit determinant. We can parameterize U in terms of $N^2 - 1$ real numbers α^a . Related to the elements of $SU(N)$ through exponentiation, t^a are the $N^2 - 1$ generators of the group $SU(N)$. Acting on ψ , t^a are hermitian, traceless $N \times N$ matrices. The generators t^a must satisfy the commutation relations given in equation 1.3, where f^{abc} are structure constants unique to the group. f^{abc} is odd under exchange of any two of the indices a , b , and c .

$$[t^a, t^b] = if^{abc}t^c \quad (1.3)$$

The derivative operator compares the values of a field at closely separated points in spacetime, but gauge transformations may be entirely unrelated from point x to point $x + \epsilon$. To ensure that the derivative of ψ transforms in the same way as ψ itself, we must redefine the derivative by adding real-valued fields that account for the derivative of the

gauge transformation. For $SU(N)$, we add $N^2 - 1$ of these fields A_μ^a , each associated with a generator t^a . The resulting *covariant derivative* D_μ is defined in equation 1.4. We must specify an arbitrary coupling constant g to uniquely define the field A_μ^a , and this constant must be experimentally determined.

$$\not{D} = \gamma^\mu D_\mu = \gamma^\mu (\partial_\mu - igA_\mu^a t^a) \quad (1.4)$$

Because they are added to the derivative operator ∂_μ , each field A_μ^a must be a four-vector under Lorentz transformations, so we identify these fields as massless, spin-1 bosons. The variation of a gauge transformation from point to point is taken into account through the transformation law for A_μ^a , given in equation 1.5.

$$A_\mu^a(x)t^a \rightarrow U(x) \left(A_\mu^a(x)t^a + \frac{i}{g}\partial_\mu \right) U^\dagger(x) \quad (1.5)$$

Finally, having introduced new dynamic fields in our modification of the derivative operator, it is necessary to add kinetic terms for these fields to the Lagrangian. The proper form for these terms is $-(F_{\mu\nu}^a)^2/4$, where $F_{\mu\nu}^a$ is defined as in equation 1.6. This form is manifestly Lorentz-invariant, and it can be shown to be invariant under $SU(N)$ gauge transformations as well.

$$F_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a + gf^{abc}A_\mu^b A_\nu^c \quad (1.6)$$

Thus, we have constructed an $SU(N)$ gauge-invariant theory including fundamental gauge bosons A_μ^a and a fundamental fermion ψ . Interactions between fermions and bosons arise from the covariant derivative, and (for $N > 1$) the kinetic terms for A_μ^a result in interactions between bosons. This theory only has two free parameters: the mass m of the fermion and the coupling constant g in the covariant derivative. As we add more fermions and multiple gauge symmetries, we arrive at the 19 free parameters of the standard model.

1.2.1 Quantum Chromodynamics

Quantum chromodynamics (QCD) is the piece of the standard model that describes the strong nuclear interaction. It is derived by imposing an $SU(3)$ gauge symmetry on the SM

Lagrangian, where quarks transform as three-component complex vectors according to the fundamental representation of the $SU(3)$ group. As there are likewise three primary colors, we poetically refer to these three quark components as red, green, and blue (hence the name *chromodynamics*). Similarly, there are three colors of anti-quarks: anti-red, anti-green, and anti-blue. Since the leptons are not involved in the strong interactions, leptonic SM fields transform in the trivial representation of $SU(3)$.

The $SU(3)$ gauge symmetry of QCD requires $3^2 - 1 = 8$ massless, spin-1 bosons, called *gluons*. The relevant interaction terms from the resultant Lagrangian are shown in equation 1.7. The first term describes the interaction between a quark q , an anti-quark \bar{q} , and a gluon G , with interaction strength proportional to the QCD coupling constant g_{QCD} . The interaction between three gluons, described by the second term, is also linear in g_{QCD} , while the four-gluon interaction in the third term is quadratic in g_{QCD} .

$$\begin{aligned} \mathcal{L}_{\text{int}}^{\text{QCD}} = & g_{\text{QCD}} \bar{q} \gamma^\mu G_\mu^a t_{\text{QCD}}^a q \\ & - \frac{g_{\text{QCD}}}{2} f_{\text{QCD}}^{abc} (\partial_\mu G_\nu^a - \partial_\nu G_\mu^a) G^{b,\mu} G^{c,\nu} \\ & - \frac{g_{\text{QCD}}^2}{4} f_{\text{QCD}}^{abc} f_{\text{QCD}}^{ade} G_\mu^b G_\nu^c G^{d,\mu} G^{e,\nu} \end{aligned} \quad (1.7)$$

Two important phenomenological properties of QCD are not immediately apparent from the fundamental Lagrangian: confinement and asymptotic freedom. As mentioned in the previous section, in low-energy effective QCD, quarks are confined inside colorless hadrons. Although confinement in QCD has escaped rigorous mathematical proof, it is consistent with all experimental observations, and theoretical results using nonperturbative methods are in strong support of these observations. While interactions in a $\bar{p}p$ collision may result in the production of individual quarks or gluons, these particles become confined inside hadrons before we can measure their trajectories, and it is nontrivial to infer the kinematics of the original quarks and gluons from our measurements of the resulting hadrons.

A further complication arises from a singularity in the theory for the emission of soft or collinear gluons from any colored particle. Before hadronization, a quark or gluon emits an indeterminate number of gluons, which in turn may split into quark-antiquark pairs,

and so on, producing a collimated jet of colored particles. The resulting quarks and gluons hadronize to form a jet of colorless particles, and we measure the total momentum of all of these hadrons to estimate the momentum of the original, final-state quark or gluon.

Confinement relies on the strong, nonperturbative nature of QCD at low energies, which is related to the existence of self-interactions between massless gluons. Another result of these self-interactions is that the effective strength of QCD interactions actually decreases at high energy scales. When the momenta involved in an interaction are particularly high, QCD becomes perturbative; eventually, at extremely high momenta, quarks and gluons behave as free, non-interacting particles. We rely heavily on this property, known as asymptotic freedom, as it allows us to treat the constituents of a proton or antiproton as effectively free. An interaction involving only one quark or gluon each from the proton and antiproton occurs independently of the other “spectator” quarks and gluons present in a $\bar{p}p$ collision.

1.2.2 The Electroweak Theory

Gauge symmetries result in massless gauge bosons such as the gluon, so at first it seems unlikely that the weak interaction could be derived from gauge invariance: the short range of weak interactions strongly suggests they are mediated by massive spin-1 bosons. Also, the bosons which mediate weak interactions must carry electric charge, indicating that the weak and electromagnetic interactions cannot be treated independently in the standard model. Indeed, these two fundamental forces are described by a single *electroweak* (EW) theory in the SM, in which fundamental $SU(2)$ and $U(1)$ gauge symmetries are spontaneously broken and mixed through the Higgs mechanism. The EW theory correctly predicted the existence of the Z boson, which was discovered by the UA1 and UA2 experiments in 1983. Existence of the Higgs boson, also predicted by the EW theory, has not yet been experimentally confirmed or excluded.

In the electroweak theory, the $SU(2)$ gauge transformations affect only left-handed fermions, while right-handed fermions are unaffected. Thus, we refer to this chiral $SU(2)$ gauge symmetry as $SU(2)_L$, L indicating left chirality. Within each SM generation, up- and down-type left-handed fermions form the two components of a weak isospin doublet,

distinguished by the third component of weak isospin, T_3 . Therefore, an arbitrary $SU(2)_L$ transformation may “rotate” neutrinos into charged leptons and up-type quarks into down-type quarks. Both left- and right-handed fermions are affected by the $U(1)$ gauge transformations, which are just local changes in the complex phases of the fields. The amount of phase change is determined by the hypercharge Y associated with each field, so we refer to this gauge symmetry as $U(1)_Y$. Under an arbitrary $SU(2)_L \times U(1)_Y$ transformation, parametrized by α^a and β , left- and right-handed fermions ψ_L and ψ_R transform according to equations 1.8 and 1.9, respectively. The matrices t_W^a are the generators of the $SU(2)_L$ gauge group.

$$\psi_L \rightarrow e^{i\alpha^a t_W^a} e^{i\beta Y_L/2} \psi_L \quad (1.8)$$

$$\psi_R \rightarrow e^{i\beta Y_R/2} \psi_R \quad (1.9)$$

Equation 1.10 shows the resultant Lagrangian for the electroweak theory. The first two terms are the kinetic terms for the $U(1)_Y$ gauge field B and the $SU(2)_L$ gauge fields W^a . The third and fourth terms are the kinetic terms for the right- and left-handed fermions, from which interactions between fermions and gauge bosons are derived. The coupling constants g_B and g_W are free parameters that must be determined by experiment. With only these four terms, we have a complete theory of independent $SU(2)_L$ and $U(1)_Y$ gauge interactions mediated by massless spin-1 bosons W^a and B . Note that the fermions in this theory are also massless: since ψ_L and ψ_R transform differently, Dirac mass terms for fermions are not gauge-invariant.

$$\begin{aligned} \mathcal{L}^{\text{EW}} = & -\frac{1}{4}(B_{\mu\nu})^2 - \frac{1}{4}(W_{\mu\nu}^a)^2 \\ & + \bar{\psi}_R^i \gamma^\mu \left(\partial_\mu - ig_B \frac{Y_R^i}{2} B_\mu \right) \psi_R^i \\ & + \bar{\psi}_L^i \gamma^\mu \left(\partial_\mu - ig_B \frac{Y_L^i}{2} B_\mu - ig_W W_\mu^a t_W^a \right) \psi_L^i \\ & + \mathcal{L}_\phi \end{aligned} \quad (1.10)$$

To better describe the world in which we live, where fermions and weak gauge bosons are

massive, we rely on additional terms, denoted \mathcal{L}_ϕ , that describe the dynamics of a complex spin-0 field ϕ . Known as the Higgs field, ϕ is a weak isospin doublet and has hypercharge $Y = 1$. Equation 1.11 shows the terms in the electroweak Lagrangian that govern the dynamics of ϕ and its interactions with other SM fields.

$$\begin{aligned} \mathcal{L}_\phi = & \frac{1}{2}|D_\mu\phi|^2 + \mu_\phi^2|\phi|^2 - \lambda_\phi|\phi|^4 \\ & - \lambda_\ell^i((\bar{\ell}_L^i\phi)\ell_R^i + \text{h.c.}) \\ & - \lambda_d^i((\bar{q}_L^i\phi)d_R^i + \text{h.c.}) - \lambda_u^i((\bar{q}_L^i i\sigma_2\phi^*)u_R^i + \text{h.c.}) \end{aligned} \quad (1.11)$$

$|D_\mu\phi|^2/2$ is the proper, gauge-invariant form of the kinetic term for a complex spin-0 field such as ϕ . Due to the transformation laws for ϕ , in this context the covariant derivative D_μ is equal to $\partial_\mu - ig_B B_\mu/2 - ig_W W_\mu^{at} a$. The second and third terms of \mathcal{L}_ϕ describe the quadratic and quartic self-interactions of the Higgs field, while the remaining terms describe gauge-invariant interactions between ϕ and the SM leptons ℓ^i and quarks q^i (u_R and d_R are right-handed up- and down-type quarks, respectively, and σ_2 is the second Pauli matrix).

Effective Gauge Boson Masses

If μ_ϕ^2 and λ_ϕ are both real and positive, then $|\phi|$ acquires a nonzero *vacuum expectation value* (VEV) $v = \sqrt{\mu_\phi^2/\lambda_\phi}$ in the ground state of the EW theory. Evaluating the kinetic term for the Higgs field at $|\phi| = v$, we can identify effective mass terms for the EW gauge bosons. We are free to choose an $SU(2)_L \times U(1)_Y$ gauge such that the VEV is realized in the real part of the $T_3 = -1/2$ component of ϕ . With this convention, we redefine the EW gauge fields in eigenstates of mass and T_3 , obtaining new fields and associated masses as given in equations 1.12 through 1.14. The W^\pm and Z bosons are the familiar mediators of the weak interactions, and A is the photon.

$$W_\mu^\pm = \frac{1}{\sqrt{2}}(W_\mu^1 \mp iW_\mu^2) \implies m_W = \frac{v}{2}g_W \quad (1.12)$$

$$Z_\mu = \frac{1}{\sqrt{g_W^2 + g_B^2}}(g_W W_\mu^3 - g_B B_\mu) \implies m_Z = \frac{v}{2}\sqrt{g_W^2 + g_B^2} \quad (1.13)$$

$$A_\mu = \frac{1}{\sqrt{g_W^2 + g_B^2}}(g_B W_\mu^3 + g_W B_\mu) \implies m_A = 0 \quad (1.14)$$

It should be noted that massive spin-1 particles have three polarization states, whereas massless spin-1 particles have only two. Thus, in the process of the W and Z acquiring effective masses, they must each also gain an extra degree of freedom. There are three degrees of freedom in the Higgs field under which $|\phi|$ is invariant, and these provide the longitudinal polarizations of the W^+ , W^- , and Z bosons. The remaining degree of freedom in the Higgs field is a massive spin-0 boson that is the quantum of variations in $|\phi|$, known as the Higgs boson. Before describing the Higgs boson in detail, I proceed to discuss the other consequences of the nonzero vacuum expectation value of $|\phi|$.

To reveal the interactions between fermions and the gauge fields W^\pm , Z , and A , we rewrite the covariant derivative in terms of these new fields, as shown in equation 1.15. Here, we define new weak isospin operators as $T^\pm = (t_W^1 \mp it_W^2)/\sqrt{2}$, which annihilate right-handed fermions and raise or lower the value of T_3 for left-handed fermions. We identify the fundamental unit of electric charge e as the coupling constant associated with the photon: $e = g_W g_B / \sqrt{g_W^2 + g_B^2}$. The charge of a particle, in units of e , is given by $Q = T_3 + Y/2$. The *weak mixing angle* θ_W describes the extent to which W^3 and B are mixed in the mass eigenstates Z and A . As θ_W is restricted to the interval $[0, \pi/2]$, it may be defined by the relation $\sin \theta_W = e/g_W$.

$$D_\mu = \partial_\mu - ig_W(W_\mu^+ T^+ + W_\mu^- T^-) - i\frac{g_W}{\cos \theta_W} Z_\mu (T_3 - Q \sin^2 \theta_W) - ieQA_\mu \quad (1.15)$$

Just as in QCD, the kinetic term for the $SU(2)_L$ gauge fields results in interactions between gauge bosons. Reformulating these interactions in terms of the fields W^\pm , Z , and A , we find cubic and quartic interactions involving W^+ , W^- pairs and individual Z bosons

or photons. Notably absent are any cubic or quartic interactions involving only Z bosons and photons, as photons couple to electrical charge Q , and as Z bosons couple to a linear combination of Q and T_3 .

Effective Fermion Masses

In addition to providing effective mass to the W and Z bosons, the Higgs VEV allows fermions to acquire mass through the gauge-invariant terms coupling left- and right-handed fermions to the Higgs field. The resultant fermion masses m_f^i are proportional to the VEV v and to the strength of the coupling λ_f^i , as shown in equation 1.16.

$$m_f^i = \frac{v}{\sqrt{2}} \lambda_f^i \quad (1.16)$$

Interestingly, the mass eigenstates of quarks are not the same as the eigenstates of the weak interaction. For example, in the interaction between a W boson and a strange quark, the other quark is in a superposition of up, charm, and top. In other words, the weak interaction does not necessarily preserve the generation of a quark. We use a unitary 3×3 matrix to account for this, called the CKM matrix after Cabibbo, Kobayashi, and Maskawa. The expected unitarity of the CKM matrix arises from the fact that all generations couple with equal strength to the $SU(2)_L$ gauge fields in the EW theory. The existence of three generations of quarks allows the CKM matrix to contain an overall complex phase, which is the source of CP violation in the standard model.

The Higgs Boson

The curvature of the Lagrangian in the vicinity of the VEV implies a Higgs boson of mass $m_H = \sqrt{2\mu_\phi^2}$, which is the quantum of variations in the magnitude of the Higgs field. Reevaluating the Lagrangian with the substitution $|\phi| = v + H$, we obtain the relevant terms $\mathcal{L}_{\text{int}}^H$ governing interactions between the Higgs boson (H) and other SM particles.

$$\begin{aligned}
\mathcal{L}_{\text{int}}^H &= \left(m_W^2 W^{\mu+} W_{\mu}^- + \frac{1}{2} m_Z^2 Z^{\mu} Z_{\mu} \right) \left(2 + \frac{H}{v} \right) \frac{H}{v} \\
&\quad - \frac{m_H^2}{2v} H^3 - \frac{m_H^2}{8v^2} H^4 \\
&\quad - \frac{m_f^i}{v} \bar{\psi}^i \psi^i H
\end{aligned} \tag{1.17}$$

The first line of equation 1.17 describes interactions involving a W or Z boson pair and one or two Higgs bosons. The strength of these interactions is proportional to the squared mass of the relevant weak gauge boson. Cubic and quartic self-interactions of the Higgs boson are described in the second line, with interaction strength proportional to m_H^2 . Finally, the last line describes the interaction between H and the fundamental fermions ψ^i , in which the interaction strength is proportional to the fermion mass m_f^i . As it represents variations in the field which provides effective mass to the particles of the standard model, the Higgs boson does not interact directly with massless particles such as photons or gluons. Moreover, higher mass implies stronger interaction with the Higgs boson.

1.3 Higgs in Experiment

There are many free parameters in the electroweak theory, all of which must be determined by experiment. In fact, various experiments in particle physics have succeeded in measuring all but one of these parameters. For instance, the Higgs vacuum expectation value is related to the Fermi constant: $v^2 = \sqrt{2}G_F$, where $G_F = \sqrt{2}g_W^2/(8m_W^2)$. The muon lifetime is inversely proportional to G_F^2 , so from precision muon decay experiments we know that $v \approx 246$ GeV [1]. The weak mixing angle θ_W is also well-constrained, as measurements of the Z boson mass and other properties of the weak neutral current are consistent with $\sin^2 \theta_W \approx 0.23$ [2]. The only free parameter in the EW theory that remains only loosely constrained is the mass of the Higgs boson.

Direct constraints on the Higgs mass have come from experiments at the Large Electron-Positron (LEP) collider and at the Tevatron. A lower bound of 114.4 GeV was determined by the LEP experiments, with 95% confidence [3]. At the Tevatron, DØ and CDF have combined the results of their Higgs searches to exclude the region $158 < m_H < 175$ GeV, also with 95% confidence [4]. In addition, indirect constraints may be placed on the Higgs mass by measuring the masses of the W boson (m_W) and the top quark (m_t). Radiative corrections to m_W yield logarithmic dependence on m_H and quadratic dependence on m_t . In figure 1.1, current measurements of m_W and m_t are shown, and hypothetical values of m_H appear as diagonal lines. The central values of m_W and m_t suggest a Higgs boson in the mass region excluded by the LEP experiments, but the uncertainties in these measurements are large enough for direct and indirect constraints on m_H to remain compatible. Combining direct Higgs searches with these indirect constraints, we are 95% confident that the SM Higgs boson has mass in the interval $[114, 158] \cup [175, 185]$ GeV [5].

The standard model allows us to unambiguously predict the properties of the Higgs boson as a function of its mass. Of particular interest to experiments at the Tevatron are the rates of various processes through which a Higgs boson may be produced in a $\bar{p}p$ collision. Figure 1.2 shows the cross sections used to determine these rates, assuming a $\bar{p}p$ center-of-mass energy of $\sqrt{s} = 1.96$ TeV. The analysis presented in this dissertation is a search for associated ZH production, which is the third most likely mode of Higgs production at the

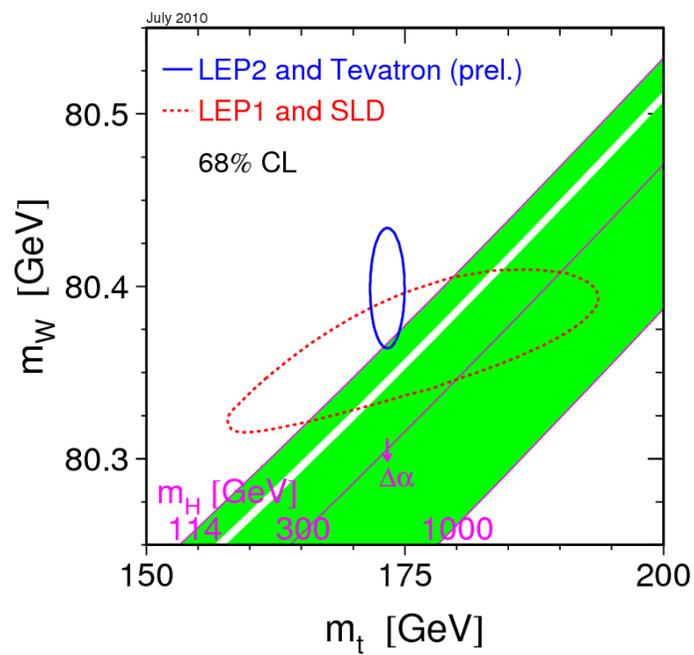


Figure 1.1: W boson mass versus top quark mass. Hypothetical values of the Higgs boson mass in the range $114 < m_H < 1000$ GeV appear as diagonal lines, while the green region represents values of m_H not directly excluded by experiment [5].

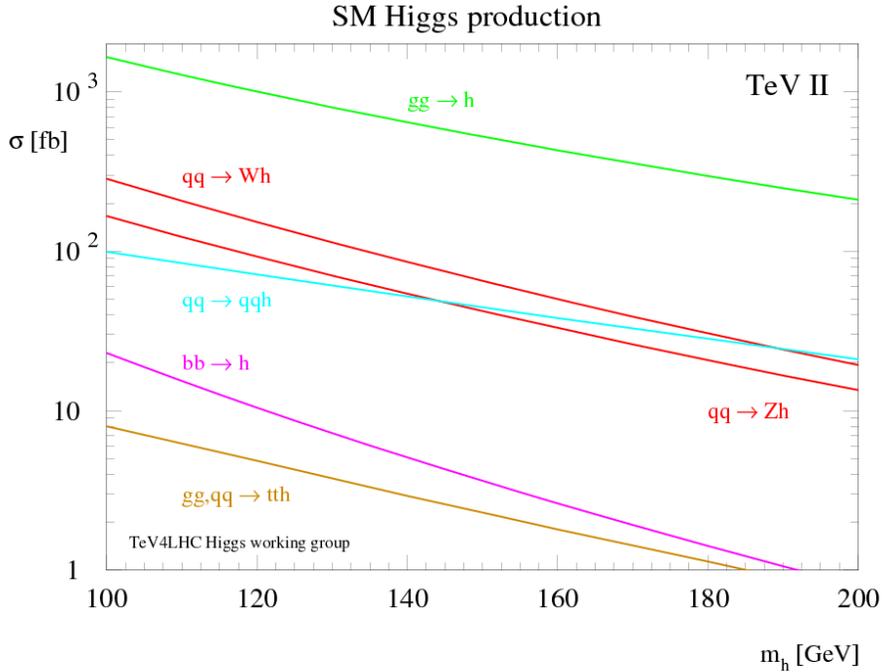


Figure 1.2: Expected cross sections of various Higgs production processes at the Tevatron during RunII, as functions of the Higgs boson mass [6].

Tevatron for $m_H < 145$ GeV.

Also necessary to obtain expected signal yields in Higgs searches, the branching ratios of the Higgs boson are shown in figure 1.3 as functions of the Higgs mass. As a general rule, the Higgs decays to a pair of the heaviest SM particles kinematically available. In other words, the dominant decay should be to the heaviest particle and antiparticle whose mass does not exceed $m_H/2$. There are some exceptions to this rule, however. The decay to W^+W^- is dominant for all Higgs masses greater than 135 GeV. The coupling of the Higgs boson to weak bosons is quadratic in the weak boson mass, while the coupling to fermions is only linear in the mass of the fermion, so decay to WW becomes dominant well before $m_H = 2m_W$, and decay to $t\bar{t}$ is never dominant despite the fact that the top quark is the

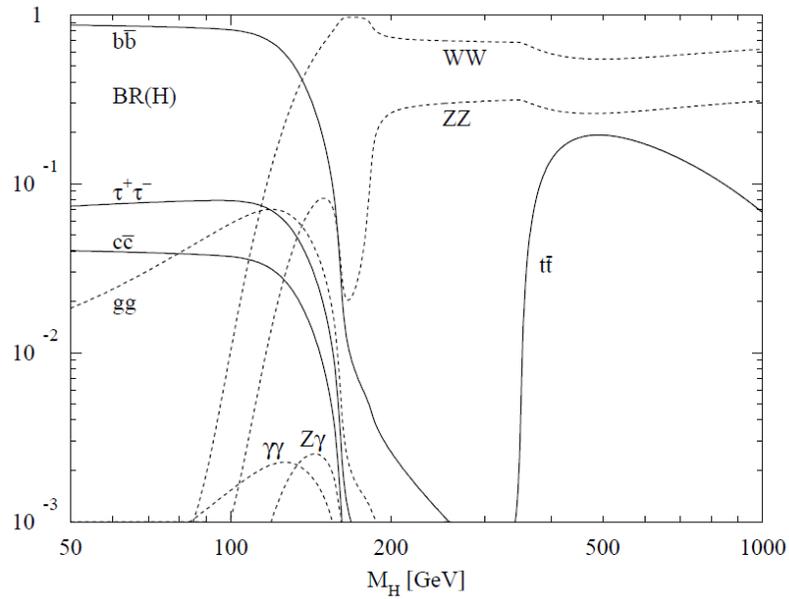


Figure 1.3: Expected branching ratios for various decays of the Higgs boson, as functions of the Higgs boson mass [7].

heaviest SM particle. Also, decay to ZZ never surpasses the W^+W^- decay mode: since the W^+ and W^- are distinguishable, decay to W^+W^- is twice as likely as decay to ZZ in the limit $m_H \rightarrow \infty$.

Chapter 2

EXPERIMENTAL APPARATUS

Previous experiments have excluded a standard-model Higgs boson with mass less than 114.4 GeV, so production of Higgs bosons under controlled experimental conditions might be possible only with a machine that creates very high energy densities. In addition, very precise measurement of this high-energy environment is necessary to confirm the existence of a Higgs boson, and highly specialized equipment is required to perform such measurements.

The Tevatron collider at the Fermi National Accelerator Laboratory (Fermilab) in Batavia, IL is one of two machines in the world theoretically capable of producing large numbers of Higgs bosons, the other being the Large Hadron Collider (LHC) at CERN in Geneva, Switzerland. The Tevatron is the final accelerator in an acceleration chain at Fermilab, in which high-quality proton (p) and antiproton (\bar{p}) beams are produced and accelerated to 980 GeV. The beams travel in opposite directions inside the Tevatron, and they are made to collide at the centers of two immense particle detectors, CDF and DØ. Both detectors are designed to precisely measure various properties of any particles that result from the $\bar{p}p$ collisions. In my search for Higgs bosons produced in association with Z bosons, I have used exclusively data produced by the Tevatron and collected by the DØ detector.

FERMILAB'S ACCELERATOR CHAIN

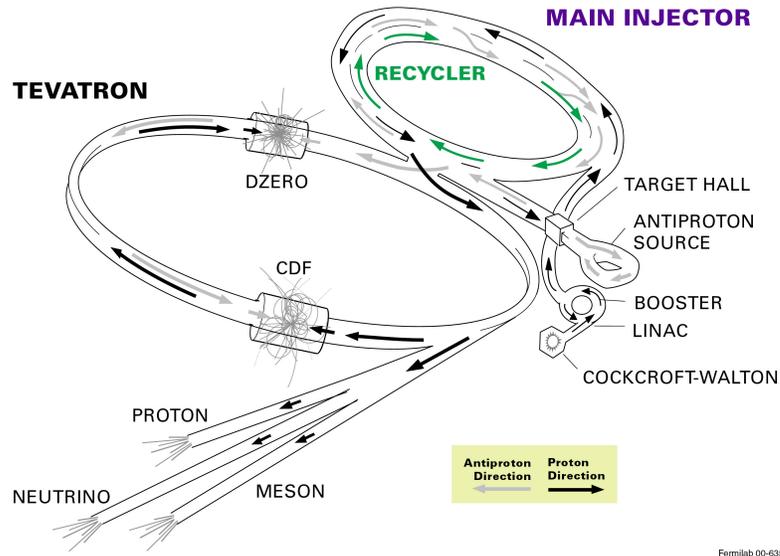


Figure 2.1: The accelerator complex at Fermilab [8].

2.1 Acceleration

Abundant, high-quality proton-antiproton ($\bar{p}p$) collisions with center-of-mass energy $\sqrt{s} = 1.96$ TeV are not easily achieved. Making high-quality, high-energy p and \bar{p} beams requires several specialized accelerators, all working in concert and with nanosecond timing. Figure 2.1 shows the paths of protons and antiprotons as they proceed through Fermilab's accelerator chain, eventually to collide inside the CDF and $D\bar{O}$ detectors.

2.1.1 Proton Beam

The protons that will eventually collide in the Tevatron begin as H_2 gas that's fed into a small magnetron ion source, roughly the size of a matchbox. Figure 2.2 is a simplified schematic of this device. A short but intense electric pulse causes a 40-A arc across the 1 mm gap between anode and cathode. The electrons in the arc spiral tightly around 1 kG magnetic field lines, ionizing the H_2 gas to form a dense plasma. Protons in the plasma

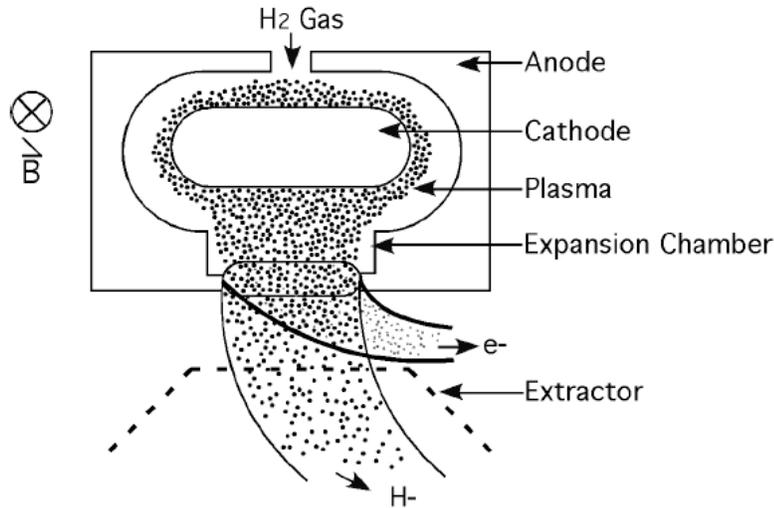


Figure 2.2: The magnetron H^- ion source [9].

are attracted to the cathode, where they gain two electrons and become H^- ions, which are then strongly repelled from the cathode and emerge through a hole in the surrounding anode. In this way, short pulses of H^- ions are created at a rate of 15 Hz.

Cockcroft-Walton

The first stage of acceleration at Fermilab is accomplished by a Cockcroft-Walton style electrostatic preaccelerator. The magnetron ion source is located inside a metallic dome with a constant electric potential of -750 kV relative to the grounded walls of the room housing the preaccelerator. The H^- ions gain 750 keV of kinetic energy as they flow through a resistive pipe connecting the dome with the wall.

Linear Accelerator

From the preaccelerator, the H^- ions travel into Fermilab's linear accelerator (Linac) at $0.04c$, or 12,000 km/s. The Linac is composed of two distinct accelerators: the older Drift Tube Linac (DTL) is 79 m long and accelerates the ions from 750 keV to 116 MeV ($0.4c$), while the Side-Coupled Linac (SCL) is 67 m long and increases their kinetic energy from

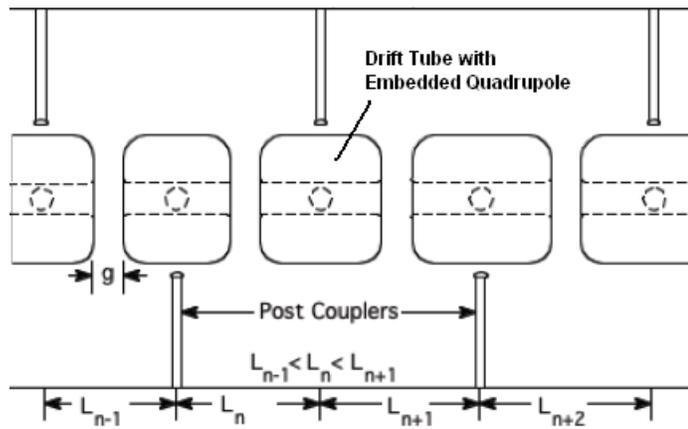


Figure 2.3: Top-down view of several accelerating cells in Fermilab's Drift Tube Linac [9].

116 MeV to 400 MeV ($0.7c$).

Five resonant RF tanks compose the DTL, each tank containing a series of drift tubes suspended from the top of the tank. The RF tanks are tuned such that the H^- ions are accelerated as they traverse the gaps between drift tubes, and the drift tubes shield the ions from the decelerating half of the RF cycle. As illustrated in Figure 2.3, each successive drift tube is longer than the last, so the time interval between gap traversals is constant despite the increasing speed of the ions. Also, the drift tubes contain quadrupole magnets that serve to focus the beam.

The second stage of the Linac has a different design: instead of large RF cavities containing several drift tubes, the SCL is a series of small resonant cavities, one cavity per accelerating cell. This design is shown in Figure 2.4. Small coupling cavities physically connect adjacent accelerating cavities, with the effect of producing a $\pi/2$ phase shift between adjacent cells. This phase shift allows the SCL to achieve twice as much acceleration over the same distance, since an appropriately-timed bunch of H^- ions entering the SCL will never encounter a decelerating electric field. This also eliminates the need for drift tubes, although it remains necessary that each accelerating cell be longer than the last. In the SCL, there are seven cavity modules, each containing many accelerating cavities. Between the modules, quadrupole magnets maintain a focused beam.

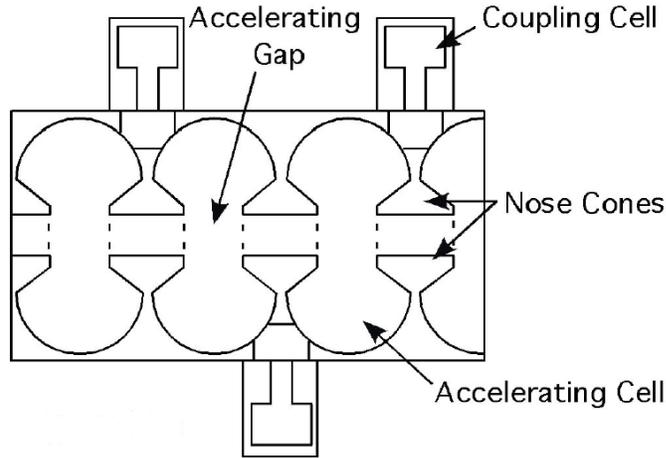


Figure 2.4: Side view of several accelerating cells in Fermilab’s Side-Coupled Linac [9].

Booster

Exiting the Linac, the H^- ions pass through carbon foil, which strips the electrons from the hydrogen nuclei, creating the proton beam [10].

Next in the accelerator chain is the Booster, a circular *synchrotron* with a radius of 75 m. The Booster accelerates the protons from 400 MeV to 8 GeV ($0.99c$). It operates at the same 15 Hz frequency as the H^- ion source, so the acceleration occurs over a time interval of only 33 ms. Synchrotrons rely on dipole magnetic fields to bend the beam in a circular trajectory, such that the same RF cavities are traversed multiple times. In order to keep the orbit radius constant, the dipole magnetic fields must increase as the beam gains energy: in the Booster, the fields vary from 740 G at injection to 7 kG at extraction. Also, to account for the increasing speed of the protons, the frequency of the accelerating cavities must increase from 38 to 53 MHz. These frequencies translate to 84 RF cycles per proton orbital period, which means that protons injected uniformly into the Booster will naturally form 84 bunches evenly spaced around the circumference of the synchrotron.

It should be noted that once protons exit the booster at 8 GeV, special relativity limits further increases in speed to at most 1%, while kinetic energy is unbounded. Because of this,

in subsequent synchrotrons the accelerating cavity frequencies do not substantially vary as protons accelerate—“acceleration” refers to increased kinetic energy rather than increased velocity.

Main Injector

The Main Injector is another synchrotron, seven times the size of the Booster. It accepts protons and antiprotons from several sources, accelerates them from 8 GeV to 120 or 150 GeV, manipulates their beam structure, and sends beam to several destinations. For the purpose of producing $\bar{p}p$ collisions in the Tevatron, the Main Injector has three important operational modes: injection of protons into the Tevatron, direction of protons to the antiproton source, and injection of antiprotons into the Tevatron.

For injection of proton beam into the Tevatron, only 7 of the 84 bunches in the Booster are extracted to the Main Injector, where they accelerate to 150 GeV. Then, through RF manipulation, the seven bunches coalesce into one high-intensity bunch, which is injected into the Tevatron. This process takes roughly three seconds from extraction to injection, and it is repeated 36 times to fill the Tevatron with 36 bunches of approximately 3×10^{11} protons each.

When using protons to produce antiprotons, the Main Injector performs two complete extractions of all 84 bunches in the Booster. This results in a beam of 84 bunches twice as intense as in the Booster, approximately 8×10^{12} protons in total. This double batch of protons is accelerated to 120 GeV, and some RF manipulation is done to minimize the momentum spread of the bunches. Then, the Main Injector delivers the proton beam to the antiproton source, which is described in the next section.

Once a sufficient number of antiprotons have been produced, the Main Injector extracts four groups of antiprotons, each group containing multiple bunches, from either the Accumulator or the Recycler. After accelerating to 150 GeV, the four groups coalesce into four high-intensity bunches, and these are injected all at once into the Tevatron. This process repeats nine times to fill the Tevatron with 36 bunches of approximately 7×10^{10} antiprotons each.

2.1.2 Antiproton Beam

To produce antiprotons at Fermilab, an intense, 120-GeV proton beam from the Main Injector collides with a target made of Inconel, an alloy of nickel, chromium and iron. The resultant nuclear reactions produce a spray of particles, including some antiprotons at approximately 8 GeV.

Particles in the spray with negative charge and kinetic energy close to 8 GeV are focused into a parallel beam by a lithium cylinder that acts as a lens. Oriented along the beam direction, the cylinder is 2 cm in diameter and 15 cm long. A current of several hundred kA flows through the cylinder, creating an azimuthal magnetic field that increases linearly with internal radius and affects the particle trajectories much as a standard convergent lens affects a beam of light. Lithium was chosen in order to be as “transparent” as possible; however, the lens absorbs roughly 18% of antiprotons through nuclear interactions.

After the lithium lens, a dipole magnet selectively bends the trajectories of negatively charged, 8-GeV particles towards the Debuncher, allowing all other particles to be absorbed by the surrounding material.

Debuncher

The antiprotons traveling into the Debuncher form a very low-quality beam, with a momentum spread of several percent and transverse size roughly equal to the width of the lithium lens. The Debuncher is a synchrotron that functions to improve the beam quality in preparation for injection into the Accumulator. It “cools” the beam, reducing the longitudinal and transverse spreads in momentum, necessarily destroying the bunch structure of the beam in the process.

While longitudinal cooling is performed through RF manipulation similar to the technique used in the Main Injector, the Debuncher has a stochastic cooling system to reduce the spread in transverse momentum. To achieve stochastic cooling, particle orbits are actively monitored at some point along the circumference of the synchrotron. An orbit with significant transverse deviation produces a correcting signal that’s amplified and sent to a dipole magnet (a *kicker*) on the other side of the synchrotron, which adjusts the orbit to

be more ideal. This process relies on very precise timing and placement to ensure that the kicker receives the signal exactly when the deviant particle arrives: incorrect timing could result in transverse heating of the beam.

The Main Injector delivers a batch of protons to the \bar{p} source approximately every two seconds, so the Debuncher must complete its beam cooling cycle in the same time interval. Just before the next batch arrives, the antiprotons are extracted from the Debuncher and proceed to the Accumulator.

Accumulator

Unlike in the Debuncher, antiprotons circulate in the Accumulator for many hours. As its name suggests, the Accumulator is a synchrotron that stores antiprotons, accumulating batches from the Debuncher until the stored \bar{p} beam reaches the desired intensity. Every two seconds, antiprotons from the Debuncher are injected into a slightly larger orbit than the already accumulated beam. Through careful RF manipulation and further stochastic cooling, the new antiprotons merge with the main beam over the course of 20 minutes.

The process of producing and storing antiprotons is not efficient: only two antiprotons are stored for every 10^5 protons striking the Inconel target. Indeed, this is the most important bottleneck in the physics reach of the Tevatron experiments. Recent large improvements to the $\bar{p}p$ collision rate in the Tevatron have come as a result of faster and more efficient \bar{p} production.

From the Accumulator, the core of the \bar{p} beam may be extracted to the Main Injector, where it is either accelerated for immediate use in the Tevatron or injected into another \bar{p} storage ring, the Recycler.

Recycler

Large numbers of antiprotons in the Accumulator tend to become unstable over time, which can lead to the sudden loss of the entire \bar{p} beam. To accumulate antiprotons beyond this limit, the stored beam is often transferred to the Recycler, another synchrotron with improved long-term stability.

The Recycler employs stochastic and electron cooling to further reduce the longitudinal and transverse spreads of the \bar{p} beam. This allows for more antiprotons to accumulate in the Recycler without causing instability, and it results in a more intense \bar{p} beam for injection into the Tevatron. Electron cooling is achieved by injecting a very high-quality e^- beam that overlaps with the \bar{p} beam and has the same average momentum. The electrons and antiprotons interact, transferring momentum until they reach thermal equilibrium. During this process, thermal energy flows from the hot \bar{p} beam to the cold e^- beam. Dipole magnets remove the electrons while having a minimal effect on the antiprotons, leaving the \bar{p} beam with significantly reduced momentum spread.

After many hours of accumulating antiprotons, the \bar{p} beam from the Recycler may be extracted to the Main Injector, where it is accelerated to 150 GeV and injected into the Tevatron.

2.1.3 The Tevatron

The final and largest accelerator at Fermilab is the Tevatron, a synchrotron with a radius of 1 km. It is also the only accelerator at Fermilab to utilize superconducting magnets, necessary to maintain a focused beam and circular orbit at energies as high as 980 GeV. The superconductor chosen for the Tevatron is niobium-titanium, which has a critical temperature of 10 K and a maximum magnetic field of 15 T. In practice, the Tevatron uses liquid helium to cool the NbTi to 4.6 K, and up to 4 kA of current produces dipole magnetic fields as high as 4.4 T. The superconducting cables are made of 23 strands 0.7 mm in diameter, each strand containing 2050 NbTi filaments embedded in a copper matrix. In total, there are roughly 1000 superconducting dipole, quadrupole, and higher-order magnets in the Tevatron, and up to 350 MJ of energy is stored in their magnetic fields.

The p and \bar{p} beams in the Tevatron each consist of 36 bunches in three trains of 12 bunches. There is a 396 ns time interval between bunches in a train, and 2.617 ms between the trains. As stated earlier, the protons circulate clockwise, the antiprotons counterclockwise. The protons and antiprotons share the same beam pipe, but electrostatic *separators* ensure that the p and \bar{p} orbits do not overlap (except where collisions are desired): the orbits

form a double helix inside the beam pipe, nominally separated by only 5 mm. Beginning at 150 GeV ($0.99999c$), the Tevatron simultaneously accelerates protons and antiprotons to 980 GeV ($0.9999995c$), just 150 m/s below the speed of light.

After acceleration, special quadrupole magnets bring the beams to sharp focus at the two collision points. This increases the transverse density of the bunches while controlling the longitudinal position of $\bar{p}p$ interactions, optimizing the interaction rate and assuring symmetry with respect to the particle detectors. Finally, the separators allow the beams to overlap at the collision points, *collimators* scrape away potentially damaging beam halo, and the detectors can begin to record the products of 1.96-TeV $\bar{p}p$ collisions.

2.2 Collision

The vast majority of $\bar{p}p$ collisions result in the proton and antiproton scattering at small angles with respect to the beam line; in our search for the Higgs boson, these events are ignored. Still others are glancing shots, in which the proton or antiproton disintegrate, but very little momentum is exchanged transverse to the beam direction. A small fraction of collisions include *hard scattering*, in which a constituent of the proton annihilates with a constituent of the antiproton, producing any number of other particles with significant transverse momenta. If the standard model is correct, occasionally this type of collision should produce a Higgs boson.

The constituents of protons and antiprotons are quarks and gluons, collectively called *partons*. We refer to the partons directly involved in the hard scatter as the *initial state*, and the particles emerging directly from the hard scatter are collectively known as the *final state*. Because hard scattering generally causes the incident proton and antiproton to disintegrate, any partons not directly involved in the hard scatter evolve into showers of hadrons close to the beam line. This is called the *underlying event*.

Many of the particles in the final state are unstable and decay rapidly, often before leaving the beam pipe. The only particles in the standard model with lifetimes long enough to be directly measured by our detectors are photons, electrons, muons, and several light hadrons such as p , n , π , and K . (Neutrinos also have long lifetimes, but our equipment is much too small to detect them directly.) Thus, the presence of short-lived particles in the final state must be inferred by measuring the kinematics of their long-lived decay products.

2.2.1 Cross Section and Luminosity

It's impossible to predict what will happen in any given $\bar{p}p$ collision. When a hard scatter occurs, which partons constitute the initial state is completely random, although some partons are more likely than others. In addition, a specific initial state could result in a wide variety of final states, so the particles directly created through hard scattering in $\bar{p}p$ collisions are doubly random. (This is in contrast to e^+e^- collisions, where the initial state is certain.) We can't say for sure what will result from a single collision, but we can

make predictions concerning large numbers of collisions. In particular, the standard model allows us to calculate the average production rate of any specified final state. For collider experiments, it is useful to separate any such rate R into a *cross section* σ , inherent to the incident and final-state particles, and a *luminosity* L , dependent on macroscopic properties of the colliding beams.

$$R(\bar{p}p \rightarrow X, p \text{ beam}, \bar{p} \text{ beam}) = \sigma(\bar{p}p \rightarrow X)L(p \text{ beam}, \bar{p} \text{ beam}) \quad (2.1)$$

The cross section is so named because it has units of area, and macroscopic collision rates are proportional to the cross-sectional area of the colliding objects. By convention, we report cross sections in units of *barns*, abbreviated “b.” A barn is defined as 10^{-28} m^2 , which is roughly the cross-sectional area of a uranium nucleus, famously described as “big as a barn” by American physicists conducting nuclear research during World War II. Indeed, a barn is quite large in comparison to the cross sections of general interest in particle physics; for example, the cross section for top quark pair production in Tevatron collisions is roughly 7 *picobarns* ($1 \text{ pb} = 10^{-12} \text{ b}$), and the expected cross section for Higgs production at the Tevatron is at most several hundred *femtobarns* ($1 \text{ fb} = 10^{-15} \text{ b}$).

As the other piece necessary to compute a production rate, luminosity accounts for the effect of beam structure on the frequency of collisions. The number of $\bar{p}p$ collisions in a given period of time depends on numerous aspects of the p and \bar{p} beams, such as the number of bunches to have arrived at the collision point, the number of particles per bunch, and the transverse and longitudinal extent of the bunches. Equation 2.2 shows how the luminosity depends on these quantities.

$$L = \frac{f n_B n_p n_{\bar{p}}}{2\pi(\sigma_p^2 + \sigma_{\bar{p}}^2)} F \quad (2.2)$$

Here, f is the orbital frequency of the beams, n_B is the number of bunches per beam (36), n_p and $n_{\bar{p}}$ are the number of protons and antiprotons per bunch, and the parameters σ_p and $\sigma_{\bar{p}}$ are the transverse Gaussian widths of the beams at the collision point. F is a factor less than one that accounts for dependence of the luminosity on the longitudinal extent of the bunches and the length of the interaction region. We report luminosity in

units of $\text{cm}^{-2}\text{s}^{-1}$, while the *integrated luminosity* over a period of time is given in units of inverse barns, such as pb^{-1} or fb^{-1} . Using this convention, if we know the integrated luminosity of our data, then we can immediately translate a cross section to an expected number of events for a particular final state. For example, in 1 fb^{-1} of data I would expect approximately $1 \text{ fb}^{-1} \times 7 \text{ pb} = 7000$ collisions in which hard scattering produced a pair of top quarks.

As the p and \bar{p} beams collide in the Tevatron, the luminosity decreases over time because of degradation in beam quality and because particles are gradually lost from the beam. Of course, the beams lose a proton and antiproton with every collision, but the actual decrease in luminosity is much faster than one would expect from this effect alone. The largest contribution to beam loss comes from transverse expansion of the beam, which leads to decreased $\bar{p}p$ interaction rate and causes particles to collide with the collimators. This expansion has many sources: scattering of the beam with residual gas in the imperfect vacuum of the beam pipe, interaction between particles in a beam, interaction between beams, and electrical noise in the magnets and RF cavities, to name a few. The total effect of beam loss is that the luminosity initially present in the Tevatron degrades by approximately 80% over the course of a day. Balancing beam loss with the \bar{p} accumulation rate, Tevatron operators maximize integrated luminosity by choosing the optimal time to dump the degraded beams and inject new p and \bar{p} batches.

2.2.2 Coordinate System

A well-chosen coordinate system is of utmost importance in approaching any problem in physics, and collider experiments are no exception to this rule. Naturally, we choose the interaction point as the origin of our coordinates. The polar axis is chosen to coincide with the direction of incident protons, since we expect rotational symmetry about the beam line. Cartesian x and y coordinates extend away from the center of the Tevatron and away from the center of the Earth, respectively, and the azimuthal angle ϕ is defined to be zero along the x -axis and $\pi/2$ along the y -axis. ϕ is constrained to the interval $[0, 2\pi)$, while the polar angle θ must be in $[0, \pi]$.

While the incident proton and antiproton each have 980 GeV of kinetic energy, each parton in the initial state carries only a fraction of the momentum of its parent proton or antiproton. Because these fractions are generally different between the two initial-state partons, the rest frame of the hard scatter is boosted along the z -axis with respect to the particle detectors. Therefore, in the absence of angular correlations between initial and final states, final-state particles would be distributed evenly in longitudinal rapidity y , as defined in Equation 2.3. In addition, the “angle” Δy between two particles is invariant under Lorentz boosts in the z direction.

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right) \quad (2.3)$$

The rapidity is inconvenient as a coordinate because it depends on the mass of the particle. However, nearly all directly-detected particles have masses that are negligible in comparison to their momenta. In this limit, y reduces to a quantity which depends only on θ , as shown in Equation 2.4.

$$\eta = \lim_{\frac{m}{p} \rightarrow 0} y = -\ln \left(\tan \left(\frac{\theta}{2} \right) \right) \quad (2.4)$$

We thus define the *pseudorapidity* η , and use it in place of the polar angle θ . η is zero if $\theta = \pi/2$, and η approaches $+\infty$ ($-\infty$) in the direction of the p (\bar{p}) beam. Using these coordinates, the angle ΔR between two particles is defined in Equation 2.5.

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} \quad (2.5)$$

Every initial-state parton has very little momentum transverse to the beam. As a consequence, conservation of momentum forces the vector sum of the transverse momenta of final state particles to be approximately zero. This is a very powerful constraint, allowing us to infer mis-measurement and the existence of undetectable particles such as neutrinos in the final state. For this reason, our detectors are specifically designed to measure the magnitude of transverse momentum (p_T) for each particle, in addition to direction and energy. The four quantities p_T , η , ϕ and E uniquely determine each particle’s four-momentum.

2.3 Detection

Much as a camera measures the positions and momenta of photons emitted from an interesting scene, collider-based particle detectors aim to measure the properties of every particle emerging from a high-energy collision. The data associated with a particular collision is effectively a sophisticated, three-dimensional snapshot of the interaction region.

The DØ detector first recorded $\bar{p}p$ collisions during RunI of the Tevatron, 1992–1996 [11]. From early 1996 to March 2001, Fermilab’s accelerator-based physics program received extensive upgrades. During this transition, the Main Injector and Recycler were built, increasing Tevatron luminosity by an order of magnitude. Also, the maximum energy increased from 1.8 to 1.96 TeV, and the number of colliding bunches increased from 6 to 36. Because of the increased luminosity and frequency of collisions, it was necessary to upgrade nearly every component of the particle detectors. In addition, a powerful solenoidal magnet was added to DØ in order to measure the momenta of charged particles before their absorption in the calorimeter. RunII of the Tevatron began in 2001 and continues to this day. As I have analyzed only data collected during RunII, in this section I provide a detailed description of the RunII DØ detector [12].

Since we expect collisions to exhibit cylindrical symmetry, the accuracy and precision of our measurements should ideally be invariant under rotations in ϕ . For this reason, much of the DØ detector has a cylindrical design. A side-view of DØ is shown in Figure 2.5. Subdetector components of DØ form nested cylinders all centered on the interaction point. This is most evident in the design of the inner tracking system, described in subsection 2.3.1. Outside the inner trackers are the calorimeters, described in subsection 2.3.2. Outside the calorimeters is another tracking system described in subsection 2.3.3, with the specific purpose of measuring muons. The muon tracker departs from the cylindrical theme, instead exhibiting a rectangular design.

In addition to measuring the properties of particles created in hard scattering, DØ locally monitors the luminosity of the p and \bar{p} beams in order to determine expected production rates. DØ’s luminosity monitor is described in subsection 2.3.4.

Equally important as the subdetectors are the electronics and computers through which

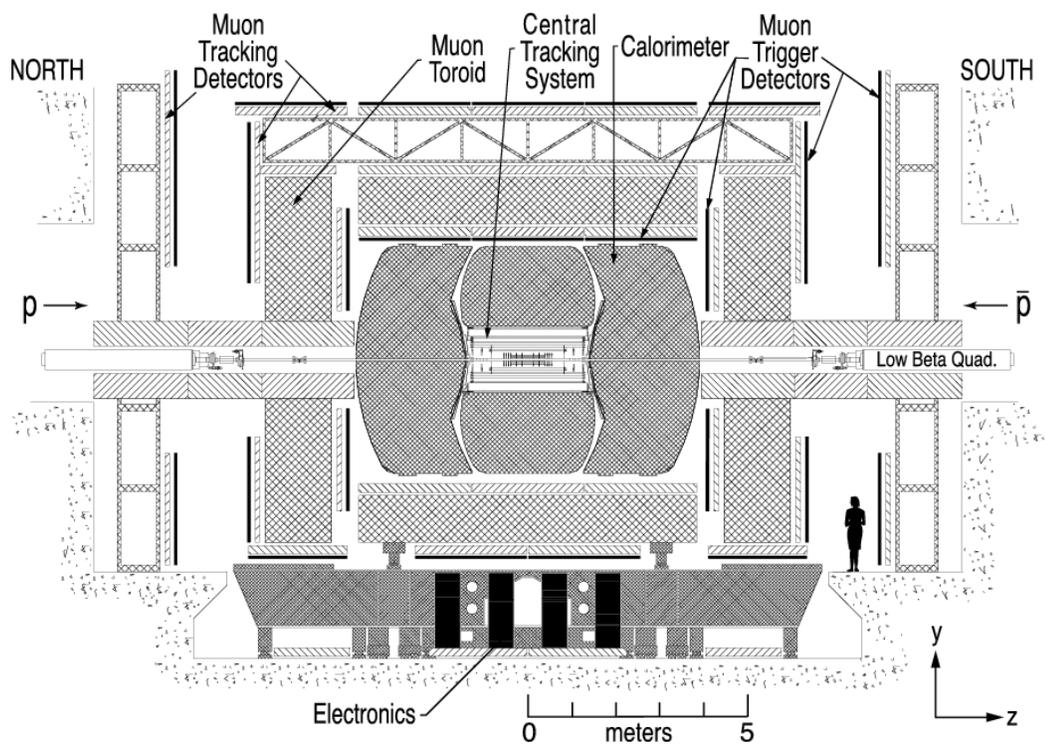


Figure 2.5: Profile of the DØ detector, as seen from inside the Tevatron ring [12].

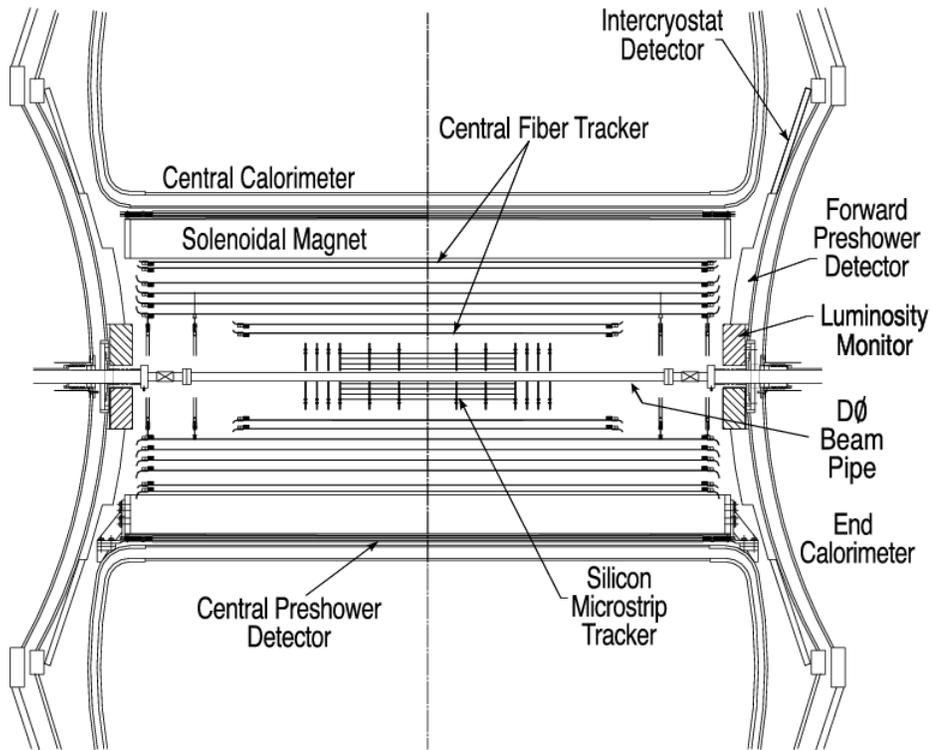


Figure 2.6: DØ’s innermost subdetectors [12].

data flows on its way to being recorded for use in analyses such as a Higgs search. The various stages of filtering and processing that occur along this path are described in subsection 2.3.5.

2.3.1 Inner Tracking

Particles emerging from $\bar{p}p$ collisions with sufficiently high p_T will pass through the beryllium beam pipe and into the first layers of the DØ detector, shown in figure 2.6. The *silicon microstrip tracker* (SMT) and *central fiber tracker* (CFT) measure the positions of charged particles as they travel away from the interaction point. A series of *hits* left by the same particle in multiple tracking layers allows us to determine the particle’s most likely trajectory, or *track*. For more details on the determination of tracks, see section 3.1.

A superconducting solenoid magnet surrounds both trackers, immersing them in a 2-T magnetic field parallel to the beams. This causes the trajectory of any charged particle to

curve in the transverse plane, and the direction of curvature is used to distinguish between positive and negative charge. Using equation 2.6, we can determine p_T from the radius of curvature ρ of the track. (Actually, curvature determines the ratio p_T/q , but $|q| = e$ for almost every charged particle we detect.)

$$p_T = \rho \left| q\vec{B} \right| \quad (2.6)$$

Extrapolating the tracks of charged particles backwards, we can precisely determine the location of the hard scatter, called the *primary vertex* (PV). Using DØ's inner trackers, we can locate the PV to within $35 \mu\text{m}$ along the z direction and within $15 \mu\text{m}$ along the transverse directions. The distribution of primary vertices shows the extent of the *interaction region*, where p and \bar{p} bunches overlap. The interaction region is approximately Gaussian with transverse variance 0.1 mm and longitudinal variance 25 cm , so the PV may deviate from the center of the detector by large fractions of a meter along the z -axis. This deviation affects our interpretation of data from the detector, so precise determination of the PV is of fundamental importance.

Precise measurement of the PV also allows us to distinguish secondary vertices, a telltale sign of relatively long-lived particles such as B mesons. Such particles often do not exist for enough time to penetrate into the detector, but vertex resolution much smaller than the beam pipe radius allows us to indirectly detect their presence. This topic is discussed at length in section 3.5.

Silicon Microstrip Tracker

A solid-state diode under reverse bias can function as an effective particle detector. Applying sufficient voltage to deplete the conduction band, very little current nominally flows through the diode. If an energetic, charged particle passes through the diode in this state, it excites electrons into the conduction band, creating a pulse of current that serves as a signal of the particle's presence.

In silicon, the gap between valence and conduction bands is roughly 1 eV , so a particle of several GeV loses only a small fraction of its energy while producing an electrical signal

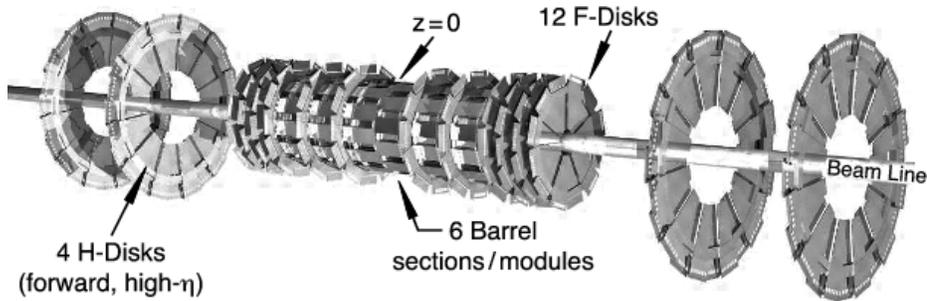


Figure 2.7: $D\bar{O}$'s silicon microstrip tracker as it existed during RunIIa [12]. For RunIIb, another barrel (Layer 0) was added inside the other barrels, and the outermost two H-disks were removed.

in a silicon detector. For high-energy physics detectors, typically such diodes are several hundred μm thick, which strikes a balance between maximizing electrical signal strength and minimizing the drift time of electrons across the diode. In traversing a typical silicon detector layer, a 1-GeV particle deposits 0.01% of its energy in promoting at least 20,000 electrons to the conduction band, and those electrons drift to the cathode in roughly 10 ns.

Adding long, thin strips of p-doped silicon to a large, n-doped wafer creates a series of elongated diodes, called *microstrips*. An array of such wafers, of various sizes and shapes, composes $D\bar{O}$'s silicon microstrip tracker [13]. The microstrips are truly microscopic, with only 50–150 μm separating adjacent strips, which translates directly to very good resolution perpendicular to the strips. In order to determine the position of a hit in the direction parallel to the strips, many wafers are double-sided. Highly n-doped strips are added on the side opposite to the p-doped strips, with a significant angle between strips on opposite sides. Thus, a hit may be localized in all three spatial dimensions by considering signals from both sides of the wafer.

Figure 2.7 illustrates the structure of the SMT. Rectangular silicon wafers are arranged in barrels that wrap around the beam line, each barrel consisting of multiple concentric layers. Wedge-shaped wafers are arranged in disks transverse to the beam line, located between barrels and at high $|z|$. This configuration is intended to maximize the probability that particles from the hard scatter will have normal incidence with the various SMT layers,

thereby minimizing the average hit resolution.

Layers 1–4 of the SMT barrel modules are each 12 cm long, with approximate radii of 3, 5, 7 and 9 cm from the interaction region. Six modules extend barrel coverage to $|z| = 37.8$ cm, and wafer composition varies with layer and module. In all barrel modules, layers 2 and 4 consist of double-sided wafers, with a 2° angle between axial (parallel to the z axis) and stereo microstrips. In the two barrel modules at high $|z|$, layers 1 and 3 are single-sided (axial); in the four central barrels, layers 1 and 3 are double-sided with a 90° stereo angle.

During the transition from RunIIa to RunIIb in 2006, an additional layer of silicon was installed inside the original four SMT barrel layers [14]. Intended to alleviate the effects of SMT aging and to provide improved secondary vertex resolution, layer 0 is entirely single-sided.

There are two types of SMT disks: F-disks and H-disks. Twelve F-disks are located at $|z| = 12.5, 25.3, 38.2, 43.1, 48.1,$ and 53.1 cm. The F-disks have an inner radius of 2.57 cm and outer radius 9.96 cm, and each disk comprises 12 double-sided silicon wedges, with a stereo angle of 30° . Four H-disks contribute to SMT coverage far from the interaction point, at $|z| = 100.4$ and 121.0 cm. The H-disks have an inner radius of 9.5 cm and outer radius 26 cm, and each disk comprises 24 pairs of single-sided silicon wedges. Each pair is stacked to effectively serve as a double-sided wedge, with a stereo angle of 15° . For RunIIb, the outermost two H-disks were removed, and their readout channels are now used for layer 0.

Central Fiber Tracker

While the SMT excels at locating vertices, $D\mathcal{O}$'s central fiber tracker provides a superior measurement of transverse momentum. This is primarily due to the larger radial extent of the CFT: its eight concentric cylinders occupy the space between 20 and 52 cm from the beam line. Longitudinally, the first two CFT cylinders extend only to $|z| = 83$ cm to accommodate the SMT H-disks, while the remaining six cylinders extend to $|z| = 1.26$ m. Each CFT cylinder is made of two doublet layers of scintillating fibers $835 \mu\text{m}$ in diameter. The first doublet in each cylinder is axial, while the second doublet is stereo at $+3^\circ$ (u) or

-3° (v). The first cylinder has a u stereo layer, and subsequent cylinders alternate between u and v .

Within a doublet layer, the two layers of fiber are offset by half of the fiber spacing, which avoids potentially insensitive regions between fibers and provides hit resolution significantly better than the radius of the fibers. Assuming that the location of each fiber is known to better than $50\ \mu\text{m}$, the resolution of a single doublet layer is approximately $100\ \mu\text{m}$.

Each scintillating fiber is made of polystyrene doped with primary and secondary fluorescent dyes. The primary fluor is paraterphenyl (pT) at 1% concentration by weight. An energetic, charged particle passing through the fiber excites the polystyrene molecules, and that energy is rapidly transferred to pT molecules via a non-radiative, dipole-dipole interaction. The excited pT molecules decay in just a few nanoseconds, emitting photons of wavelength $340\ \text{nm}$. Since polystyrene is opaque to photons of this wavelength, a secondary, wavelength-shifting fluor is necessary. 3-hydroxyflavone (3HF) serves this purpose at a concentration of 1500 ppm. 3HF molecules absorb the 340-nm photons and emit at $530\ \text{nm}$, a wavelength to which polystyrene is transparent.

The scintillating fibers are optically connected to clear fibers at one end of the CFT, while the other end of each fiber is coated with 90%-reflective aluminum. The clear fibers are chemically and structurally similar to the scintillating fibers, but do not contain any fluorescent dyes. Scintillation photons travel up to 5 m in the scintillating fibers (a round-trip in the outer six cylinders), and up to 12 m in the clear fibers. To maximize transmission by total internal reflection over this distance, all fibers are dual-clad. The index of refraction varies from 1.59 in the fiber bulk to 1.49 in the first cladding and 1.42 in the second cladding, which results in attenuation lengths of roughly 5 m in the scintillating fibers and 8 m in the clear fibers.

It should be noted that the production of scintillation light in the CFT is entirely passive. In contrast to the SMT, the CFT in principle requires no electrical power to signal the presence of energetic, charged particles. In practice, however, a CFT hit is composed of at most 40 photons emerging from a clear fiber. To reliably detect such a small number of photons, DØ employs extremely sensitive, solid-state devices called *visible light photon counters* (VLPC's) [15]. Cooled by liquid helium to a temperature of 9 K, each VLPC

operates with a quantum efficiency of more than 75% and a signal gain of 22,000–65,000.

2.3.2 Calorimetry

The properties of a track can be used to determine three components of a particle’s four-momentum; to measure the fourth component, energy, we use calorimeters. In addition, trackers are only sensitive to charged particles, while calorimeters provide sensitivity to nearly every particle type. Since various neutral particles, such as photons and neutrons, are often produced in high-energy $\bar{p}p$ collisions, calorimeters are required for a complete record of hard-scatter events.

While the design of a tracker aims to minimize the impact on a particle’s trajectory, particle detection in a calorimeter is destructive. Calorimeters are designed to absorb all of the energy from an incident particle, and to produce an electrical signal from which the absorbed energy may be inferred. The resulting energy measurement is most precise when the incident particle has very high energy, whereas momentum resolution from tracking degrades at high p_T , so tracking and calorimetry are complementary.

The processes involved in absorbing a particle’s energy vary greatly depending on the type of particle. High-energy positrons and electrons lose energy primarily through bremsstrahlung. The radiated photons may convert to electron-positron pairs, which radiate photons in turn, and so on until the resultant particles each have very little energy. We refer to such cascades as electromagnetic (EM) showers. Compared to electrons, sufficiently energetic photons travel somewhat farther into the calorimeter before converting and producing pairs of overlapping EM showers. Hadrons travel much farther until they interact with nuclei to produce wide sprays of various other hadrons, forming hadronic showers. With a finely-segmented calorimeter, these signatures are easily distinguishable, so we often rely on calorimeters for particle identification.

It’s useful to define some terminology for describing interactions between high-energy particles and bulk matter. A *radiation length*, denoted X_0 , is the mean distance over which a relativistic electron radiates all but $1/e$ of its energy through bremsstrahlung. An equivalent definition is $7/9$ of the mean free path of a high-energy photon before conversion to e^+e^- .

The *nuclear interaction length* is an analogous quantity that is more relevant for hadrons. Denoted λ_I , it is defined as the mean free path of a relativistic hadron prior to inelastic interaction with a nucleus.

Preshower Detectors

Located between DØ's inner trackers and liquid-argon calorimeters, the preshower detectors combine tracking and calorimeter technology to improve particle identification and to help account for the fact that many EM showers begin in the solenoid and in other inner-tracker material. There are three preshower detectors in DØ: the *central preshower* (CPS), between the solenoid and the central liquid-argon cryostat, and two *forward preshower* (FPS) detectors, located near the beam pipe on the inner faces of the endcap liquid-argon cryostats. The CPS covers $|\eta| < 1.3$, and the FPS covers $1.5 < |\eta| < 2.5$.

The sensitive components of the preshower detectors are very similar to those of the CFT. Doublet layers of scintillating polystyrene strips send photons to an array of VLPC's, providing electronic evidence of charged particles. The cross section of each strip is a rounded equilateral triangle, 6 mm to a side. In the center of each strip, a wavelength-shifting fiber identical to those in the CFT collects scintillation light, and a clear fiber transports the light to the readout hardware. The two layers in a doublet are offset and inverted such that the triangular strips fit together tightly, eliminating dead space between strips. Just as in the CFT, the overlap between strips in a doublet can yield position measurements more precise than the width of an individual strip.

The CPS contains three cylindrical doublet layers surrounding a lead cylinder 7/32 inches thick. The first scintillating layer is axial, while the second and third are stereo at $+23.774^\circ$ and -24.016° , respectively. The lead cylinder provides approximately $1 X_0$ in the transverse plane, while the solenoid is $0.9 X_0$ thick, so particles pass through at least two radiation lengths before producing scintillation in the CPS. A large fraction of electrons and photons begin showering in the solenoid or the lead absorber, while heavier particles such as muons and hadrons are unlikely to interact strongly with either. Thus, the properties of a cluster of hits in the CPS may be used to identify an EM particle (photon, electron) and

to gauge roughly where an EM shower began.

Each FPS detector is a disk comprising eight 22.5° wedges. Each wedge contains four scintillating doublet layers, two on either side of a $2X_0$ lead absorber, and there is a stereo angle of 22.5° between the two layers on each side of the absorber. With this design, the inner two layers act as a coarse tracking detector, while the outer two serve the same purposes as the three scintillating layers in the CPS.

Liquid Argon Calorimeter

A large amount of material is necessary to fully absorb the energy of relativistic particles, especially hadrons. Several interaction lengths are necessary to contain a typical shower. For example, in iron, $X_0 \approx 2$ cm, and $\lambda_I \approx 17$ cm, so an iron calorimeter would need to be at least 1 m thick to fully capture hadronic showers. To reduce size, a substantially denser material is often chosen as the absorber, such as the depleted uranium in DØ's calorimeter.

In addition to the absorption of relativistic particles, a calorimeter must contain material that serves as an *active medium* to provide a signal indicating the amount of absorbed energy. Ideally, the same material would serve as both absorber and active medium, but this is quite difficult and expensive in practice. More commonly, a homogeneous calorimeter is approximated by alternating layers of absorber and active medium, so the active layers sample particle showers at regular intervals. This design is appropriately called a *sampling calorimeter*. It should be noted that the lead absorbers and scintillating strips of the preshower detectors are essentially highly granular sampling calorimeter layers, extending DØ's calorimetry to just before the liquid argon cryostats.

The primary sampling calorimeter in DØ uses liquid argon as the active medium. As shown in figure 2.8, liquid argon fills the gaps between absorber plates and readout boards. The absorbers are electrically grounded, and the readout board is kept at a positive potential of roughly 2 kV. Showers begin in the absorber plate, and any resultant charged particles ionize the liquid argon, sending a large number of electrons towards the readout board, arriving after a maximum drift time of 450 ns. This creates a pulse of current that relates directly to the energy in the shower. Taking into account these electrical signals from many

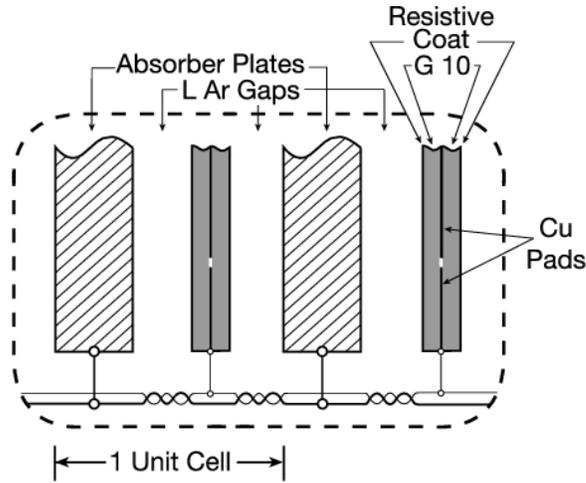


Figure 2.8: Small-scale structure of DØ's liquid argon sampling calorimeters [12].

adjacent calorimeter cells, we can accurately determine the total energy deposited by an incident particle.

DØ's liquid argon calorimeter is divided into three parts: one central calorimeter (CC), covering up to $|\eta| \approx 1$, and two endcap calorimeters (EC) that extend coverage to $|\eta| \approx 4$. Each calorimeter is enclosed by a cryostat responsible for maintaining the liquid argon at 90 K.

The CC and EC are further divided into electromagnetic (EM), fine hadronic (FH), and coarse hadronic (CH) sections, in order of proximity to the interaction region. Depleted uranium is the absorber for the EM calorimeter, with plates 3 (4) mm thick in the CC (EC). The FH sections use a uranium-niobium alloy in plates 6 mm thick, while the CH absorber plates are 46.5 mm thick and made of copper (CC) or stainless steel (EC). There are four EM readout layers in both the CC and EC, with a total thickness of roughly $20 X_0$. Outside the EM layers, there are three (four) FH readout layers in the CC (EC), providing roughly $3 (4) \lambda_I$. Outside the FH layers, a single CH readout layer amounts to $3.2 \lambda_I$ in the CC and as much as $6.0 \lambda_I$ in the EC.

The various calorimeter readout layers provide longitudinal granularity to our measurements of particle showers. In addition, transverse granularity is necessary to isolate individ-

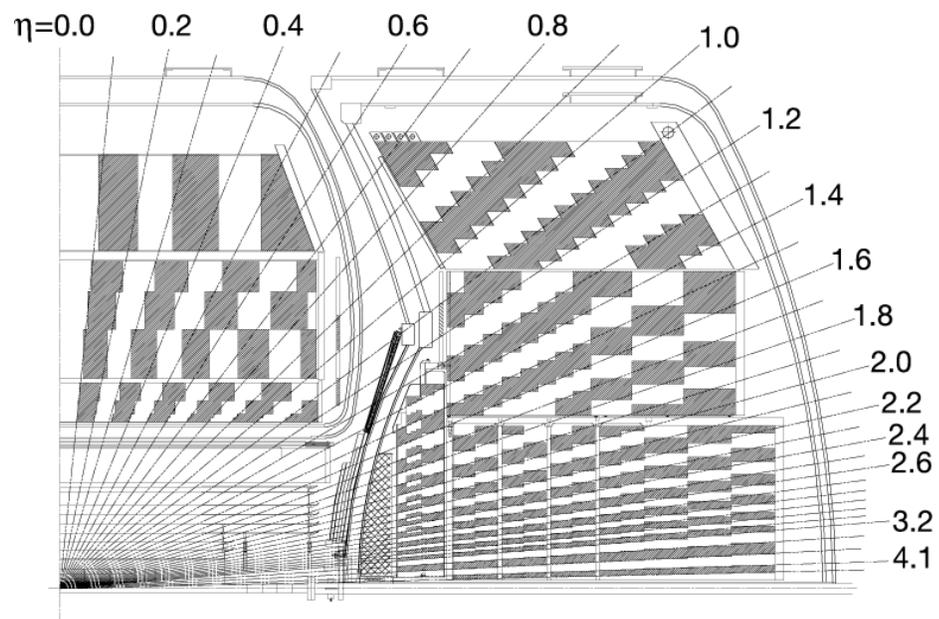


Figure 2.9: Segmentation of the DØ calorimeter, showing pseudo-projective cell towers [12].

ual showers and make use of identifying properties such as shower width. To achieve this, the readout boards are segmented into transverse cells 0.1×0.1 in $\eta \times \phi$, which is roughly the transverse size of a typical EM shower. In order to provide the best granularity where we expect maximum EM shower activity, the third EM readout layer is further segmented into 0.05×0.05 cells. As shown in figure 2.9, transverse readout cells in subsequent layers are offset to form pseudo-projective towers of cells, all originating at the nominal collision point. With this design, a particle created in hard scattering near $z = 0$ often deposits most of its energy in a single calorimeter tower.

Inter-Cryostat Detector

Because of the cryostat walls between the CC and EC calorimeters, the pseudorapidity region $0.8 < |\eta| < 1.4$ has poor calorimeter coverage and a large amount of insensitive material. As a result, the resolution of energy measurements is significantly degraded in the *inter-cryostat region* (ICR). To partially address this issue, an extra active layer was added in the ICR. The *inter-cryostat detector* (ICD) are two rings of 16 scintillating tiles each, mounted on the inner faces of the EC cryostats and covering $1.1 < |\eta| < 1.4$. Each scintillating tile is 0.5 inches thick and subtends a solid angle of 0.3×0.4 in $\eta \times \phi$. To achieve the same granularity as the liquid-argon calorimeter, each tile is divided into 12 0.1×0.1 subtiles, with independent readout for each subtile. Scintillation light produced in an ICD subtile travels through wavelength shifting fibers to a photomultiplier tube (PMT), where the pulse of light becomes an electrical signal.

2.3.3 Muon Tracking

DØ's various calorimeters will succeed in absorbing most particles resulting from high-energy $\bar{p}p$ collisions, but a few types of particles manage to penetrate through the many layers of dense material. Of course, the entirety of the detector is virtually transparent to neutrinos, but also high-energy muons will generally pass through the calorimeter. Although muons are electrically charged, their large mass suppresses bremsstrahlung for muons below several hundred GeV. Inelastic nuclear interactions are exceedingly rare for muons, since

they are neutral under the strong nuclear force, so muons deposit energy primarily through ionization. A typical muon loses roughly 3 GeV in the calorimeter, spread evenly along its trajectory.

A high-energy muon produces a track in the inner tracker and a small but collimated energy deposit in the calorimeter. Although this is a unique signature, muon identification is much improved by the addition of particle detectors outside the calorimeter. Also, while an inner track could be matched simply to a hit in the muon detectors, additional matching criteria can substantially improve this technique. Thus, an outer muon tracking system, composed of several sensitive layers and a dedicated magnet, provides an independent measurement of each muon's momentum and charge. The resolution of the muon tracker cannot compete with that of the SMT and CFT, but an outer muon track provides the information necessary to accurately identify the inner track belonging to a particular muon.

To bend the trajectories of muons emerging from the calorimeter, DØ employs a rectangular iron toroid magnet. The magnetic field of 1.9 T points roughly along ϕ , so an outer muon track bends to higher or lower η , depending on the charge of the muon. To measure this deflection, we compare track segments from several tracking layers before and after the toroid. There are four tracking layers before the toroid, collectively called layer A. After the toroid, layers B and C have three tracking layers each.

In the muon tracker, a muon's trajectory is measured by means of *proportional drift tubes* (PDT's). The design of a PDT is simple: a very thin wire is suspended in the center of a metal tube, with the wire held at a large positive electrical potential relative to the tube. A high-energy particle passing through the tube ionizes the gas inside, and the resulting free electrons drift toward the wire, producing an electrical pulse. With a very thin wire, the large electrical field immediately surrounding the wire causes drifting electrons to accelerate enough to further ionize the gas, resulting in a substantial signal gain. The pulse profile can be used to infer the particle's distance of closest approach to the tube axis, and the relative timing of the pulse as received on opposite ends of the tube can provide longitudinal hit resolution.

In the central part of the muon tracker, layers A, B and C are composed of relatively coarse PDT's with a maximum electron drift time of 500 ns. Their pseudorapidity coverage

extends approximately to $|\eta| = 1$. In the forward region up to $|\eta| = 2$, layers A, B and C have finer segmentation, using appropriately named *mini drift tubes* (MDT's) with a maximum drift time of 60 ns. Fine segmentation in the muon tracker is especially important in the forward region, where particle densities are highest. To optimize momentum resolution in both central and forward regions, the PDT's and MDT's are aligned to be roughly parallel to the toroidal magnetic field.

In addition to drift tubes, the muon tracker utilizes scintillating tiles to provide precise timing information for muon tracks. With response times of a just a few nanoseconds, scintillator hits can be spatially matched to PDT hits in order to associate the muon track with the correct collision. Without scintillator information, 500 ns of drift time in the central PDT's is longer than the 396 ns between successive bunch crossings, so there could be some ambiguity as to the origin of some central muons. Also, since the timing of each bunch crossing is precise to within nanoseconds, scintillator hits may be used to reject muons detected at unexpected times, likely originating from the beam halo or cosmic rays. For these purposes, two layers of scintillating tiles are used in the central part of the muon tracker, and three scintillation layers are used in the forward region.

2.3.4 Luminosity Measurement

Any analysis of the data collected by DØ is fundamentally a comparison between observation and expectation. At the most basic level, we need to be able to predict an expected number of events with a particular kinematic signature. As discussed in section 2.2.1, both the cross section and the beam luminosity are required to make this prediction. Theoretical calculations often provide the cross section, but the luminosity of the beams colliding at DØ must be measured using a dedicated particle detector.

DØ's *luminosity monitor* (LM) is composed of two disks mounted on the inner faces of the EC cryostats, occupying the space between the beam pipe and the FPS detectors. Each disk is made of 24 wedges of plastic scintillator, and scintillation light is detected by PMT's mounted on each wedge. Designed to be as sensitive as possible to inelastic $\bar{p}p$ collisions, the LM covers the forward pseudorapidity region $2.7 < |\eta| < 4.4$, where particle multiplicity

is particularly high for the majority of inelastic collisions.

The LM measures the timing of scintillator hits to within 0.3 ns, with the dominant contribution to this uncertainty coming from ambiguity in the location of a hit along the radial length of a wedge. We rely on this precision to determine the z -coordinate of the primary vertex, using the time difference between hits in the two LM disks. If $|z_{\text{PV}}| > 1$ m ($\Delta t > 6.7$ ns), then it's very unlikely (roughly 4σ) that the hits resulted from a $\bar{p}p$ interaction, so the hits are rejected as background from the beam halo.

$$L = f \frac{\bar{N}_{\text{LM}}}{\sigma_{\text{LM}}} \quad (2.7)$$

We calculate the luminosity L using equation 2.7. \bar{N}_{LM} is the average number of inelastic collisions detected by the LM per bunch crossing; σ_{LM} is the effective cross section for inelastic $\bar{p}p$ collisions, taking into account the imperfect acceptance and efficiency of the LM [16]; and f is the bunch crossing frequency. Since there are often multiple inelastic collisions within a single bunch crossing, we directly measure the fraction of bunch crossings without coincident LM hits and use poisson statistics to infer \bar{N}_{LM} .

2.3.5 Triggering

The average bunch crossing frequency at DØ is 1.7 MHz. Due to technological and financial constraints, it is not possible to record every bunch crossing. Furthermore, the vast majority of $\bar{p}p$ collisions offer very little of interest to most analyses at DØ. For these reasons, we use a sophisticated series of filters to quickly choose the most promising collisions and ignore the rest. We refer to these filters as *triggers*. Approximately one event out of 17,000 passes the three levels of triggering at DØ, so we record events for further analysis at a rate of 100 Hz. Figure 2.10 provides a diagram of data flow through the trigger system.

Level 1 Trigger

The first level of triggering uses coarse information from the calorimeter, CFT, preshower detectors, and muon tracker to accept data at a rate of 2 kHz. Level 1 (L1) buffers hold enough data to allow $3.7 \mu\text{s}$ for the subdetectors to provide the necessary information for

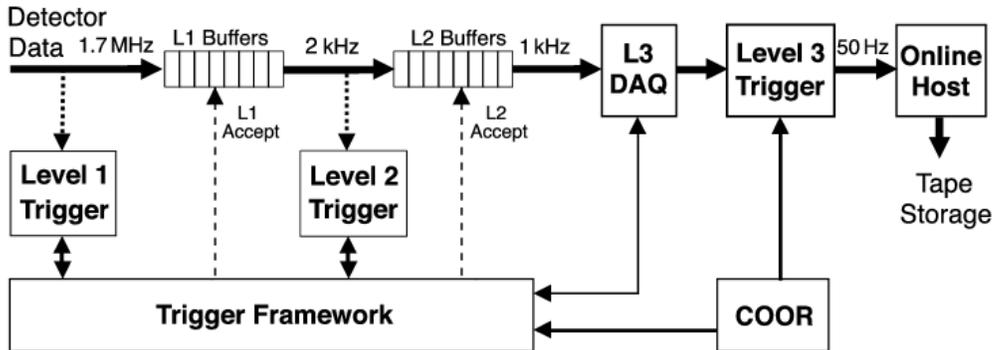


Figure 2.10: Flow of data through the DØ trigger system [12].

a L1 decision, and for that decision to be made. The time necessary for readout from L1 to the level 2 trigger causes some events to be dropped immediately after every L1 accept, which places an upper limit on the L1 accept rate. 2 kHz was chosen to keep this deadtime below 5%.

L1 calorimeter (L1Cal) triggers use fast estimates of the EM and hadronic energy in $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$ towers to construct simple quantities such as the total transverse energy (E_T) or the number of towers above a given E_T threshold. Thresholds may be applied to groups of adjacent towers to roughly trigger on jets.

The L1 central track trigger (L1CTT) compares hits in each 4.5° transverse sector of the CFT to a library of approximately 20,000 predefined tracks. The six highest- p_T track candidates in each sector are matched to clusters in the CPS, and thresholds may be applied to quantities such as the total number of tracks and the track p_T . In addition, thresholds may be applied to the number of clusters in the FPS.

L1Muon triggers match hits between the various muon scintillator and PDT layers and between outer muon tracks and L1CTT tracks. Similarly, the L1 calorimeter-track (L1CalTrk) triggers match L1Cal information to L1CTT tracks. Various p_T thresholds may be applied to the matched objects. L1CalTrk triggers were added during the transition from RunIIa to RunIIb.

Level 2 Trigger

At the second level of triggering, the increase in decision time over L1 allows for the use of more detailed information and more sophisticated algorithms. Subdetector-specific processors work in parallel to construct objects for use in the global level 2 (L2) trigger, where for the first time, events are filtered according to correlations across the entire DØ detector. The L2 filters decrease data flow by a factor of two, so the L2 accept rate is approximately 1 kHz.

While L1 relied exclusively on the CFT and preshower detectors for tracking, L2 also makes use of hits in the SMT. The L2 silicon track trigger (L2STT) takes input from the SMT and from L1CTT, refining L1 tracks with SMT information. This improves the track p_T resolution and enables identification of long-lived particles through measurement of the transverse distance between tracks and the beam line.

Level 3 Trigger

The level 3 (L3) trigger is the only stage of triggering without a dedicated data path; instead of using a subset of data to make trigger decisions, L3 can use the full information content of each event accepted by L2. Using a large amount of this information, the L3 trigger makes decisions based on objects such as jets, electrons and muons that are nearly as sophisticated as those used in offline analyses. For the first time, tracks are used to locate the primary vertex, and the location provides improved calorimeter-based E_T measurements. Identification of secondary vertices is also possible, so b -jet tagging can be implemented in the L3 trigger. Filters may be applied to single objects or to variables relating several objects, resulting in an accept rate of approximately 100 Hz. Events accepted by the L3 trigger are recorded for future analysis.

L3 trigger processing is accomplished by a linux farm with multiple CPU's in each farm node. Currently, the L3 farm contains 288 nodes, each with either four or eight CPU's. When an event is accepted by the L2 trigger, the many fragments of the event are sent from all detector subsystems to the same farm node, where the fragments are merged into a complete event. Completed events then proceed through object construction and filtering,

with each event assigned to a particular CPU and all CPU's working in parallel. Thus, currently the L3 farm can process over 1600 events simultaneously. Assuming no deadtime at L3, the 1 kHz input rate to this massively parallel system translates to a maximum decision time of more than one second.

Chapter 3

RECONSTRUCTION

An essential and non-trivial first step in analyzing the data collected by DØ is to translate the raw information provided by the detector to objects more closely related to the physics we aim to study. This process has many steps; for example, in the CFT alone, hits within a doublet layer must first be interpolated, then axial and stereo hits must be combined to form clusters, and clusters in various layers must be connected to form a track. Often considering data from various other subdetectors as well, the *reconstruction* of particle trajectories and energy deposits becomes quite complicated, and accurately identifying the particle type from this information adds further complexity.

In this chapter, I give a condensed description of the reconstruction of the objects and particles of physical interest in the $ZH \rightarrow \mu\mu bb$ analysis. First, reconstruction of charged particle tracks is described in section 3.1. Section 3.2 explains how these reconstructed tracks are used to identify the primary vertex. Connecting inner tracks to hits in the outer muon tracker, section 3.3 describes muon reconstruction. Section 3.4 explains the reconstruction and identification of jets from calorimeter signals, and section 3.5 describes how we use tracks to identify jets originating from b quarks.

3.1 Tracking

A hit in any tracking element of $D\phi$ provides information on both the position and the amount of ionization or scintillation light produced by a high-energy charged particle. Particles traversing a scintillating fiber or a PDT near the edge of its active volume result in smaller signals than particles traversing near the center of the fiber or tube. Similarly, a particle traveling through a silicon wafer at an oblique angle will ionize many more silicon atoms than a particle at normal incidence. When a particle produces multiple hits in a single tracking layer, we interpolate between these hits, weighting each hit according to the charge collected by the silicon strip, VLPC, or drift tube. With this technique, we can effectively improve the hit resolution beyond the physical dimensions of the detectors. The information contained in the amount of collected charge is used throughout the process of track reconstruction to yield measurements as precise as possible.

To identify tracks from a very large collection of hits, we employ two optimized track-finding algorithms: the *histogramming track finder* (HTF) [17] and the *alternative algorithm* (AA) [18]. The use of two independent algorithms is important to achieve maximal efficiency. The HTF and AA excel in somewhat different regions of phase space, and they rely on somewhat different assumptions.

Assuming tracks originate from the interaction point, the HTF performs a Hough transformation to a phase space defined by the curvature $1/p_T$ and the initial azimuthal angle ϕ of helical tracks. Because there exists an infinite spectrum of tracks able to connect the interaction point with a single hit, the HTF maps the transverse coordinates of a hit in the SMT or CFT to a line in this phase space. If several hits belong to the same track, the corresponding lines should intersect at a single point $(1/p_T, \phi)$. The coordinates of this point are the transverse properties of the track responsible for these hits. Our implementation of this algorithm divides $1/p_T \times \phi$ space into discrete bins and fills the resultant histogram with tracking hits. Then, we identify track candidates by selecting narrow peaks in the histogram.

Tracks originating from decays of long-lived particles may not be found with the HTF algorithm, since the intersection of hits from such tracks may spread over several $1/p_T \times \phi$

bins, as the assumption that the track originates from $(x, y) = (0, 0)$ is false. For these tracks, the AA outperforms the HTF. Starting with a seed track of three SMT hits, the AA extrapolates the track outwards to the next layer and attempts to find hits within a narrow window around the extrapolation. Hits are added to the track if they increase the fit χ^2 by less than 16. Layers in which no such hits are found are counted as misses, and multiple hits in a window spawn multiple track candidates. The algorithm proceeds iteratively through all the layers of the inner tracker, or until it counts three missed layers. Since the origin of the track was never assumed to be close to the interaction point, SMT-based AA tracks may come from secondary vertices. The alternative algorithm can also begin with a seed of three CFT hits, in which case the PV is added as an assumed fourth hit, and extrapolation proceeds inwards from the CFT.

The HTF and AA provide idealized, coarse track candidates. These algorithms do not account for energy lost to ionization in the tracking layers and in the bulk material that supports the trackers. Also, multiple scattering may cause random directional changes to the particle trajectories. To deal with these effects, we apply a Kalman filter that fits a sophisticated, non-helical track model to the hits associated with each HTF or AA candidate [19]. We reject track candidates incompatible with this model by placing an upper bound on the fit χ^2 , thereby reducing the occurrence of fake tracks, while tracks passing the Kalman filter benefit from improved precision and accuracy in comparison to the coarse HTF and AA tracks.

3.2 The Primary Vertex

The primary vertex is of critical importance to our interpretation of the detector signals that result from a $\bar{p}p$ collision. To varying degrees, every detailed measurement of the particles in a hard-scattering event relies on precise determination of the PV location. In particular, translating a calorimeter energy deposit to an E_T measurement depends on the particle's pseudorapidity, and the true value of η can only be determined with precise knowledge of z_{PV} . Also, the fit of a track to hits in DØ's various tracking detectors may be improved by constraining the track to pass through the PV, and precise resolution of the distance between displaced tracks and the PV allows for efficient identification of long-lived particles.

Finding the PV is complicated by the existence of other vertices in the event. Especially at high luminosity, it's common for several $\bar{p}p$ interactions to occur in the same bunch crossing, but usually only one interaction involves a hard scatter. We refer to the others as *minimum bias* (MB) interactions, and particles from such vertices are a significant instrumental background. We reduce this background through accurate identification of the PV and association of every detected particle with the correct vertex.

Reconstruction of the PV has two major steps: identify all vertices consistent with $\bar{p}p$ collisions [20], then choose the vertex least consistent with the MB hypothesis [21]. First, tracks are selected using these criteria:

- $p_T > 0.5$ GeV;
- at least two SMT hits; and
- $DCA/\sigma_{DCA} < 5$, where DCA is the *distance of closest approach* of the track to the beam line, and σ_{DCA} is its associated uncertainty.

The selected tracks are clustered according to the z -coordinate where the DCA is realized, z_{DCA} . Starting with the highest- p_T track, other tracks are added to its z -cluster if they satisfy $|z_{DCA} - \bar{z}| < 2$ cm, where \bar{z} is the mean z -coordinate of the tracks already in the cluster. In the next iteration, the algorithm builds another cluster around the highest-

p_T unclustered track, and so on until all tracks are contained in several non-overlapping z -clusters. At this point, any trivial clusters containing only one selected track are rejected.

Kalman fits are performed to locate the vertex within each z -cluster. If the fit χ^2 per degree of freedom (χ^2/n_{DF}) is larger than 10, then the track contributing most to the χ^2 is removed. This process repeats until $\chi^2/n_{\text{DF}} < 10$ or until the cluster contains only one track. In this way, the list of z -clusters becomes a list of PV candidates.

To choose the correct PV from the list of candidates, we rely on the observation that tracks from MB interactions generally have smaller p_T than tracks from hard-scatter vertices. More formally, we use the expected p_T distribution of tracks matched to MB vertices, obtained from Monte Carlo simulation, to calculate a MB probability for each track. Individual track probabilities are combined to yield an MB probability for each vertex. To take into account high- p_T displaced tracks from the decay of heavy, long-lived particles, we consider all tracks satisfying $\Delta z < 2$ cm relative to each vertex. Finally, the PV is identified as the vertex with the lowest MB probability.

3.3 Muons

We can identify muons by virtue of the fact that they penetrate the calorimeter and leave hits in the outer muon tracker. Outer muon tracks are composed of straight-line segments in the three tracking layers A, B, and C. While trajectories curve in the toroid between layers A and B, segments in layers B and C belonging to the same track should be collinear. For this reason, in muon reconstruction the B and C layers are merged. The loose muon definition used in the $ZH \rightarrow \mu\mu bb$ analysis requires at least one scintillator hit and two PDT or MDT hits in either the A layer or the merged BC layer [22]. To reject cosmic muons, the scintillator hit must match the bunch crossing time plus time of flight to within 10 ns.

Momentum measurement with the outer tracker is only possible if we identify segments both before and after the toroid. Since we do not require hits on both sides of the toroid in the $ZH \rightarrow \mu\mu bb$ analysis, we rely entirely on the inner tracker to measure muon momentum. Every muon candidate in the outer tracker must be directionally matched to an inner track, which provides all the kinematic information for the reconstructed muon. We reject inner tracks with $DCA > 0.02$ cm if the track has at least one SMT hit, and reject $DCA > 0.2$ cm otherwise.

Muons are common constituents of jets, as they result from leptonic and semileptonic decays of hadrons. In order to reject such muons in favor of those produced in the decay of a Z boson, we apply isolation criteria to reconstructed muons. We express isolation criteria in terms of two variables that are calculated during reconstruction: the *track halo* and the *calorimeter halo*. The track halo is the scalar p_T sum of all other tracks in the cone defined by $\Delta R < 0.5$ with respect to the muon. The calorimeter halo is the scalar E_T sum of all calorimeter cells satisfying $0.1 < \Delta R < 0.4$ with respect to the muon. To account for the fact that a high- p_T muon is less likely to be the product of a hadron decay, this analysis employs *scaled* isolation variables obtained by dividing the track and calorimeter halos by the p_T of the muon.

3.4 Jets

Jets are arguably the most complicated reconstructed objects. Resulting from the soft and collinear showering of final-state quarks and gluons, each jet contains a multitude of particles of various types, and each particle type leaves its own characteristic signature in the calorimeter. The fraction of jet energy carried by particles of a specific type is randomly distributed about an average value that depends on the particle type, jet flavor, and jet energy. This randomness, combined with large fluctuations in subsequent hadronic showers in the calorimeter, fundamentally limits the precision with which we can measure jet kinematics.

An important but difficult aspect of jet reconstruction is the concept of *infrared* (IR) *safety*. While perturbative calculations in QCD provide a good model of the interactions and evolution of high- p_T quarks and gluons, proper modeling of low- p_T quarks and gluons would require a non-perturbative treatment. Currently, these non-perturbative effects are not well understood, so they are likely to be poorly modeled in simulation, and we have no way to quantify the uncertainty in this aspect of jet modeling. We therefore design our jet algorithm to be IR-safe: reconstructed jets should be independent of these poorly understood, non-perturbative effects. While jet reconstruction at $D\emptyset$ is not truly IR-safe, efforts towards approximate safety have lent substantial complexity to our jet algorithm.

Since a random but significant fraction of particles in a jet are neutral, jet reconstruction is necessarily based on calorimeter information. First, we build pseudo-projective towers from calorimeter cells containing a sufficient amount of energy, then we combine towers into *preclusters*, which serve as seeds in the formation of *protojets*. Protojets are split or merged to avoid double-counting of energy, and the resultant jets must pass several quality cuts before use in a typical analysis. Finally, selection criteria may be applied to tracks associated with calorimeter jets in order to reject those jets that are the products of MB interactions. Selected jets must also be corrected to account for nonlinear calorimeter response and other shortcomings of the detector and jet algorithm. The jets used in the $ZH \rightarrow \mu\mu bb$ analysis are products of these steps; the final step of b -tagging is described in the next section.

We must ensure that the reconstructed transverse kinematics of a jet are independent of the jet's longitudinal boost, so in this section ΔR is defined using the rapidity y instead of the pseudorapidity η . This distinction is important, as the preclusters used to reconstruct jets generally have non-negligible mass.

3.4.1 Identification

To begin, we describe calorimeter cells using four-momenta relative to the reconstructed PV. Cells are approximated as pointlike energy deposits at the centers of the actual cells, which results in the four-momenta being massless ($E = |\vec{p}|$). In the first of many steps to reduce the impact from calorimeter noise, *zero-suppression* rejects cells with energies less than 2.5 times the RMS width (σ) of cell-specific electronic noise. Since real particle showers generally cover several cells, we use two algorithms to reject isolated cells. The NADA algorithm rejects any particularly energetic cell if its neighbors are not also sufficiently energetic [23], while the T42 algorithm rejects cells where $E < 4\sigma$ unless an adjacent cell has $E > 4\sigma$ [24]. Also, T42 rejects any cells with negative energies as a result of online pedestal subtraction. Cells satisfying these criteria are selected to build towers.

A tower is the four-momentum sum of all selected cells in one of the pseudo-projective towers described in section 2.3.2. The four-momenta are defined with respect to the PV, but cells in a tower extend outward from $z = 0$, so a tower containing more than one cell has nonzero mass if $z_{\text{PV}} \neq 0$.

From multiple towers, we build preclusters using a simple cone algorithm. In this stage of jet reconstruction, coarse hadronic cells are removed from the towers because the CH layers are the most noisy in the calorimeter. The CH-subtracted towers are sorted by p_{T} , and we begin constructing preclusters starting with the highest- p_{T} tower. In order to initiate a precluster, a tower must have $p_{\text{T}} > 0.5$ GeV. With this criterion satisfied, other towers are added to the precluster if they have $p_{\text{T}} > 1$ MeV and fall within a $\Delta R < 0.3$ cone about the initial tower. After the first precluster is built, this algorithm repeats on the remaining unclustered towers until the tower list is depleted. Finally, preclusters with $p_{\text{T}} < 1$ GeV or containing only one tower are rejected.

The selected towers and preclusters are inputs to DØ's RunII Cone Algorithm, essentially an iterative midpoint cone algorithm with cone radius $\Delta R = 0.5$ [25]. This algorithm begins by clustering towers around precluster seeds. Starting with the highest- p_T precluster, we build a protojet as the sum of all towers satisfying $\Delta R < 0.5$ from the precluster seed. The protojet is then used as the seed for a new protojet, and this clustering procedure iterates until ΔR between old and new protojets is less than 0.001. After a stable protojet is produced, the algorithm moves on to the next highest- p_T precluster seed satisfying $\Delta R > 0.25$ from all existing protojets. New protojets are required to differ by at least 1% in p_T and 0.005 in ΔR from all existing protojets. The algorithm proceeds as described until the list of precluster seeds is depleted. Finally, a second set of protojets is seeded by the midpoints between protojet pairs satisfying $0.5 < \Delta R < 1.0$. Each midpoint is defined as the p_T -weighted average of (y, ϕ) for the protojet pair. Consideration of protojets built from midpoint seeds alleviates sensitivity to soft gluon emission between two closely-separated jets, but our jet definition remains sensitive to soft gluon emission in the vicinity of three or more jets.

With the full list of protojets from both precluster and midpoint seeds, many towers belong to multiple protojets. The final stage of the RunII Cone Algorithm avoids double-counting of energy by assigning each tower to only one jet, resolving any ambiguities in tower ownership. If two protojets share at least one tower, the shared p_T is compared to the smaller protojet p_T . If the shared p_T is greater than half the smaller protojet p_T , then the smaller protojet is merged into the larger. Otherwise, the shared towers are split between the two protojets, each tower belonging to the closest protojet. This is also an iterative process: after a split or merger, the list of protojets is updated, and the process continues until there is no overlap between protojets. Finally, protojets with $p_T > 6$ GeV are identified as jets.

3.4.2 Selection

To minimize the misidentification of electrons, photons, and calorimeter noise as jets, we place further requirements on the jet quality [26]. Because the hadronic calorimeter layers

are more noisy than the EM layers, we reject jets with less than 5% of their energy in the EM layers. The coarse hadronic layers are the most noisy, so we reject jets with more than 40% of their energy in the CH layers. To discriminate against electrons and photons, we reject jets with less than 5% of their energy in the hadronic layers. A persistent problem in the $D\bar{O}$ calorimeter has been the phenomenon of *hot cells*, in which an isolated cell occasionally experiences much more readout noise than usual. To mitigate the impact of hot cells on jet reconstruction, we require that no single cell in a jet contains 90% of the jet energy, and that the most energetic cell in a jet contains less than ten times the energy of the second most energetic cell. We refer to jets that pass all of these criteria as *good jets*.

In the high-luminosity collision environment of RunIIb, we must also be concerned about the effect of MB vertices on jet reconstruction. The most common high- p_T product of a MB interaction is a jet, but we are only interested in jets originating from the PV. To reject MB jets, we select tracks pointing toward the location of the jet in the calorimeter and require at least two of these tracks to originate from the PV. Jets that pass this requirement are said to be *vertex-confirmed*.

3.4.3 Calibration

Having selected the jets to be analyzed in our data, we must correct the reconstructed jet energies to more accurately reflect the total energy of the particles constituting each jet. This correction is known as the *jet energy scale* (JES), and it has several distinct pieces to account for distinct sources of inaccuracy in the measurement of jet energy [27]. Equation 3.1 shows how the particle-level energy E_{JES} is obtained from the directly measured jet energy E_{raw} via the JES correction.

$$E_{\text{JES}} = \frac{E_{\text{raw}} - O(\eta, N_{\text{MB}})}{R(E_{\text{raw}}, \eta) \times S(E_{\text{raw}}, \eta)} \quad (3.1)$$

$O(\eta, N_{\text{MB}})$ is an offset energy that covers the effects on the measured jet energy from electronic noise, radioactive decays of the uranium in the calorimeter, low-energy particles from MB interactions, and energy remaining in the calorimeter from previous bunch crossings. The offset is typically close to 1 GeV in the CC and grows to several GeV in the EC.

It increases linearly with the number of MB vertices, N_{MB} .

$R(E_{\text{raw}}, \eta)$ is the response of the calorimeter to a particle jet, parametrized as a function of the measured energy and pseudorapidity region in the detector. As shown in equation 3.2, R may be further decomposed into the response to the EM part of the jet, R_{EM} , and the response to everything else, R_{else} . (The potential η -dependence of each response is suppressed.)

$$R(E_{\text{raw}}) = f_{\text{EM}}(E_{\text{raw}}) \times R_{\text{EM}} + (1 - f_{\text{EM}}(E_{\text{raw}})) \times R_{\text{else}} \quad (3.2)$$

To a good approximation, R_{EM} and R_{else} are independent of the jet energy. In fact, assuming the calorimeter is already calibrated to reproduce the Z mass peak in $Z \rightarrow e^+e^-$ data, $R_{\text{EM}} = 1$. The factor f_{EM} is the fraction of the jet's energy deposited by electrons and photons. This EM fraction increases with jet energy because every nuclear interaction above a certain threshold typically results in π^0 mesons carrying away 1/3 of the incident hadron's energy, and each π^0 immediately decays to two photons. Increasing the jet energy generally allows for more nuclear interactions above the π^0 threshold, so f_{EM} increases asymptotically to 1. Because energy is lost in nuclear interactions and through hadron decays involving muons or neutrinos, $R_{\text{else}} < 1$. The total jet response R increases with jet energy, and typical values range between 0.6 and 0.8 for jet energies between 20 and 200 GeV.

The final piece of the JES, $S(E_{\text{raw}}, \eta)$ is a correction to the response that accounts for the effects of calorimeter showers extending across the boundary of the jet cone. Showers initiated by particles within the jet cone may not be entirely contained by the cone, and likewise particles outside the cone may deposit energy inside the cone boundary. The showering correction S corrects for these two effects simultaneously. Because jet width decreases with increasing jet energy, the highest values of S occur at low energy. Due to high particle densities in the forward regions, contamination from particles outside the jet cone can be especially problematic for forward jets, so S decreases with increasing jet $|\eta|$. The effect of the showering correction is always less than 5%.

Figure 3.1 shows the ratio of raw jet E_T to corrected jet E_T , having applied all components of the JES correction. The asymptotic behavior of the response is apparent: the

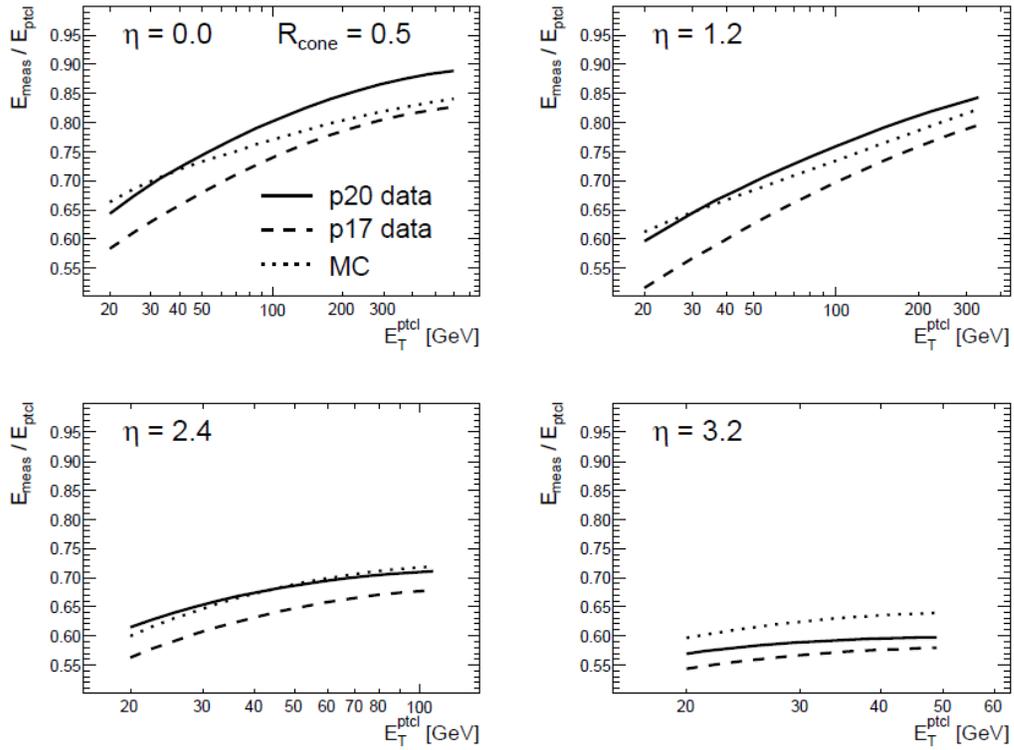


Figure 3.1: The ratio of the raw jet E_T (E_T^{meas} in the figure) to the JES-corrected jet E_T (E_T^{ptcl} in the figure), versus the corrected jet E_T , for four values of jet η : 0.0, 1.2, 2.4, and 3.2. The solid curves represent jets in the RunIIb data, the dashed curves represent jets in the RunIIa data, and the dotted curves represent simulated jets [28].

response generally increases with jet energy, but the slope of the curve decreases as the response approaches unity. Another striking feature of these plots is the pronounced decrease in response in the forward region. This is due to more stringent energy thresholds for forward calorimeter cells, intended to reject energy deposits from the beam halo.

3.5 b Tagging

Low-mass Higgs bosons ($m_H < 135$ GeV) are expected to decay primarily to a pair of bottom (b) quarks. The properties of b quarks are quite unique, and we exploit this uniqueness to optimize Higgs sensitivity. Due to their mass being substantially less than that of the W boson, and because the CKM matrix elements V_{ub} and V_{cb} are small in magnitude, decays of b quarks are heavily suppressed. More correctly, speaking in terms of the low-energy degrees of freedom in QCD, hadrons with valence b -quarks have relatively long lifetimes. Bottom hadrons often travel a few mm from the interaction point before decaying, and any charged decay products will produce tracks substantially displaced from the primary vertex. Furthermore, the large mass of the b quark makes b hadrons more massive than other typical jet constituents. Finding displaced tracks and high-mass secondary vertices matched to reconstructed jets allows us to identify jets originating from b quarks [29]. Rejection of other jet flavors (u , d , s , c , and g) dramatically reduces background to a Higgs signal, so the *tagging* of b jets is one of the most important elements of any low-mass Higgs search, such as the $ZH \rightarrow \mu\mu bb$ analysis.

Reconstructed jets are calorimeter-based objects, but b tagging relies on tracking information. Thus, the first step in b tagging is to associate tracks with the energy deposits that define a jet. To do this, we first build track jets using the Snowmass jet algorithm [30] with cone radius 0.5, clustering tracks instead of calorimeter towers. The tracks selected for track jets must have at least one SMT hit, $p_T > 0.5$ GeV, $\text{DCA} < 2$ mm, and $|z_{\text{DCA}} - z_{\text{PV}}| < 4$ mm. Additionally, track jets are required to contain at least two tracks. Angular matching is performed between calorimeter and track jets, and $\Delta R < 0.5$ qualifies a jet as *taggable*. With each taggable jet, we proceed to identify displaced tracks and secondary vertices in the associated track jet.

The presence of a b hadron in a jet leaves many identifiable signatures in the associated track jet. To fully exploit all of the available information, $D\mathcal{O}$ combines the discrimination power of several track jet variables in an artificial *neural network* (NN). Trained with large numbers of simulated b jets and light (u , d , s , and g) jets, the NN b -tagger provides an output with substantially more discrimination power than any of its input variables. Ranked in

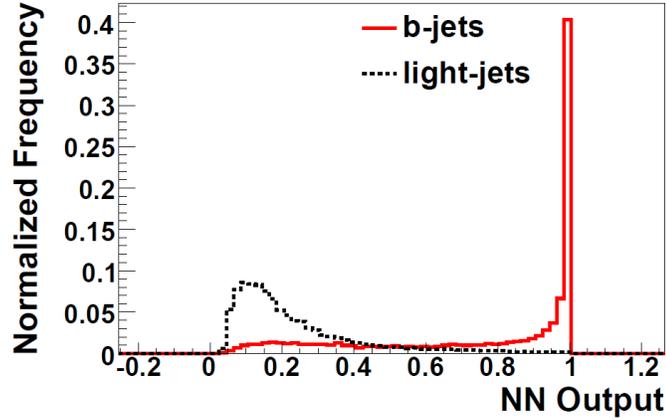


Figure 3.2: The output of the NN b -tagger for simulated b jets and light-flavor jets [29].

order of importance, the following seven variables are the inputs to the NN:

1. the significance of the transverse distance from the PV to the most significant secondary vertex;
2. a combined variable counting the number of tracks that pass various thresholds on the impact parameter significance;
3. the probability that all tracks in the jet originate from the PV;
4. the fit χ^2/n_{DF} of the most significant secondary vertex;
5. the number of tracks emerging from the most significant secondary vertex;
6. the invariant mass of the most significant secondary vertex; and
7. the number of secondary vertices found within the track jet.

As shown in figure 3.2, the NN output peaks at 1 for b jets and near 0 for light jets. Because c hadrons have lifetimes approaching those of b hadrons, the NN output for c jets

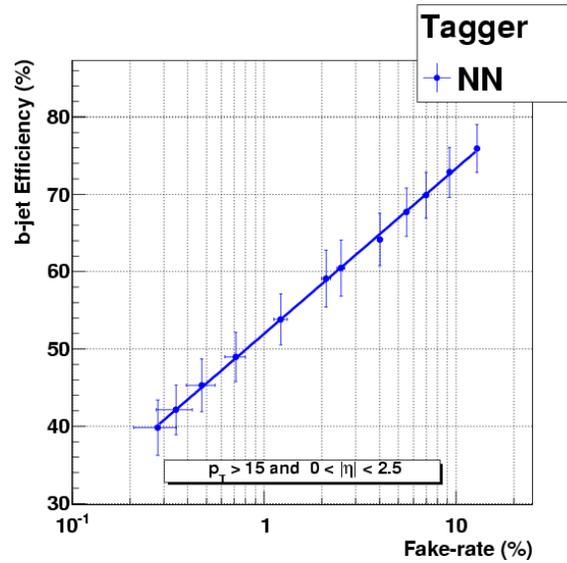


Figure 3.3: The efficiency of the NN b -tagger versus the light jet misidentification rate for jets with $p_T > 15$ GeV and $|\eta| < 2.5$.

has peaks near both 0 and 1. A jet is tagged and subsequently treated as a b jet if its NN output is larger than a specified cut value. For the $ZH \rightarrow \mu\mu bb$ analysis, we use loose and tight tag definitions, corresponding to cut values of 0.2 and 0.85. The b jet identification efficiency and light jet misidentification rate both depend strongly on the jet p_T and η , but figure 3.3 shows average values for jets with $p_T > 15$ GeV and $|\eta| < 2.5$. The loose cut is 70% efficient with a misidentification rate of 7%, while the tight cut is 45% efficient with a misidentification rate of 0.5%.

Chapter 4

SIMULATION

Fundamentally, a scientific experiment is a comparison between theoretical predictions and observed data. In particle physics, this translates to comparing expected and observed distributions of kinematic variables such as muon p_T or the angle between two jets resulting from a $\bar{p}p$ collision. Ideally, we would derive the expected distributions analytically and compare the observed data to smooth curves representing the theoretical prediction, but in practice analytical models including detector effects are unfeasible. Instead, we simulate large numbers of individual $\bar{p}p$ collisions and fill histograms to construct the expected kinematic distributions. To approximate smooth distributions, our samples of simulated events should be large enough to have negligible statistical uncertainty in comparison to the data. Using the integrated luminosity of the data and the expected cross section of the simulated process, we normalize each simulated distribution to put the data and simulation on equal footing.

It is useful to have a modular approach towards simulation, breaking the procedure into several independent pieces. For example, the specific process responsible for a final-state quark is effectively unrelated to the showering of the resulting hadrons in the calorimeter, so these two aspects of simulation may be treated independently. Beginning with a specified hard-scatter interaction, each simulated event is constructed in multiple steps, as follows.

1. *Monte Carlo* (MC) generators (section 4.1) model the hard-scatter process and the final state's evolution into long-lived particles.
2. To simulate the detector (section 4.2), we first use MC methods to model particle interactions with bulk material and magnetic fields. Then, data-based *minimum bias overlay* is added to each simulated event, modeling beam conditions and various luminosity-dependent detector effects. The simulated detector signals are digitized to

mimic the raw detector data.

3. Physical objects such as muons and jets are reconstructed with the same software as is used for data.
4. Finally, corrections are applied to the reconstructed objects (section 4.3) in simulated events in order to account for shortcomings in the detector simulation.

4.1 Monte Carlo Event Generators

Monte Carlo event generators work by assigning a probability distribution to the many-dimensional phase space associated with a specified theoretical process, then randomly sampling that distribution. This technique is applied to the hard scatter, in which the momenta of initial and final-state particles are randomly sampled, and to other processes such as initial- or final-state radiation, hadronization, and decay of short-lived particles. The two MC generators used in the $ZH \rightarrow \mu\mu bb$ analysis are PYTHIA and ALPGEN, described in the following subsections. To model initial-state parton kinematics for both MC generators, we used the CTEQ6L1 parton distribution functions [31].

4.1.1 PYTHIA

PYTHIA is a leading-order (LO) MC event generator [32], meaning that theoretical processes are simulated using only the lowest-order terms in the corresponding perturbation series. A hard-scatter process is specified, and a random sampling of the corresponding phase space determines the kinematics of initial- and final-state particles. Radiation of photons and gluons from the initial state and the final state are modeled using the *parton shower* method. With this method, jet substructure is well-modeled, but reconstructed jet multiplicity is less accurate. In other words, parton showers model soft and collinear radiation very well, but hard radiation is better modeled by including it as part of the hard scatter. For hadronization of final-state quarks and gluons, PYTHIA uses a model that is conceptually very similar to the heuristic explanation of confinement given in section 1.1.1. Effectively, quarks are joined by gluon “strings” to form hadrons. Stretching a string results in a linear increase in potential energy associated with the string, until that energy becomes large enough to produce a quark-antiquark pair that breaks the string, forming two new hadrons.

4.1.2 ALPGEN

ALPGEN is also a leading-order generator [33], but unlike PYTHIA, it employs a *matrix element* method to model both the hard scatter and initial- and final-state radiation. When

using ALPGEN, we specify a hard-scatter process and an associated number of radiated quarks or gluons. By generating several MC samples with different specified numbers of radiated partons, we achieve an accurate model of jet multiplicity, for example in Z +jets events. For brevity, I'll denote these samples as $Z+n_{\text{LP}}$, where LP stands for "light parton." In the $ZH \rightarrow \mu\mu bb$ analysis, we use four such ALPGEN samples, corresponding to $n_{\text{LP}} = 0, 1, 2,$ and 3 .

Since ALPGEN only provides final-state partons, we use PYTHIA to simulate the formation of jets through showering and hadronization. In this process, parton showering in PYTHIA may result in additional jets, but we may have already accounted for additional final-state radiation by generating an ALPGEN sample with a higher value of n_{LP} . To avoid double-counting of radiated jets, we use the MLM matching scheme to determine whether jets correspond to the original final-state partons from ALPGEN [34]. While we allow additional final-state radiation in the ALPGEN sample with the highest n_{LP} , in every other ALPGEN sample we reject events containing any jets not matched to ALPGEN partons.

Parton showering in PYTHIA may also produce heavy-flavor (HF) jets, but we already account for HF production by generating $Z + b\bar{b} + n_{\text{LP}}$ and $Z + c\bar{c} + n_{\text{LP}}$ ALPGEN samples. Therefore, to avoid double counting of HF jets, we reject events containing b jets in the $Z + c\bar{c} + n_{\text{LP}}$ samples and events containing b or c jets in the $Z + n_{\text{LP}}$ samples. This procedure is known as HF skimming.

4.2 *Detector Simulation*

As long-lived particles travel through the detector, they interact with bulk material, both active and inactive, and with magnetic fields. Indeed, without these interactions there would be no way of discerning the presence of these particles! We must therefore simulate the interactions between long-lived MC particles and the various components of the detector. Also, beam conditions and detector effects related to those conditions are difficult to simulate, so we overlay actual minimum-bias data onto our MC events to account for these effects. Finally, we have no direct simulation of the $D\bar{O}$ trigger system, so we apply trigger efficiencies as event weights to model imperfect trigger acceptance.

4.2.1 *GEANT*

To simulate interactions between particles and detector components, we rely on a detailed model of detector composition and geometry, as well as precise knowledge of the solenoidal and toroidal magnetic fields. The software we use to evolve particles through this detector model is based on GEANT3 [35]. It uses Monte Carlo methods, randomly sampling the phase space associated with several distinct interactions: ionization of bulk material, EM and hadronic showering, and charged particle deflection in a magnetic field. Simulated ionization of active detector material is translated to detector signals, providing the raw information that is used to reconstruct physical objects.

4.2.2 *Minimum Bias Overlay*

Every event generated by MC represents only the collision of a single proton with an antiproton. In actuality, the proton and antiproton are constituents of beams containing a huge number of other protons and antiprotons, and multiple $\bar{p}p$ collisions often occur during each bunch crossing. These additional collisions also produce detector signals, and this minimum-bias background is dependent on instantaneous luminosity. To model this and other luminosity-dependent effects, we collect data using a minimum-bias trigger and overlay these MB events onto our MC events. In this way, we account for inefficiencies due to detector occupancy at high luminosity, for residual signals from preceding bunch crossings,

and for particles in the beam halo interacting with the detector. Also, although this effect is not luminosity-dependent, the minimum bias overlay is used to model electronic noise.

Each minimum-bias event corresponds to a particular value of the instantaneous luminosity. Overlay of MB events onto MC events therefore gives the MC sample a distribution of instantaneous luminosity, called the *luminosity profile*. The data used in our $ZH \rightarrow \mu\mu bb$ analysis also has a luminosity profile, but it may not agree with the profile of the MB events. To ensure proper modeling of luminosity-dependent effects, the luminosity profile is reweighted to agree with the profile in our data.

4.2.3 Triggering

In the absence of a direct trigger simulation, we derive trigger probabilities as functions of the kinematics of reconstructed objects, evaluate these functions for every reconstructed object in an event, and combine the resulting probabilities into a global trigger probability for the event. Application of these event probabilities serves to account for the possibility that MC-generated particles might not have passed the $D\emptyset$ trigger, though the parametrization of the trigger probabilities is a coarse approximation of the actual criteria in the trigger.

This technique is complicated by the fact that there are changes to the $D\emptyset$ triggers from time to time. To properly account for this, we keep track of the integrated luminosity associated with each trigger version. Trigger probabilities are derived for each version, and the event weight applied to MC is a weighted average of the probability associated with each version, as shown in equation 4.1. L_v and P_v are the integrated luminosity and the probability associated with trigger version v .

$$P = \frac{\sum_v L_v P_v}{\sum_v L_v} \quad (4.1)$$

In the $ZH \rightarrow \mu\mu bb$ analysis, as described in section 5.3, we use the probabilities derived for a single muon trigger suite [36] in our modeling of the inclusive trigger efficiency. To account for the geometry of the inner and outer trackers, these probabilities are functions of muon η and ϕ .

4.3 Corrections to Reconstructed Objects

Following the detector simulation and application of MB overlay, every simulated event passes through the same reconstruction software as is used for the data. However, reconstructed objects in simulated events tend to be identified at a higher rate than those in data, and the kinematics of identified objects tend to have better resolution in simulated events than in data. These discrepancies are due to shortcomings of the detector simulation; for example, the simulation does not include the effects of aging from radiation exposure, and some inactive material may not be present in the detector model used by GEANT. As a result, it is necessary to apply corrections to the reconstructed objects in simulated events.

4.3.1 Primary Vertex

Our MC generators assume that the interaction region is Gaussian along the z direction, whereas the z_{PV} distribution in data has significant non-Gaussian features. To bring the simulated and actual z_{PV} distributions into agreement, we directly reweight the simulated distribution. The reweighted distributions are overlaid with the data in figure 4.1. In general, one must be careful with direct reweighting, as it has the potential to hide interesting new physics. Fortunately, we do not expect any new physics to be distinguished by z_{PV} , as the location of the primary vertex should be unrelated to the hard-scatter process occurring in the associated collision.

4.3.2 Muons

The muon p_T resolution is significantly better in simulation than in data. To correct for this, we apply a random, Gaussian smear to the curvature of the inner track associated with each simulated muon, such that the width of the Z mass peak in $Z \rightarrow \mu\mu$ events is well-modeled.

Also, muon identification and track-matching efficiencies are higher in simulation than in data [22]. To obtain the muon ID efficiency, a *tag-and-probe* method is used in which the existence of an identified muon (the *tag*) implies the presence of a second muon (the *probe*), whether identified or not. The efficiency is just the fraction of events where the probe muon

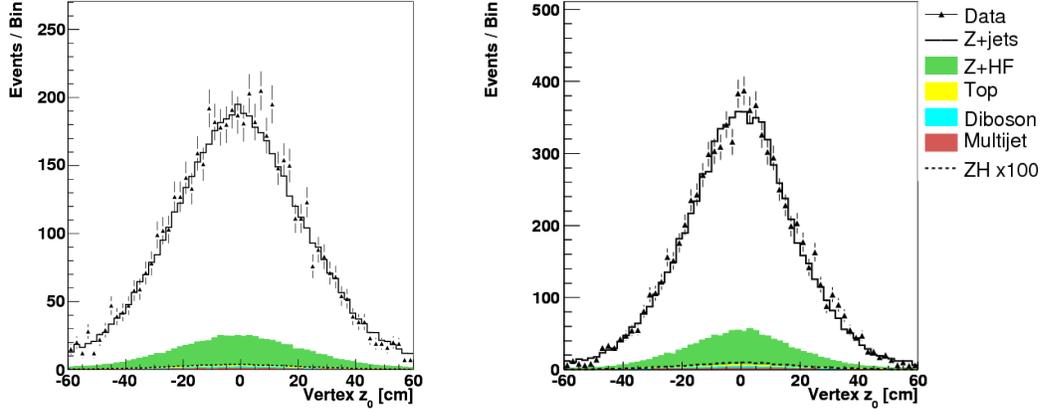


Figure 4.1: Distribution of the z coordinate of the primary vertex in data and simulation, after correcting for non-Gaussian features. To the left is RunIIa; to the right, RunIIb.

is identified, parametrized by the kinematics of the muon. To resolve the discrepancy in muon ID between simulation and data, we apply a scale factor for each identified muon in the simulation, where the scale factor is defined as the ID efficiency in data divided by the efficiency in simulation. Likewise, we derive and apply track-matching efficiency scale factors to account for more efficient tracking in the simulation.

4.3.3 Jets

Like muons, jets have better resolution in simulation than in data, and identification of simulated jets is also more efficient. Additionally, the JES-corrected energy of a simulated jet is somewhat shifted with respect to the equivalent jet in data. To correct for all of these effects, we use an algorithm known as JSSR, or jet smearing, shifting, and removal [37]. First, the energies of simulated jets are randomly smeared to recover the jet energy resolution as measured in data. Then, the jet energies are shifted to account for the difference in the jet energy scale between simulation and data. Finally, jet identification scale factors are applied not as weights, but by interpreting the scale factors (always < 1) as probabilities and randomly rejecting some identified jets.

Much as the simulated muon track-matching efficiency is too high, we find that vertex confirmation of jets is more efficient in simulation than in data. We parameterize the vertex-confirmation efficiencies as functions of jet p_T and $z_{PV} \times \text{sign}(\eta_{\text{jet}})$ and calculate scale factors as the ratio of efficiency in data to efficiency in simulation. Just as for jet ID, the vertex-confirmation scale factors are interpreted as probabilities and applied by randomly rejecting some vertex-confirmed jets.

4.3.4 b Tagging

Much as the vertex confirmation efficiencies are poorly modeled, the jet taggability requirement is substantially more efficient in simulation than in data. Our method of correcting the taggability efficiency, specific to the $ZH \rightarrow \ell\ell bb$ analyses, is described in section 5.6.1.

The b -tagging neural net output is also poorly modeled, with any cut resulting in exaggerated b -jet efficiency and underestimated light-jet misidentification in simulation. In the $ZH \rightarrow \mu\mu bb$ analysis, we choose to avoid this problem by using an indirect b -tagging method, similar to our indirect method of triggering. Tagging probabilities are derived for each flavor of jet as functions of the jet p_T and η [38]. These probabilities are known as *tag rate functions*, or TRF's. Event weights are determined by evaluating the appropriate TRF for each jet in each simulated event, as described in section 5.6.2, and these weights are applied to the simulated events in order to model b tagging.

Chapter 5

ANALYSIS

There are several standard-model Higgs production modes accessible at the Tevatron. By almost an order of magnitude, the most common mode is gluon fusion to a single, unaccompanied Higgs boson. However, low-mass Higgs bosons ($m_H < 135$ GeV) decay mainly to $b\bar{b}$, and the $b\bar{b}$ final state is dominated by background from various QCD interactions. To exploit the dominant decay mode of low-mass Higgs bosons (see section 1.3), we search for associated production of the Higgs with other particles whose presence serves to reduce the problematic QCD background. So-called *Higgsstrahlung* processes, in which a Higgs boson is produced in association with a W or Z boson, are particularly useful. Leptonic decays of the W or Z distinguish these final states from multijet production via QCD. The Tevatron, as a $\bar{p}p$ collider, is especially suited to WH and ZH production through valence quark-antiquark annihilation.

At $D\bar{O}$, several analysis teams search for WH and ZH production. Each team focuses on one out of three distinct final-state topologies: $WH \rightarrow \ell\nu bb$, $ZH \rightarrow \nu\nu bb$, and $ZH \rightarrow \ell\ell bb$. (The symbol ℓ denotes an electron or muon, but ν may be any flavor of neutrino.) WH is expected to have the highest production rate, and ZH production should occur roughly 0.6 times as often. The branching ratio of W decay to an electron or muon and a neutrino is 22%, while $BR(Z \rightarrow \nu\nu) = 20\%$ and $BR(Z \rightarrow \ell\ell) = 6.7\%$. Thus, for every ten $WH \rightarrow \ell\nu bb$ events collected by $D\bar{O}$, we expect approximately five $ZH \rightarrow \nu\nu bb$ events and two $ZH \rightarrow \ell\ell bb$ events.

Assuming each low-mass Higgs search achieves the same ratio of signal to background (S/B), the relative signal yields suggest that $ZH \rightarrow \ell\ell bb$ would contribute approximately 5% to overall Higgs sensitivity. Actually, the $\ell\ell bb$ final state contributes close to 10%; in other words, each $\ell\ell bb$ event is twice as sensitive to Higgs production as an analogous $\ell\nu bb$ or $\nu\nu bb$ event. This is due to several unique features of the $ZH \rightarrow \ell\ell bb$ signal. The presence

of two charged leptons in the final state suppresses the QCD multijet background, as it is quite difficult for two jets to be both misidentified as leptons. The absence of neutrinos makes this final state unambiguous: we can directly measure all four objects expected in every $\ell\ell bb$ event. This allows us to fully reconstruct the dilepton and dijet mass resonances that are expected in our signal but often absent from the backgrounds. Finally, the lack of ambiguity also enables the use of polar angles to distinguish between the s -channel signal and t -channel backgrounds.

There are many SM processes other than $\bar{p}p \rightarrow ZH$ that yield $\ell\ell bb$ events at $D\mathcal{O}$, and a few of these occur at substantially higher rates than ZH production. Zbb production is the largest of these physical backgrounds, with 100 Zbb events expected for every ZH ($m_H = 115$ GeV) event. Top quark pair ($t\bar{t}$) and ZZ production are much smaller than Zbb but significant compared to ZH , with expected S/B ratios of 25 and 5, respectively. We tackle these backgrounds by exploiting the $H \rightarrow b\bar{b}$ dijet mass resonance, the dominant t -channel nature of Zbb and ZZ , and the presence of high-energy neutrinos in $t\bar{t} \rightarrow \ell\nu\ell\nu b\bar{b}$ events.

Instrumental backgrounds, caused by misidentification of various objects and particles in the final state, must also be considered. Although the multijet background is dramatically reduced by the selection of two charged leptons, the total multijet cross section is so high that events with two jets misidentified as leptons remain more common than ZH events. The production of a Z boson in association with light-flavor jets (Zjj) occurs much more frequently than Zbb , so Zjj remains an important background despite impressive b -tagging rejection rates. Similarly, c jets are often misidentified as b jets, so Zcc and $ZW(\rightarrow cq)$ also have significant event yields after b tagging. Fortunately, these instrumental backgrounds may be further reduced in much the same way as the physical backgrounds, as they are all kinematically distinct from the ZH signal.

Using as much information as possible, we rely on a multivariate discriminant to separate the ZH signal from the backgrounds described above. With so many independent variables each lending substantial discrimination power, we can afford to select events based on very loose criteria, thus maximizing signal acceptance. At the final stage of the analysis, we use the full information available in the shape of the final discriminant, again avoiding any cuts

through which we could lose ZH events. With this strategy, we exploit the unique features of the $\ell\ell b\bar{b}$ final state to overcome a small expected $ZH \rightarrow \ell\ell b\bar{b}$ event yield and contribute significant sensitivity to the Higgs search at DØ.

The $ZH \rightarrow \mu\mu b\bar{b}$ analysis described here is the single most sensitive component of the search for $ZH \rightarrow \ell\ell b\bar{b}$ at DØ [39]. As part of a larger, coordinated effort, many aspects of this analysis were developed jointly and are shared with the other leptonic channels ($\mu\mu_{\text{TRK}}$, ee , and ee_{ICR}). Most notably, all $\ell\ell b\bar{b}$ channels share the same jet selection, MC normalization procedure, b tagging, kinematic fit, and random forest discriminant. Although channel-specific limits are useful for comparison, the leptonic channels are combined to set more stringent limits on ZH production.

5.1 Data

The data used in the $ZH \rightarrow \mu\mu bb$ analysis are divided into two epochs: RunIIa data were collected between April 2002 and February 2006, while RunIIb data were collected between June 2006 and December 2008. The integrated luminosity of the RunIIa data is 1.074 fb^{-1} , and RunIIb comprises 3.077 fb^{-1} . In total, 4.151 fb^{-1} of data are analyzed. To maximize signal acceptance, there is no explicit trigger requirement on the data used in this analysis; we select events accepted by any subset of DØ's inclusive trigger suite.

5.2 Pre-Tag Selection

This section details the selection criteria for muons and jets, before b tagging. This *pre-tag* selection provides a large sample of events with which we can test the accuracy of our background model and correct it if necessary. Since the subset of b -tagged events is quite small, the risk of bias from this procedure against a potential signal is minimal.

Before selecting muons or jets, we first require that the hard scatter occur in the bulk of the interaction region, where substantial SMT coverage allows for high-quality tracking and enables high-efficiency b tagging. The reconstructed primary vertex must be no farther than 60 cm from the center of the detector along the z axis.

A detector-based pseudorapidity η_{det} will be useful for the purposes of this chapter and in subsequent discussions. Because the PV can deviate from the center of the detector by as much as 60 cm, the physical pseudorapidity η (measured with respect to the PV) may differ substantially from η_{det} (measured with respect to the center of the detector), and η_{det} is more appropriate when describing detector effects.

5.2.1 $Z \rightarrow \mu^+ \mu^-$

Muons are reconstructed as described in section 3.3. Each event is required to contain at least two loose muons, with each muon matched to a central track of $p_{\text{T}} > 10$ GeV. Selected muons have $|\eta_{\text{det}}| < 2$, which corresponds to the boundary of the outer muon tracker. In addition, the inner track associated with each muon must satisfy $|z_{\text{DCA}} - z_{\text{PV}}| < 1$ cm.

Due to our method of modeling the inclusive trigger (see section 5.3), at least one muon is required to be within the coverage of the single muon trigger suite. The single muon triggers rely on the CFT, so one muon must satisfy $|\eta_{\text{det}}| < 1.5$. In addition, this central muon must have $p_{\text{T}} > 15$ GeV because the single muon trigger efficiencies are not well-understood for muons with $p_{\text{T}} < 15$ GeV. These requirements were found to have negligible impact on the signal acceptance, as the Z boson lends a significant boost to the muon p_{T} , and s -channel processes are especially likely to produce central particles in the final state.

To reconstruct the Z boson, we first build Z candidates from every pair of muons in the event. In each pair, the two muons are required to have opposite charge. To reject

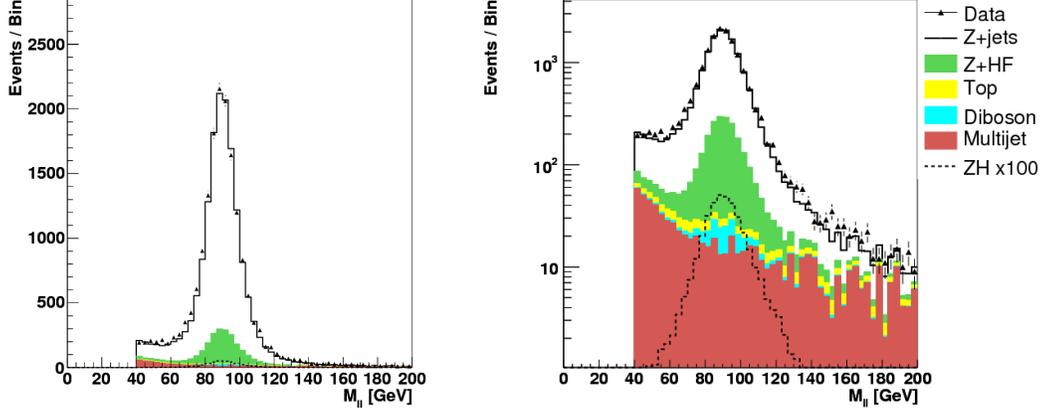


Figure 5.1: The dimuon invariant mass spectrum, in linear (left) and log (right) scales, requiring two jets and combining RunIIa and RunIIb.

cosmic-ray muons, we place a cut on a variable describing the collinearity of the muons. If the muons are collinear, then $\zeta = \pi - (\phi(\mu_1) - \phi(\mu_2)) + |\pi - (\theta(\mu_1) + \theta(\mu_2))|$ should be close to zero, so we require $\zeta > 0.05$. To reject muons produced in jet showering, we compute a scaled isolation variable for each muon as the sum of the track halo and calorimeter halo, divided by the muon p_T . For each Z candidate, the product of the scaled isolation variables of the two muons must be less than 0.03. This product-based approach to muon isolation is intended to maximize signal acceptance: one muon is allowed to be in close proximity to a jet, provided that the other muon is sufficiently isolated.

Events are required to have at least one Z candidate satisfying the criteria above. If an event has multiple candidates, then we choose the Z candidate closest to the mass of the Z boson, 91.2 GeV. Finally, to reduce non- Z backgrounds, the invariant mass of the chosen Z candidate must be between 60 and 150 GeV, although the region below 60 GeV is used during the background normalization procedure. The dimuon invariant mass spectrum is shown in figure 5.1.

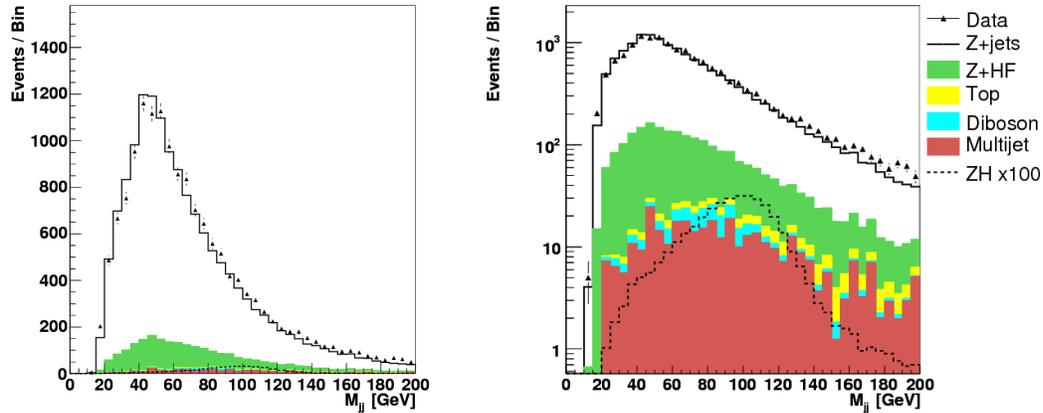


Figure 5.2: The invariant mass of the leading two jets, in linear (left) and log (right) scales, combining RunIIa and RunIIb. The expected Higgs resonance (multiplied by 100) can be seen as the wide peak in the dashed curve, close to 100 GeV.

5.2.2 Jets

For an energy deposit in the calorimeter to be considered as a jet in this analysis, the good jet criteria from section 3.4 must be satisfied, and the jet must have $p_T > 15$ GeV and $|\eta_{\text{det}}| < 2.5$. In RunIIb data, the jet must also be vertex-confirmed. Selected events contain at least two such jets, with the further criterion that at least one jet has $p_T > 20$ GeV. These cuts on jet p_T serve to reduce non-resonant backgrounds such as Zbb , and they have little impact on the jets expected from Higgs decay. When referring to the selected jets, the highest- p_T jet is called the *leading* jet, the second-highest- p_T jet is the *second* jet, and so on. Shown in figure 5.2, the dijet invariant mass is a powerful variable for separating the resonant Higgs signal from the non-resonant background.

Process	NLO/LO Scale
ZW	1.062
ZZ	1.030
$t\bar{t}$	1.434

Table 5.1: NLO corrections to MC background yields, excluding Z/γ^* +jets.

5.3 Monte Carlo

This section provides analysis-specific details of the simulated events (see chapter 4) used in the search for $ZH \rightarrow \mu\mu bb$. For more general aspects of our methods of simulation, see chapter 4.

With the exception of the multijet background, all significant background processes are simulated with Monte Carlo generators. We generate Z/γ^* +jets and $t\bar{t}$ using ALPGEN, while ZZ and ZW are generated with PYTHIA. Here, Z/γ^* +jets refers to all jet flavors and specifically to the decay $Z/\gamma^* \rightarrow \mu\mu$. Potential background from $Z/\gamma^* \rightarrow \tau\tau$ was measured to be at most 0.1%, and its non-resonant shape is similar to multijet, so we account for it in the normalization of the multijet background.

The signal process $ZH \rightarrow \mu\mu bb$ is likewise simulated using PYTHIA. We consider eleven signal hypotheses, varying the Higgs mass between 100 and 150 GeV in 5-GeV steps. Potential contributions to the signal from $H \rightarrow cc$ and $H \rightarrow \tau\tau$ decays are neglected, as they are expected to contribute at most 2% more signal at $m_H = 115$ GeV.

As the detector and the beam conditions changed significantly between RunIIa and RunIIb, it is necessary to use separate MC samples for the two data epochs. Although the MC events may be equivalent at generator level, the detector simulation is necessarily different for RunIIa and RunIIb, as is the minimum-bias overlay.

Although PYTHIA and ALPGEN are leading-order MC generators, we can correct these MC samples to account for several significant effects arising from higher-order Feynman diagrams. Most importantly, we scale the total event yield of each MC process to reflect a higher-order calculation of the cross section. All Z/γ^* +jets MC events are scaled up by a

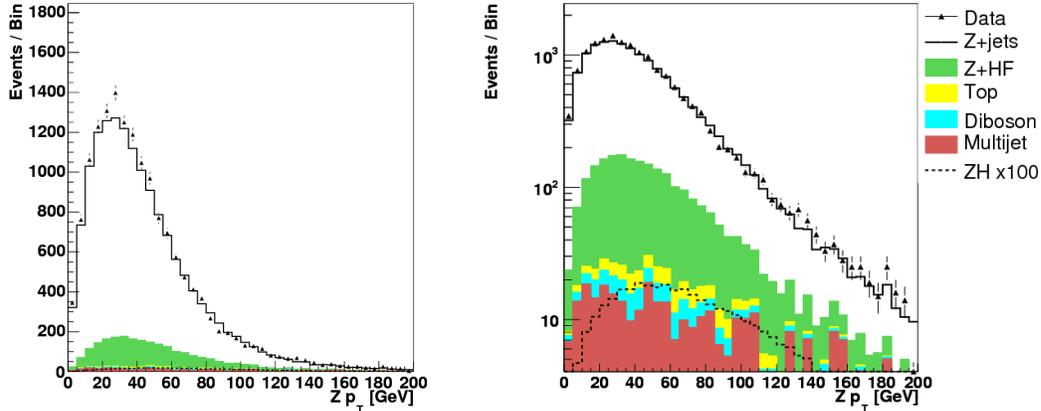


Figure 5.3: The p_T of the Z candidate after $Z p_T$ reweighting, in linear (left) and log (right) scales, requiring two jets and combining RunIIa and RunIIb.

factor of 1.3 to match NNLO calculations [40]. From comparison of NLO and LO heavy-flavor production, Zcc and Zbb MC events are further scaled by factors of 1.67 and 1.52, respectively, so our Zcc and Zbb event yields are roughly twice as large as those predicted by ALPGEN. Similarly, the other background MC yields are corrected according to NLO expectations, using the scales listed in table 5.1 [41, 42]. The $ZH \rightarrow \mu\mu bb$ signal cross sections are scaled to match NNLO calculations [43].

In addition to the event yields, we must correct some kinematic distributions of particular importance to the $ZH \rightarrow \mu\mu bb$ analysis. Both the $Z p_T$ distribution and the angle between jets are known to be poorly modeled in $Z/\gamma^* + \text{jets}$ events generated with ALPGEN. To correct the $Z p_T$, we apply event weights parametrized by the $Z p_T$ that were derived from measurements in $Z \rightarrow ee$ data [44]. There is some dependence of this reweighting on the number of jets, so we use the correction derived specifically for events with two or more jets. The $Z p_T$ distribution after reweighting is shown in figure 5.3.

To address the discrepancy in the angle between jets, a reweighting was derived using the combination of $\mu\mu$ and ee pre-tag events selected for the $ZH \rightarrow \ell\ell bb$ analysis [45]. The non- $Z/\gamma^* + \text{jets}$ backgrounds (including multijet) were subtracted from the data, and the resulting distributions were used as templates for the reweighting, which is assumed to be

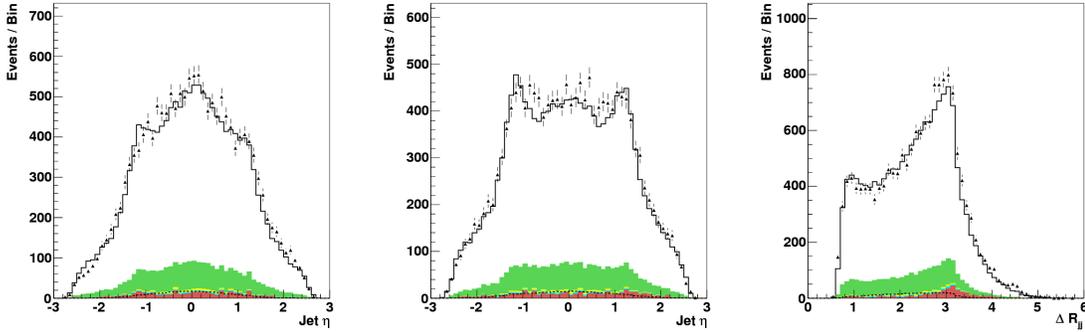


Figure 5.4: From left to right: leading jet η , second jet η , and ΔR between the two jets after jet angle reweighting, combining RunIIa and RunIIb.

a separable function of the leading jet η , second jet η , and ΔR between the jets. The η corrections are parametrized as symmetric fourth-order polynomials, and the ΔR correction is a third-order polynomial. Figure 5.4 shows the pseudorapidities of and ΔR between the leading two jets after applying the jet angle reweighting.

As the final correction to the pre-tag MC, we must account for imperfect acceptance of muons by the inclusive trigger suite. To do this, we first apply well-understood single-muon trigger efficiencies to the MC [36]. Then, a data-based correction serves to boost the MC yield in the kinematic regions where the inclusive trigger suite is substantially more efficient than the single-muon triggers.

To derive the correction from single-muon to inclusive triggering, we require that the leading central muon is matched to a single-muon trigger object in our data, and compare the resulting muon kinematics with those of the inclusively triggered data. The correction is assumed to be a separable function of the jet multiplicity (N_{jets}), the triggered muon η_{det} , and the untriggered muon η_{det} . The dependence on N_{jets} is used to account for the overall gain in trigger efficiency, as the inclusive efficiency is observed to increase with the number of jets because additional objects provide more triggering opportunities. The remaining two components of the trigger correction, dependent on triggered and untriggered muon η_{det} , are measured in the high-statistics sample of events with zero jets. As a cross-check,

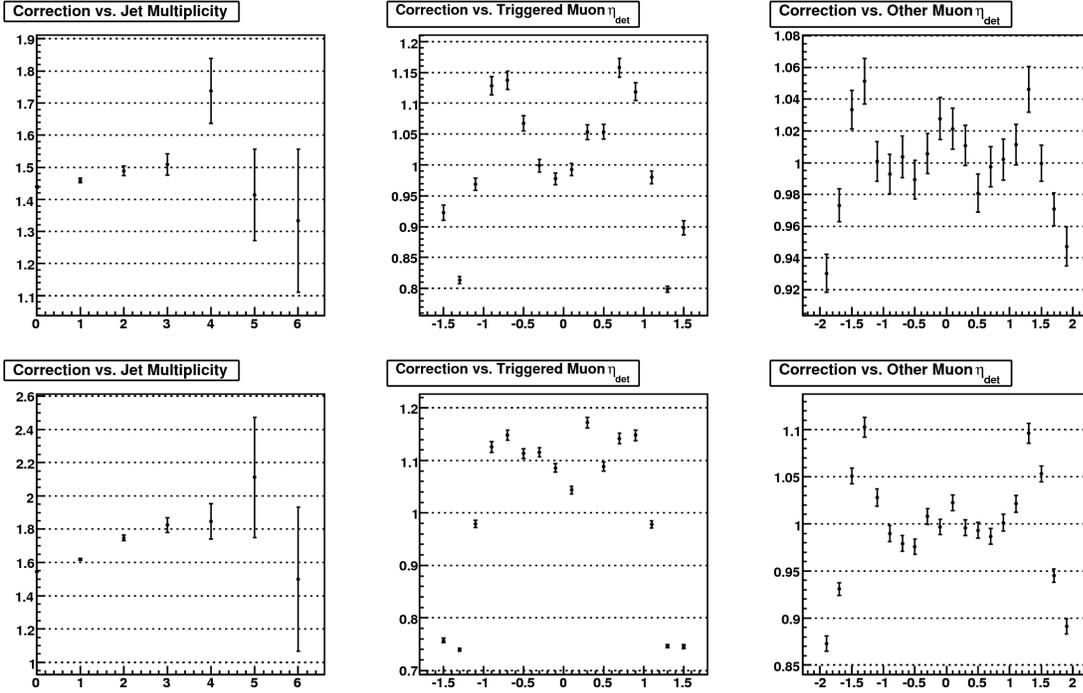


Figure 5.5: The components of the single-muon to inclusive trigger correction in RunIIa (top) and RunIIb (bottom). From left to right: the dependence on N_{jets} , triggered muon η_{det} , and untriggered muon η_{det} .

these components were also derived using events with at least two jets, and the resulting corrections were found to be consistent with those derived using zero-jet events.

Instead of fitting the components of the trigger correction with smooth functions, we use the same form as the single-muon trigger efficiencies, which are discrete with bins of width 0.2. The components of the RunIIa and RunIIb trigger corrections are shown in figure 5.5.

The single muon trigger efficiencies are likewise applied to the leading central muon in every MC event, and the data-based correction from single-muon to inclusive triggering is evaluated and applied as an additional event weight. Using this approach, we observe a clear improvement in the modeling of our muon η_{det} distributions, and we account for the fact that the inclusive trigger suite rejects a small but significant number of dimuon events that would otherwise pass our selection criteria. The muon η_{det} distributions after applying

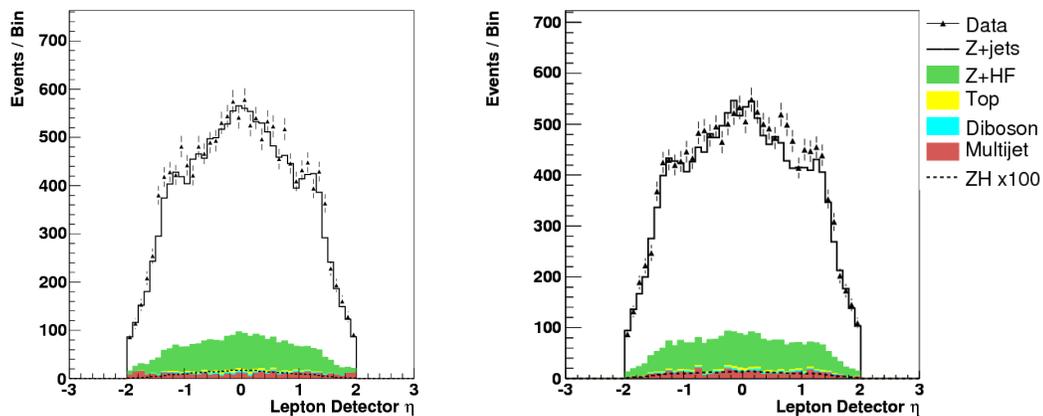


Figure 5.6: The leading muon η_{det} (left) and second muon η_{det} (right) after the trigger correction, requiring two jets and combining RunIIa and RunIIb.

the trigger correction are shown in figure 5.6.

5.4 Multijet Sample

Jets are very rarely misidentified as muons, so to properly simulate the instrumental multijet background would require an unfeasibly large multijet MC sample. Also, it is likely that the detector simulation would not be adequate to model properly the misidentification of jets as muons. Instead, we obtain a representative sample from the data. Our multijet selection uses the same criteria as given in section 5.2, with the exception of the isolation requirement for Z candidates. Every Z candidate in the multijet sample must fail the isolation cut and pass all other selection criteria.

While this data-based multijet background model is preferable to any MC-based solutions, the model nonetheless has relatively little statistical significance. Large statistical fluctuations are apparent in many kinematic distributions, for example the dijet mass distribution shown in figure 5.2. This is a weak point in the analysis, but the small size of the multijet background mitigates any negative impact on our results, and our limit-setting technique accounts for any statistical errors in our background models.

It should be noted that the “multijet” sample includes a small number of events containing single muons not attributed to jets, such as W +jets events. Because these events are kinematically similar to actual multijet events (with respect to the variables used in the $ZH \rightarrow \mu\mu bb$ analysis), we do not distinguish between the two. Thus, “multijet” subsequently refers to any event in which at least one jet is misidentified as a muon.

To complete the pre-tag multijet background model, the yield of the multijet sample must be scaled to match the number of multijet events expected to pass all selection requirements. This scale factor is determined as part of the background normalization procedure, described in section 5.5.

5.5 Background Normalization

In order to improve the accuracy of the background model, scale factors adjusting the pre-tag contribution of each background are fit to the observed pre-tag data. These scale factors are found to be consistent with the magnitudes of the relevant theoretical and experimental uncertainties, so effectively we are using our pre-tag data to further constrain these uncertainties. The fit adjusts the normalization of dilepton invariant mass ($m_{\ell\ell}$) histograms from each background sample to obtain the best agreement with the pre-tag data. This fit includes all $ZH \rightarrow \ell\ell bb$ leptonic channels in order to take into account correlations from one leptonic channel to another. For example, a correction to the $Z/\gamma^*+2\text{jet}$ cross section should be common across all leptonic channels.

The multijet samples must also be normalized to match the number of multijet events which actually satisfy the pre-tag selection criteria. The low $m_{\ell\ell}$ region is dominated by multijet events, so it serves to constrain the multijet normalization. As the cuts differ for each leptonic channel and the efficiency may depend on jet multiplicity, a separate multijet scale factor (α) is fit for each channel and jet-multiplicity bin.

The peak in the $m_{\ell\ell}$ distribution from Z production is used to constrain the product of lepton efficiency, integrated luminosity, and Z/γ^* production cross section in each jet bin. As the inclusive Z/γ^* cross section, dominated by the zero-jet bin, is known to much better accuracy than the $Z/\gamma^*+2\text{jet}$ cross section, the zero-jet bin is used to constrain the product of lepton efficiency and luminosity, which applies to every MC sample, while the two-jet bin is used to constrain the $Z/\gamma^*+2\text{jet}$ cross section.

The normalization fit is performed simultaneously for all jet-multiplicity bins and leptonic channels. The contribution of each background component is scaled by appropriate fit parameters: an N_{jets} -independent lepton efficiency correction multiplies all MC background, and an N_{jets} -dependent correction to the Z/γ^* cross section multiplies all Z/γ^* samples. Using dilepton invariant mass histograms with mass-bin index m , for each lepton channel i and jet multiplicity bin j , tallied separately for data (D_m^{ij}), $Z/\gamma^*+2\text{jets}$ MC (Z_m^{ij}), non- Z/γ^* MC (O_m^{ij}), and multijet (Q_m^{ij}), the combined fit minimizes the χ^2 given in equation 5.1.

$$\chi^2 = \sum_i \sum_j \sum_m \frac{1}{D_m^{ij}} \left(D_m^{ij} - \alpha^{ij} Q_m^{ij} - k_L k_\epsilon^i \left(k_Z^j Z_m^{ij} + O_m^{ij} \right) \right)^2 \quad (5.1)$$

The parameters obtained from the fit account for the luminosity (k_L), the lepton efficiency for each leptonic channel (k_ϵ^i), the Z/γ^* cross section for each jet-multiplicity bin (k_Z^j), and the multijet normalization for each channel and jet-multiplicity bin (α^{ij}).

There are too many parameters to determine directly from data, so some external constraints are necessary. The most important constraint is on the $Z/\gamma^*+0\text{jet}$ cross-section. For determining the nominal normalization parameters we fix $k_Z^0 = 1$, which is equivalent to assuming that the inclusive Z/γ^* cross-section is known exactly. We reexamine this external constraint when assessing systematics, as described in section 5.10. Also, the luminosity factor k_L appears everywhere alongside the lepton efficiency k_ϵ^i . As we have no need to separate these two effects, we merely fix $k_L = 1$, absorbing any luminosity error into the factor k_ϵ^i . Were we provided no luminosity estimate whatsoever, the analysis could proceed unchanged: we would lose only the cross-check that the factor k_ϵ^0 is near one.

The multijet background is most significant in the region of low dilepton invariant mass, $m_{\ell\ell} < 60$ GeV. In order to achieve a normalization optimized for the tighter signal region $60 < m_{\ell\ell} < 150$ GeV, the combined fit is applied in two stages. First, all parameters are free and fit in the extended range $40 < m_{\ell\ell} < 150$ GeV. Then, the multijet factors α are fixed and the fit is repeated in the tighter range $60 < m_{\ell\ell} < 150$ GeV.

The normalization results are not consistent between RunIIa and RunIIb. For this reason, we fit RunIIa and RunIIb separately. To account for these discrepancies, additional systematic uncertainties are applied, as discussed in section 5.10. The fit results for the $\mu\mu$ and ee channels are presented in table 5.2. For both normalization fits, $\chi^2/n_{\text{DF}} \approx 2$.

RunIIa		
	$\mu\mu$	ee
k_ϵ^i	0.951 ± 0.003	1.09 ± 0.01
α^{i0}	1.00 ± 0.07	1.07 ± 0.05
α^{i1}	0.25 ± 0.03	1.01 ± 0.12
α^{i2}	0.014 ± 0.004	0.47 ± 0.07

RunIIb		
	$\mu\mu$	ee
k_ϵ^i	0.894 ± 0.002	0.96 ± 0.01
α^{i0}	1.89 ± 0.10	0.67 ± 0.01
α^{i1}	0.20 ± 0.02	0.61 ± 0.01
α^{i2}	0.011 ± 0.003	0.55 ± 0.01

	RunIIa	RunIIb
k_Z^0	1	1
k_Z^1	0.963 ± 0.007	0.917 ± 0.005
k_Z^2	1.09 ± 0.01	1.00 ± 0.01

Table 5.2: Combined normalization fit results and statistical uncertainties.

5.6 *b* Jet Selection

In order to better isolate the $H \rightarrow bb$ signal from light-jet background processes, we require b -tagged jets in each event, using the neural network b tagger described in section 3.5. First, we select events with at least two loose b -tagged jets, where “loose” means NN output greater than 0.2. I will refer to this b -tagging requirement as our “2L” selection. The sample of 2L-selected events provides the bulk of our sensitivity to a Higgs signal, but we also use a second, orthogonal tagging requirement to maximize signal acceptance. From the sample of events failing the 2L requirement, we select events with exactly one tight b -tagged jet, where “tight” means NN output greater than 0.85. I’ll refer to this second b -tagged sample as our “1T” selection.

To optimize the discrimination power of dijet variables, as the Higgs boson is expected to decay to two b jets, we use the available b -tagging information when selecting the two jets from which we construct the dijet system. When tagging information does not unambiguously select two jets, the p_T of each jet is used to resolve the ambiguity. If an event has more than two tagged jets, we form the dijet system from the two highest- p_T tagged jets. If an event has only one tagged jet, we take the highest- p_T untagged jet as the second b jet.

The remainder of this section is devoted to our handling of b tagging in our background and signal models. In subsection 5.6.1, I discuss how we correct for mismodeling of the taggability efficiency in our simulated events. Subsection 5.6.2 explains how we indirectly tag simulated jets to retain statistical significance in our MC samples. Finally, subsection 5.6.3 describes the method of indirect b tagging used to retain statistics in the modeling of our post-tag multijet background.

5.6.1 *Taggability*

For a jet to be tagged as a b -jet, it must first pass a taggability requirement, as described in section 3.5. It should be noted that we only require taggability for jets that are b -tagged. Because tracking is more efficient in simulation, the taggability efficiency is higher in simulated events than in data, so we must apply scale factors to our MC events to properly model the impact of the taggability requirement. Also, these scale factors tend to be highly

sample-dependent, so we derive them using our own selected pre-tag data and MC samples in the $\mu\mu$ and ee leptonic channels. To alleviate p_T bias, both the leading and second jets in the event are used to derive these analysis-specific scale factors.

Our treatment of the taggability scale factors differs somewhat between RunIIa and RunIIb, since we don't require vertex confirmation in RunIIa. We first discuss the simpler case of the RunIIb correction. Figure 5.7 shows the taggability rates in RunIIb data and MC, expressed as functions of jet $\eta_s = \eta \times \text{sign}(z_{PV})$ and of jet p_T . Also shown are the ratios of data to MC, which define the taggability scale factors. There is no significant dependence on p_T , so the scale factor is evaluated as a function of η_s only. We parametrize this dependence using a fourth order polynomial. An ‘‘anti-taggability’’ scale factor for the inverted taggability requirement is also determined, as shown in figure 5.8. Anti-taggability scale factors are required to account for mismodeling of the rate for jets to fail the taggability criteria.

As shown in figure 5.9, the taggability scale factor for RunIIa does vary with jet p_T . To model this dependence, taggability and anti-taggability scale factors are computed independently for jets with $p_T < 30$ GeV and $p_T > 30$ GeV. The scale factors in these two p_T regions, expressed as functions of η_s , are shown in figures 5.10 and 5.11.

To evaluate the impact on our simulated event selection, we calculate the taggability scale factor for each taggable jet and the anti-taggability scale factor for each jet that is not taggable. We then apply the product of these scale factors over all jets in each event as an event weight to correct the simulated post-tag event yields.

5.6.2 Event Weights from TRF's

Limited statistics in our Monte Carlo samples necessitate a probabilistic treatment of b tagging in simulated events. Instead of directly tagging the jets, we allow for the possibility that any jet in the event could be tagged. Tag rate functions, as described in section 4.3.4, give the probability that a given jet is tagged as a function of the jet flavor, p_T , and η . (A dependence on taggability is also implied, as jets that are not taggable have zero probability to be tagged.) Because we use dijet variables that depend on which jets are

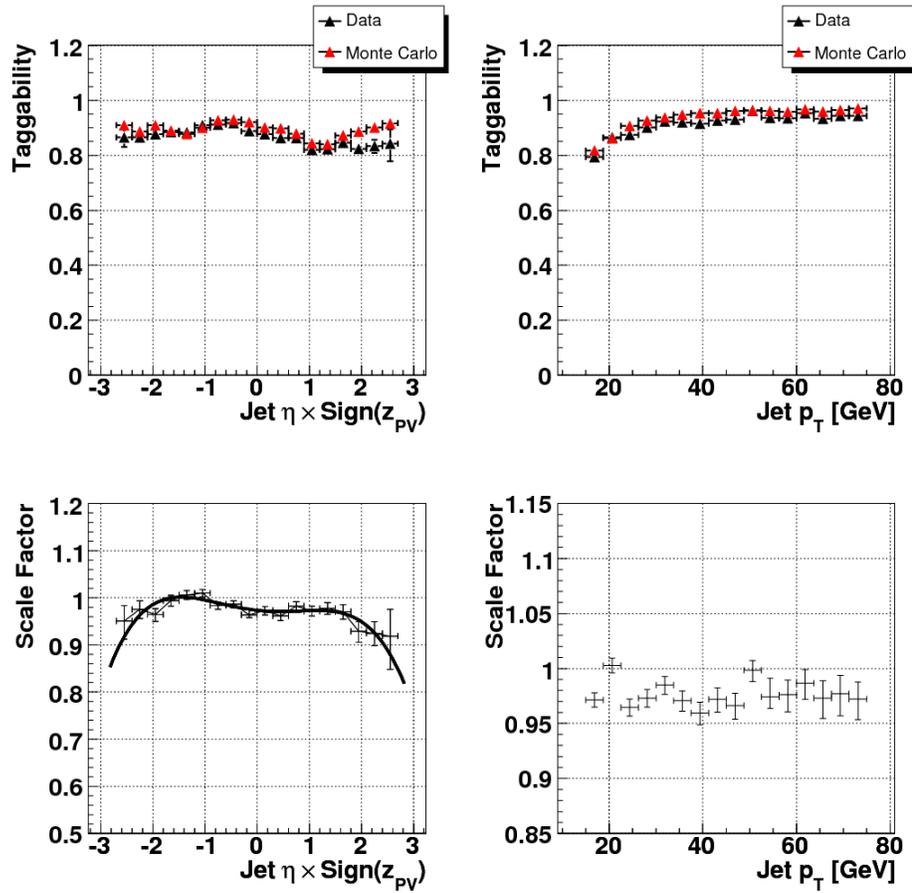


Figure 5.7: Taggability in RunIIb. The top plots show the efficiencies for data (black) and MC (red) jets versus jet η_s (left) and p_T (right). The bottom plots show the ratios of data efficiencies to MC efficiencies. The ratio, expressed as a function of η_s , is parametrized as a fourth order polynomial and used as a scale factor. No dependence on jet p_T is observed.

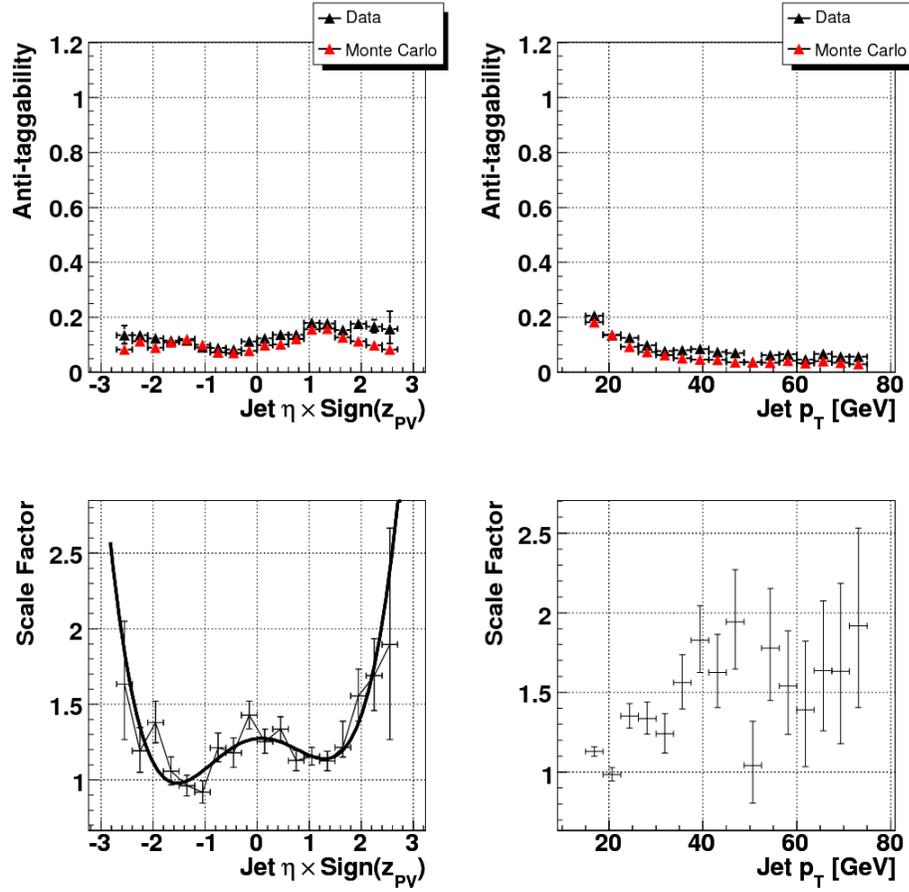


Figure 5.8: Anti-tagability in RunIIb. The top plots show the rates in data (black) and MC (red) jets versus jet η_s (left) and p_T (right). The bottom plots show the ratios of data efficiencies to MC efficiencies. The ratio, expressed as a function of η_s , is parametrized as a fourth order polynomial and used as a scale factor. No significant dependence on jet p_T is observed.

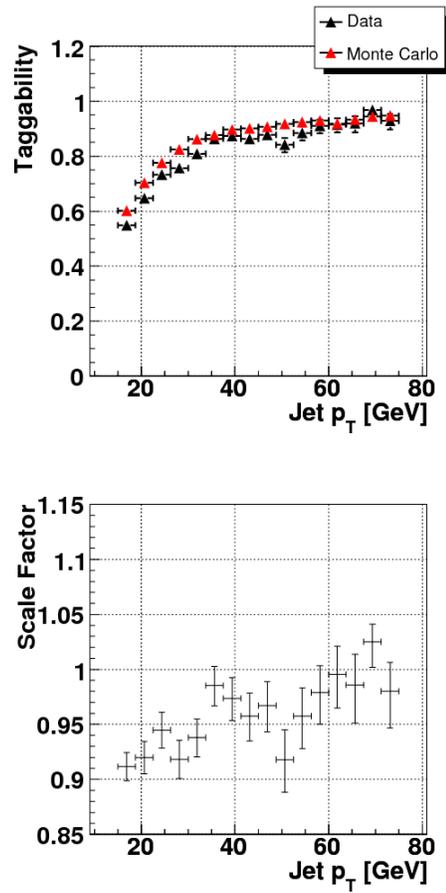


Figure 5.9: Taggability in RunIIa. The top plot shows the efficiency for data (black) and MC (red) jets versus jet p_T . The bottom plot shows the ratio of data efficiency to the MC efficiency.

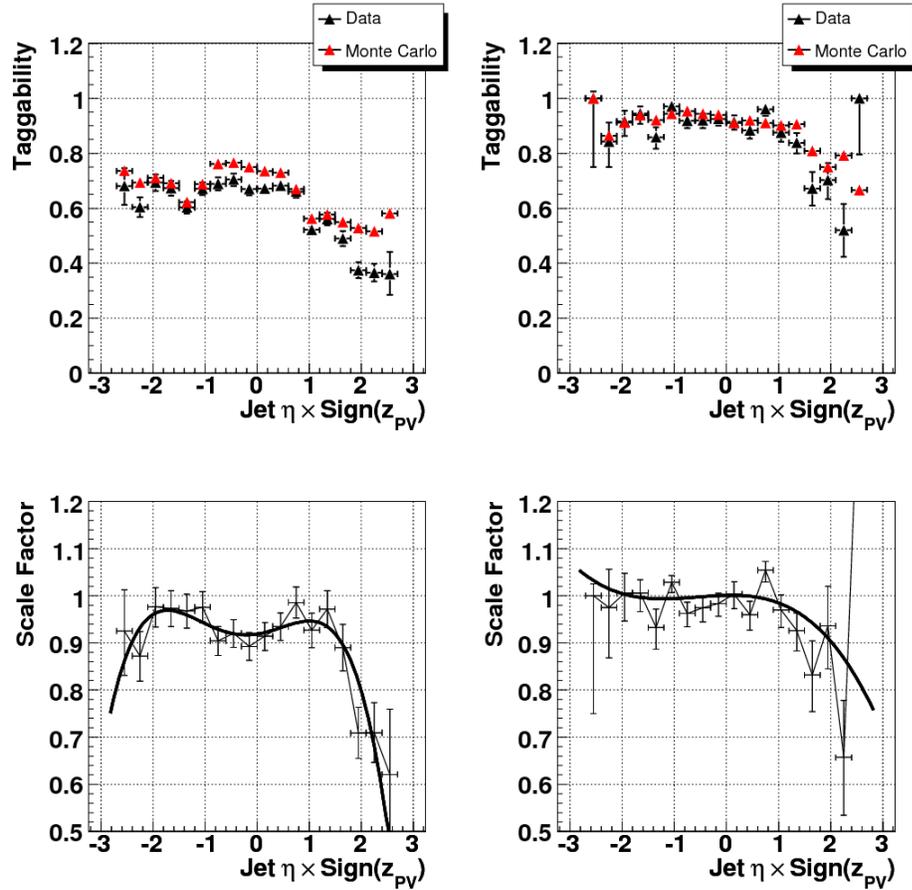


Figure 5.10: Taggability in RunIIa. The top plots show the efficiencies for data (black) and MC (red) jets versus jet η_s . The bottom plots show the ratios of data efficiencies to MC efficiencies. On the left, we require jet $p_T < 30$ GeV. On the right, we require jet $p_T > 30$ GeV.

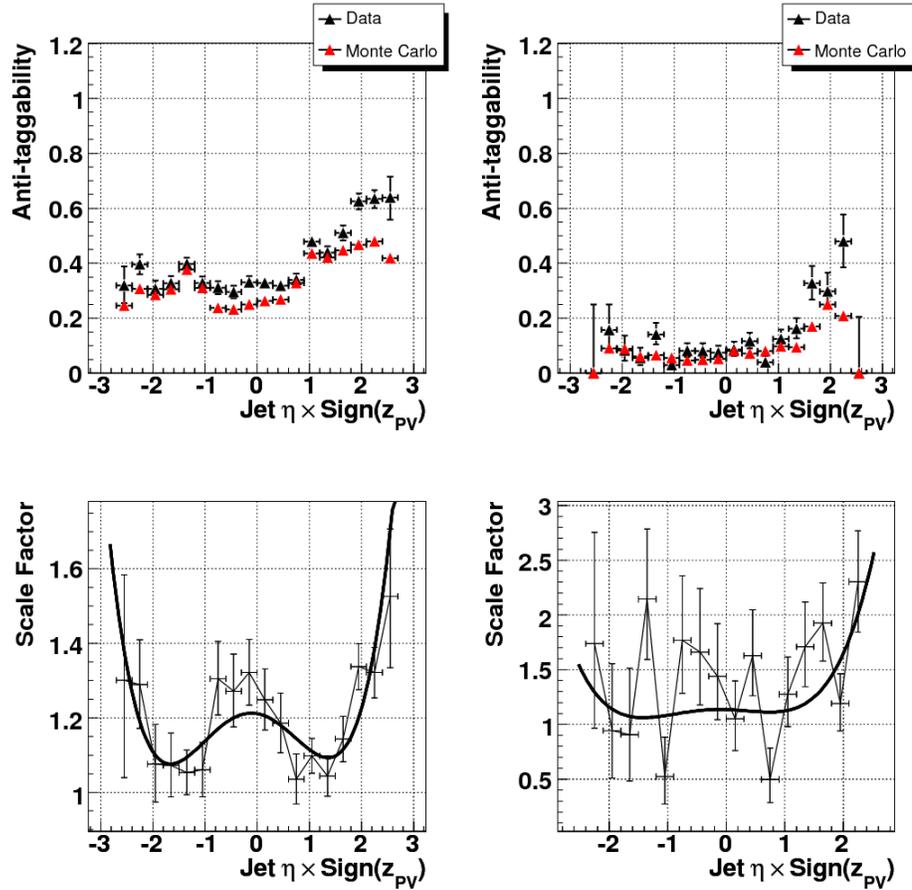


Figure 5.11: Anti-tagability in RunIIa. The top plots show the efficiencies for data (black) and MC (red) jets versus jet η_s . The bottom plots show the ratios of data efficiencies to MC efficiencies. On the left, we require jet $p_T < 30$ GeV. On the right, we require jet $p_T > 30$ GeV.

tagged, it is not sufficient to know the total probability that each event satisfies the 2L or 1T tagging criteria. Instead, we must consider all possible di- b -jet combinations and assign appropriate TRF-based probabilities to each jet pair. Thus, one MC event may provide several alternative values (with associated weights) of any dijet variable, corresponding to all potentially tagged jet pairs in the event. In the remainder of this subsection, I describe how we use TRF's to obtain the 2L and 1T weights associated with each jet pair.

To select a jet pair as 2L-tagged, the pair must comprise the two highest- p_T tagged jets. Let L_i be the probability that jet i has a loose tag. For a given jet combination $1 \leq i < j \leq N_{\text{jets}}$, where i and j are p_T -ordered jet indices, the probability that each jet pair is 2L-tagged is given in table 5.3, for events with up to four jets. Equation 5.2 gives the general formula for these probabilities, which is applicable to events with any number of jets.

$$p(i, j) = \frac{L_i L_j}{(1 - L_i)(1 - L_j)} \prod_{k=1}^j (1 - L_k). \quad (5.2)$$

To select a jet pair as 1T-tagged, one of the jets must have a tight tag, and no other jet in the event can have even a loose tag; otherwise, the event would be 2L-tagged. To complete the 1T pair, the highest- p_T non-tagged jet is selected. It is therefore impossible not to include the highest- p_T jet in a 1T pair, as it is always either tagged or the highest- p_T non-tagged jet. Define L_i as above, and let T_i be the probability that jet i has a tight tag. For a given jet combination $1 \leq i < j \leq N_{\text{jets}}$, the probability that each jet pair is 1T-tagged is given in table 5.4, for events with up to four jets. Equation 5.3 gives the general formula for these probabilities (δ_{ij} is the Kronecker delta function), which is applicable to events with any number of jets.

$$p(i, j) = \delta_{i1} \left(\frac{T_j}{1 - L_j} + \delta_{j2} \cdot \frac{T_1}{1 - L_1} \right) \prod_{k=1}^{N_{\text{jets}}} (1 - L_k) \quad (5.3)$$

5.6.3 Tagging in the Multijet Sample

Direct b tagging in the multijet sample results in a very low-statistics model of our post-tag multijet background, with statistical fluctuations obscuring much of the multijet shape in

N_{jets}	Jet Pair (ij)	2L Probability
2	12	L_1L_2
3	12	L_1L_2
	13	$L_1(1 - L_2)L_3$
	23	$(1 - L_1)L_2L_3$
4	12	L_1L_2
	13	$L_1(1 - L_2)L_3$
	14	$L_1(1 - L_2)(1 - L_3)L_4$
	23	$(1 - L_1)L_2L_3$
	24	$(1 - L_1)L_2(1 - L_3)L_4$
	34	$(1 - L_1)(1 - L_2)L_3L_4$

Table 5.3: The probability for a jet pair to be 2L-tagged, in terms of the loose TRF's L_i , in events with up to four jets.

N_{jets}	Jet Pair (ij)	1T Probability
2	12	$T_1(1 - L_2) + (1 - L_1)T_2$
3	12	$T_1(1 - L_2)(1 - L_3) + (1 - L_1)T_2(1 - L_3)$
	13	$(1 - L_1)(1 - L_2)T_3$
	23	0
4	12	$T_1(1 - L_2)(1 - L_3)(1 - L_4) + (1 - L_1)T_2(1 - L_3)(1 - L_4)$
	13	$(1 - L_1)(1 - L_2)T_3(1 - L_4)$
	14	$(1 - L_1)(1 - L_2)(1 - L_3)T_4$
	23	0
	24	0
	34	0

Table 5.4: The probability for a jet pair to be 1T-tagged, in terms of the loose (L_i) and tight (T_i) TRF's, in events with up to four jets.

Epoch	1T Scale	2L Scale
RunIIa	9.5 ± 0.4	11.7 ± 0.5
RunIIb	10.8 ± 0.3	10.8 ± 0.3

Table 5.5: The scale factors obtained by comparing direct b tagging with the application of light-jet TRF's in the multijet sample. They are significantly larger than one because the multijet sample has non-negligible heavy-flavor jet content.

many important variables. To avoid these fluctuations and preserve a statistically significant background shape, we employ an indirect method of tagging in the multijet sample. The true flavors of jets are unknown in the data-based multijet sample, but we do know that most jets have light flavor. Thus, we apply the light-jet TRF's to all jets in the multijet sample, calculating 2L and 1T event weights as described in the preceding section. These weights lead to underestimated post-tag multijet event yields, as there is in fact some heavy-flavor jet content in the multijet background. To correct for this, we apply scale factors to bring the post-tag multijet yields into agreement with the yields obtained through direct tagging in the multijet sample. Table 5.5 shows the scale factors used for this purpose, along with their statistical uncertainties.

5.7 Event Yields

Table 5.6 shows event yields in the data, the total background model, each background separately, and the ZH signal model (assuming $m_H = 115$ GeV). The first column of yields gives the number of events after selection of two muons but before the dimuon mass cut and before any jet selection. The second column gives event yields after the full selection of muons and jets but before b tagging. Finally, the third and fourth columns give event yields after the 1T and 2L b -tagging requirements, respectively. RunIIa and RunIIb yields are presented separately, since different MC samples were used for the different data epochs.

In table 5.7, the benefits of the various stages of selection are clear. In particular, S/B increases by factors of 10 and 20 upon applying, respectively, the 1T and 2L b -tagging criteria. The significance S/\sqrt{B} is a rudimentary measure of our sensitivity to a signal; in a simple counting experiment, neglecting systematic uncertainties, we would expect to see evidence of the signal if $S/\sqrt{B} > 3$. Also shown in table 5.7 are the substantial improvements to signal significance that result from b tagging.

RunIIa				
	$N_{\text{jets}} \geq 0$	$N_{\text{jets}} \geq 2$ (pre-tag)	1 tight tag	2 loose tags
data	100153	5497	112	77
background	97348 ± 151	5383 ± 15	120.89 ± 0.54	67.11 ± 0.38
Multijet	975 ± 24	28.09 ± 0.63	3.258 ± 0.076	3.036 ± 0.061
Zjj	93363 ± 149	4643 ± 14	31.097 ± 0.083	14.238 ± 0.037
$Zb\bar{b}$	768.6 ± 2.9	197.1 ± 1.6	50.32 ± 0.48	25.59 ± 0.33
$Zc\bar{c}$	2032.1 ± 6.8	432.5 ± 3.2	25.67 ± 0.23	11.72 ± 0.13
ZZ	43.75 ± 0.38	24.75 ± 0.29	1.962 ± 0.049	1.939 ± 0.060
ZW	53.72 ± 0.65	26.81 ± 0.47	0.974 ± 0.028	0.374 ± 0.015
$t\bar{t}$	48.38 ± 0.24	27.28 ± 0.18	7.461 ± 0.052	10.162 ± 0.075
$ZH(115)$	1.3059 ± 0.0074	1.0630 ± 0.0067	0.2907 ± 0.0020	0.3705 ± 0.0027
RunIIb				
	$N_{\text{jets}} \geq 0$	$N_{\text{jets}} \geq 2$ (pre-tag)	1 tight tag	2 loose tags
data	198567	8286	223	184
background	197251 ± 271	8161 ± 38	229.3 ± 1.2	160.20 ± 0.93
Multijet	6195 ± 101	43.72 ± 0.73	5.90 ± 0.11	6.380 ± 0.096
Zjj	185308 ± 251	6876 ± 37	50.98 ± 0.28	40.96 ± 0.21
$Zb\bar{b}$	1481.8 ± 6.2	341.0 ± 3.2	94.4 ± 1.0	57.02 ± 0.80
$Zc\bar{c}$	3830 ± 14	744.8 ± 6.5	56.10 ± 0.58	29.69 ± 0.35
ZZ	90.48 ± 0.92	44.52 ± 0.68	3.71 ± 0.12	3.89 ± 0.14
ZW	110.9 ± 1.4	47.88 ± 0.98	2.042 ± 0.064	0.971 ± 0.037
$t\bar{t}$	107.85 ± 0.58	56.76 ± 0.42	15.95 ± 0.12	21.16 ± 0.17
$ZH(115)$	2.986 ± 0.022	2.287 ± 0.019	0.6402 ± 0.0057	0.8164 ± 0.0077

Table 5.6: Event yields before and after b tagging. Dimuon selection is implied, though the dimuon invariant mass cut is omitted for the $N_{\text{jets}} \geq 0$ yields.

	$N_{\text{jets}} \geq 0$	$N_{\text{jets}} \geq 2$ (pre-tag)	1 tight tag	2 loose tags
S/B	1.5×10^{-5}	2.5×10^{-4}	2.7×10^{-3}	5.2×10^{-3}
S/\sqrt{B}	7.9×10^{-3}	2.9×10^{-2}	5.0×10^{-2}	7.9×10^{-2}

Table 5.7: The expected ratio of signal to background and expected significance S/\sqrt{B} before and after b tagging, using the sum of event yields from RunIIa and RunIIb.

5.8 Kinematic Fit

The DØ detector measures lepton energy more precisely than jet energy. Furthermore, the p_T of the $\mu\mu bb$ system is expected to be small in any ZH signal events and in most background events. As a result, a significant transverse momentum imbalance is usually due to poor measurement of the jet energy. To make optimal use of the available kinematic information, the energies and angles of the two muons from the Z candidate and the two jets associated with the Higgs are fit within their resolutions to new values [46], minimizing a multidimensional χ^2 function subject to the following constraints:

- $m_{\mu\mu} = m_Z$, approximated as a Gaussian with width $\Gamma_Z/2.35 \approx 1$ GeV;
- $\Sigma p_x = 0$, allowing for a $\mu\mu bb$ system boost along the x axis with a Gaussian width of 7 GeV; and
- $\Sigma p_y = 0$, allowing for a $\mu\mu bb$ system boost along the y axis with a Gaussian width of 7 GeV.

The width the $\mu\mu bb$ p_T distribution was determined from the ZH MC, as were the jet and muon angular resolutions. When there are more than two jets in the event, we vary only the two jets used to construct the dijet invariant mass. Figure 5.12 shows the b -tagged dijet invariant mass distributions before and after the kinematic fit. The dijet mass distribution of the background samples is broadened, but shifted towards lower values. Both the signal resolution and the signal to background ratio are improved. As a result of the kinematic fit, signal significance improves by 8% in the vicinity of the Higgs resonance.

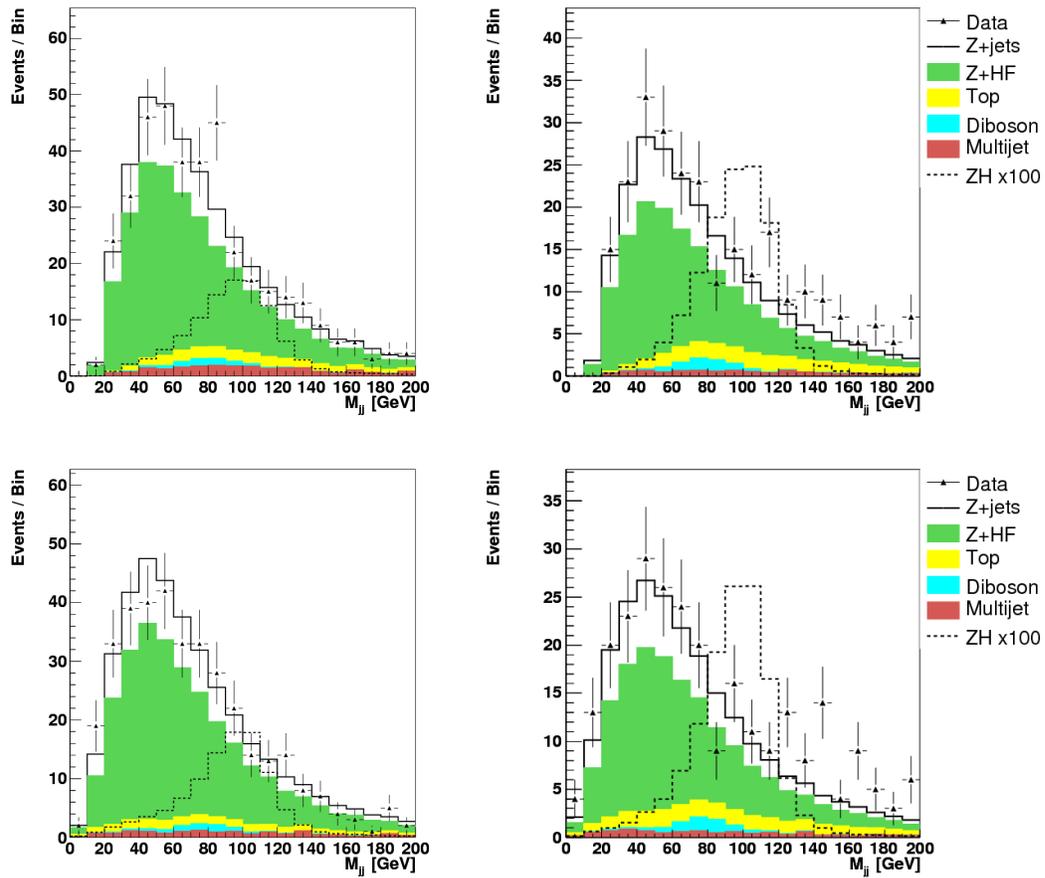


Figure 5.12: Dijet invariant mass, requiring 1 tight tag (left) and 2 loose tags (right), before (top) and after (bottom) applying the kinematic fit.

5.9 Multivariate Classifier

Given a sample of simulated signal and background, in which each event contains several kinematic variables, and the signal and background have somewhat different shapes in these variables, our task is to construct a function that takes these variables as input and provides an output that optimally separates the signal from the background. Such a function is commonly known as a multivariate *classifier*, and the process of constructing the classifier is generally called *training*. A typical classifier has an output ranging from zero to one. With ideal separation power, a cut at 0.5 would result in two samples: one of purely signal, the other of purely background.

With any type of multivariate classifier, because the signal and background samples used for training are discrete and finite, one must be cautious to minimize sensitivity to statistical fluctuations in the training sample. The phenomenon in which a classifier becomes trained to respond to statistical fluctuations is known as *overtraining*, and it generally leads to suboptimal classifiers. In the $ZH \rightarrow \mu\mu bb$ analysis, we omit the multijet background from our training procedure for precisely this reason, as the data-based multijet model suffers from sizeable fluctuations. Fortunately, this omission does not lead to large degradations in our results, as the multijet background is quite small, and several kinematic variables treat the multijet background similarly to the $t\bar{t}$ background, which we do include in our training sample.

To ensure the validity of our results, it is of utmost importance that the classifier output be well modeled. First, this requires that every variable input to the classifier is well-modeled. Also, we must avoid any mismodeling due to overtraining by using separate samples for training and evaluation of results. In fact, to avoid optimization bias as well, we must separate the evaluation sample from any events that are used to optimize the parameters of the classifier. These events, to which I'll refer as the testing sample, should also be kept separate from the training sample, so as to avoid optimization based on statistical fluctuations. In the $ZH \rightarrow \mu\mu bb$ analysis, we are careful to keep these samples separate, using 1/4 of our simulated events for training, 1/4 for testing, and the remaining 1/2 for evaluation of results.

The specific type of classifier we've chosen is a *random forest* (RF), which is a collection of *decision trees* (DT's) trained with many random subsets of the total training sample. This classifier was chosen for its stability against statistical fluctuations, its simplicity (which translates to very fast training and ample time available for optimization), and the ease with which it handles complicated correlations among a large number (~ 20) of variables.

5.9.1 Decision Trees

A decision tree divides the n -dimensional phase space defined by its n input variables into many rectangular regions, called *leaves*. Each leaf would ideally contain only signal or only background, but realistically, we associate with each leaf a *purity*, defined as the fraction of events in the leaf that are signal. The output of a DT is just the purity of the leaf that contains the input event. Thus, the DT output ranges from zero to one, with the signal distributed near one and the background near zero.

The goal of a DT is to arrive at leaves containing only signal or only background. Both signal-rich and background-rich leaves are equally desirable, so to quantify this goal we construct a figure of merit that is symmetric about purity $p = 0.5$. A very simple but effective figure of merit is known as the Gini index, defined as $G = p(1 - p)$. Using the Gini index, a perfect classifier would allow one to define a threshold such that events above the threshold have $G = 0$, as do events below the threshold, while the entire collection of events might be maximally mixed with $G = 1/4$. We may compare the effectiveness of thresholds meant to separate signal from background by calculating the average decrease in the Gini index, ΔG , from the parent sample to the two daughter samples defined by each threshold. ΔG is defined in equation 5.4, where the subscript 0 indicates the parent sample, the subscripts 1 and 2 indicate the daughters, and N is the number of events in the indicated sample ($N_0 = N_1 + N_2$).

$$\Delta G = G_0 - \frac{1}{N_0}(N_1 G_1 + N_2 G_2) \quad (5.4)$$

To begin training a decision tree, we define the root node as containing the entire training sample. Considering the full spectrum of every input variable, we find the single variable

and threshold that maximize ΔG for the root node, and we apply this cut to create two daughter nodes. Next, for each daughter node, if it has a sufficient number of events, then find the variable and threshold that maximize ΔG for this node. Apply the cut, and repeat this procedure on the resulting daughter nodes. If a node has fewer events than a specified threshold (usually ~ 100), then the node becomes a leaf (i.e. it is not further divided) with purity defined by the signal and background events it contains from the training sample.

The lower limit on the number of events in a node helps to avoid overtraining. To further reduce overtraining in the decision trees we use in the $ZH \rightarrow \mu\mu bb$ analysis, we also limit the depth of each tree to 15 nested decisions. Even so, a single DT remains sensitive to statistical fluctuations because of its structure. For example, if there are two variables that are uncorrelated but have similar separation power, a fluctuation could well determine which variable is chosen, and this choice can dramatically affect the subsequent structure of the decision tree, potentially resulting in a suboptimal classifier. Fortunately, the techniques used to build a random forest alleviate this problem.

5.9.2 The Random Forest

In a random forest, each decision tree is trained using a random subset of the total training sample. This subset is selected with a technique called *bagging*, or Bootstrap AGGREGatING. If the total training sample contains N events, then we randomly choose an event N times from the full training sample. To preserve the phase-space distribution of events for every choice, events are never removed from the training sample: we are allowed to choose the same event (randomly) multiple times. As duplicate events are removed, this process results in a random sample containing fewer than N events. In the $ZH \rightarrow \mu\mu bb$ analysis, we train 200 decision trees using bagging to select the training sample for each tree, and the output of the random forest is just the average output from all of the trees. In effect, bagging smears the training sample and results in substantially less overtraining than a single DT.

In the $ZH \rightarrow \mu\mu bb$ analysis, we found that it was optimal to further randomize the trees by considering only a random subset of variables when maximizing ΔG at each node in each tree. The optimal number of variables to consider at each node was roughly half

of the total number of variables. We interpret this as specifically smearing the correlations between variables, which can be particularly sensitive to statistical effects.

For implementation of the Random Forest classifier in the $ZH \rightarrow \mu\mu bb$ analysis, we used the TMVA software package [47].

5.9.3 Input Variables

As it is difficult to predict which variables will be useful after a series of nested cuts, from the outset we considered as many variables as possible and paid little attention to the apparent separation power of each variable. First, we narrowed the list by requiring that the pre-tag distribution of the variable be well modeled. We computed the Kolmogorov-Smirnov (KS) and χ^2 probabilities for each pre-tag distribution and rejected variables for which both probabilities were less than 5%. Then, preliminary RF's were trained in each $ZH \rightarrow \ell\ell bb$ leptonic channel in order to rank the variables according to how often they were chosen to optimize ΔG . As expected, any variables dependent on lepton type did not rank highly, so the most important 20 variables common to all channels were chosen for the training of a common random forest. Merging the leptonic channels for the training of the RF provides maximum training statistics, thus protecting us from overtraining, and also relieves some sensitivity to channel-specific systematic errors.

While the leptonic channels were merged to train the random forests, separate RF's were trained using the 1T and 2L b -tagging criteria, as the different tag selections result in different mixtures of backgrounds. For example, the $t\bar{t}$ background is more prominent in the 2L selection. Also, because jets and b tagging are significantly different between the two data epochs, separate RF's were trained for RunIIa and RunIIb. In RunIIa, we observed that the post-kinematic-fit jet p_T was not well modeled, so those two variables were omitted from the training of the RunIIa RF's.

The four distinct variables most often used by the random forests are the dijet invariant mass, the leading jet p_T , $\Delta\phi$ between the two leptons, and the scalar sum of p_T from the two jets and two leptons. Our dijet mass distributions are shown in figure 5.12, while distributions of the other three variables are shown in figure 5.13. All 20 variables used as

input to the RF are described below.

Jet Variables

These jet-related variables were used as inputs to the RF discriminant:

- dijet invariant mass, both before and after the kinematic fit;
- p_T of each jet in the dijet pair, both before and after the kinematic fit;
- p_T of the dijet system;
- $\Delta\phi$ and $\Delta\eta$ between jets in the dijet pair; and
- invariant mass and p_T of the system of all jets in the event.

Lepton Variables

These lepton-related variables were used as inputs to the RF discriminant:

- p_T of the dilepton system;
- $\Delta\phi$ and $\Delta\eta$ between leptons in the dilepton system; and
- collinearity of the leptons in the dilepton system.

Global Variables

These variables, related to both jets and leptons, were used as inputs to the RF:

- invariant mass of the $\ell\ell b\bar{b}$ system;
- p_T of the $\ell\ell b\bar{b}$ system;
- scalar sum of lepton and jet p_T 's in the $\ell\ell b\bar{b}$ system;
- $\Delta\phi$ between the dilepton and dijet systems; and

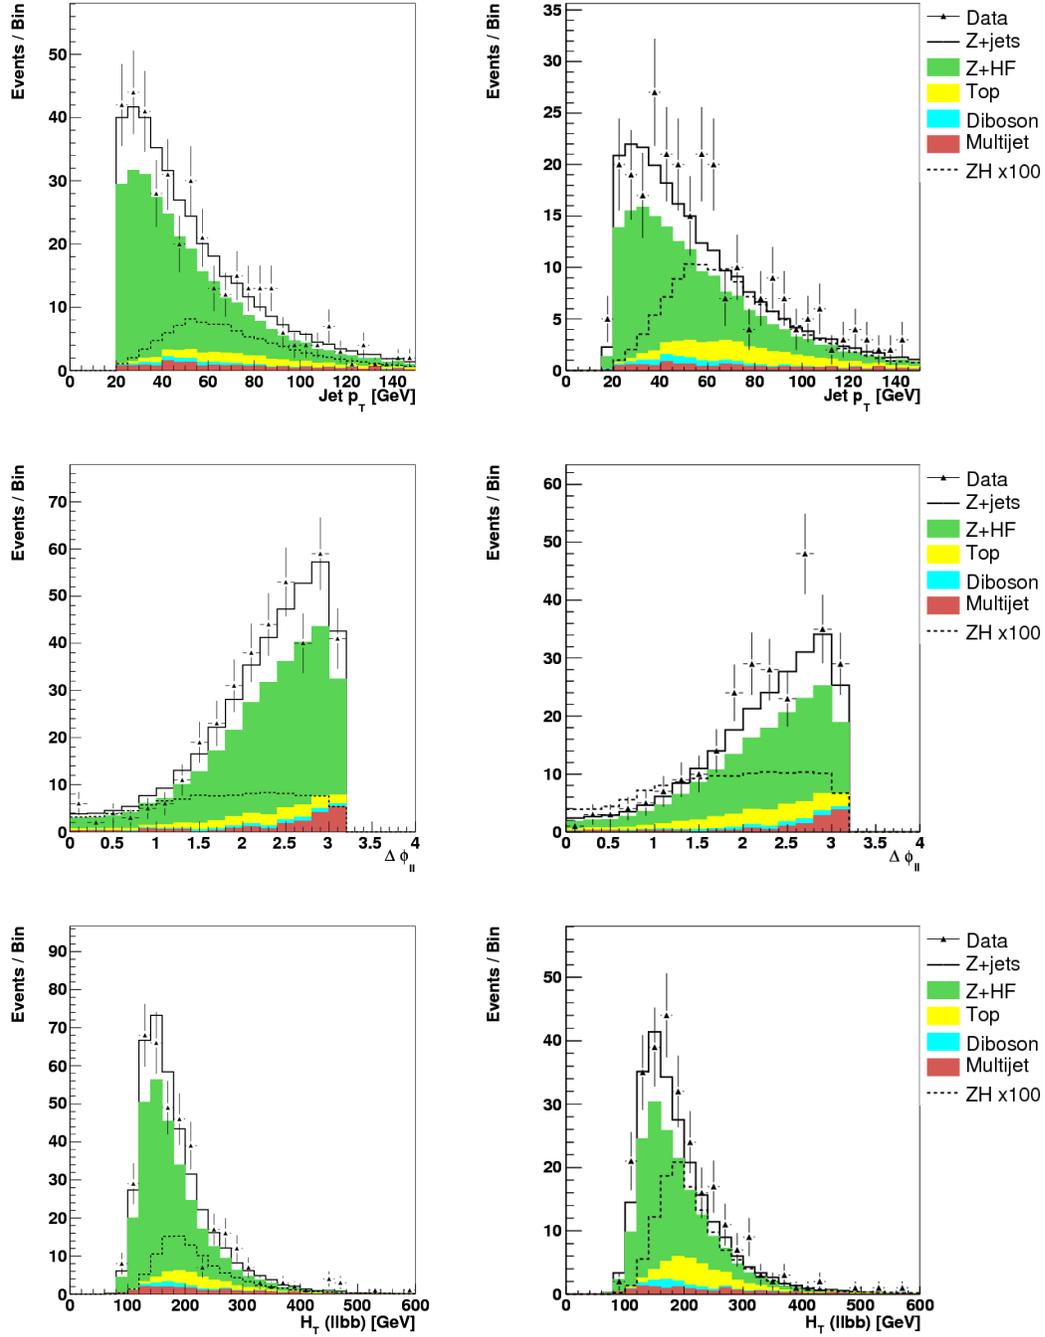


Figure 5.13: Distributions of leading jet p_T (top), $\Delta\phi$ between muons (middle), and scalar sum of p_T from the two jets and two muons (bottom). On the left are 1T-selected events; on the right, 2L-selected events.

- cosine of the angle θ^* between the beam and the dilepton system, in the rest frame of the initial state.

5.9.4 Performance

In addition to having separate random forests depending on data epoch and b -tag selection, we trained different RF's for each hypothetical Higgs mass between 100 and 150 GeV in 5-GeV steps. The mass of the Higgs boson has a large impact on the expected kinematics of the signal, so we expect that entirely different decision trees will be optimal at different values of m_H . For example, the optimal thresholds on the dijet mass should clearly depend on m_H . Figure 5.14 shows the pre-tag output distributions of the RF's trained using $m_H = 115$ GeV. In these high-statistics plots, it is clear that the RF output is well modeled. Figure 5.15 shows the same RF outputs in 1T- or 2L-tagged events. These are the distributions we use to set limits on ZH production at $m_H = 115$ GeV.

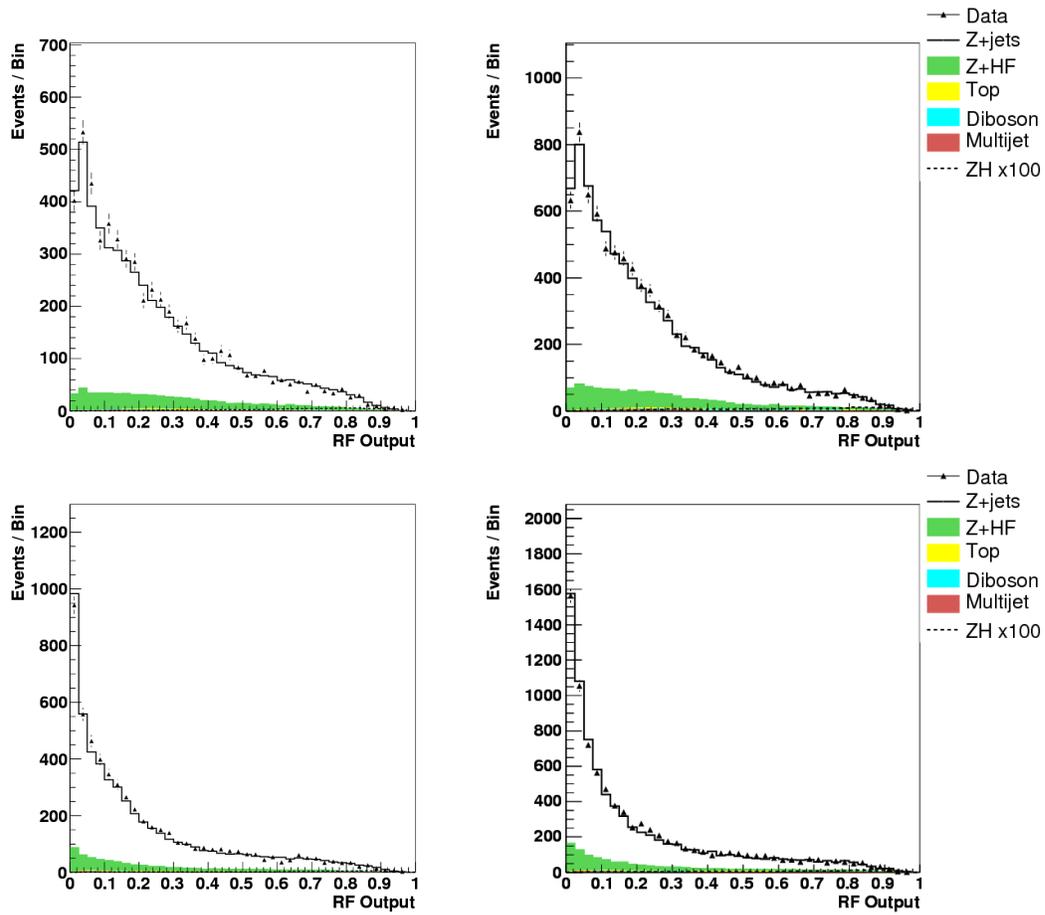


Figure 5.14: Pre-tag RF output distributions in RunIIa (left) and RunIIb (right), using the 1T-trained RF's (top) and the 2L-trained RF's (bottom).

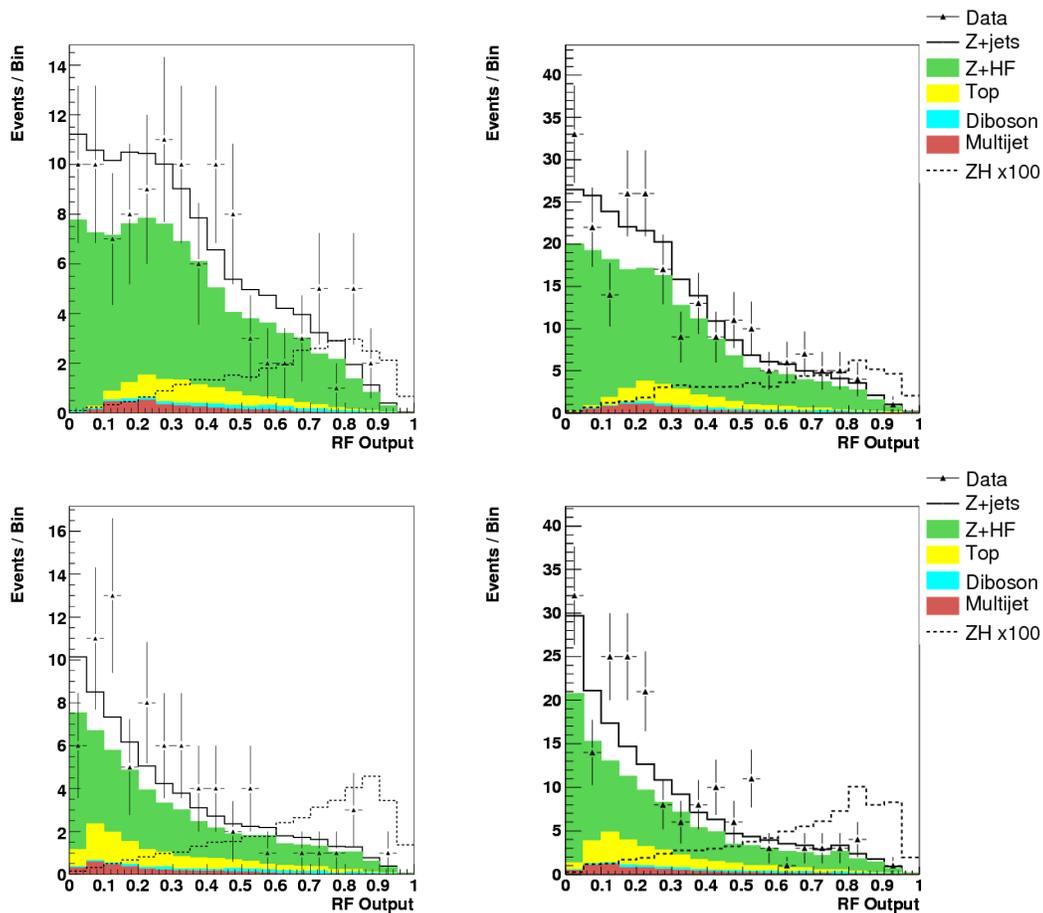


Figure 5.15: Post-tag RF output distributions in RunIIa (left) and RunIIb (right), with 1T tagging criteria (top) and 2L tagging criteria (bottom).

5.10 Systematic Uncertainties

In addition to the statistical uncertainties inherent to samples of discrete events, we must account for the systematic uncertainties in our methods of modeling the expected backgrounds and signal. Because this analysis is based purely on a comparison between the data and the expected background and signal predictions, we ignore any systematic uncertainty in our interpretation of the data. For example, the JES may not be entirely correct when applied to b jets, but so long as this error is common between the data and our background and signal models, it does not invalidate our results. If a 115-GeV Higgs were apparent in our data, we would expect to see a dijet invariant mass peak closer to 100 GeV; thus, we would not infer the Higgs mass directly from this dijet mass distribution, but instead via comparison with signal hypotheses at several values of the Higgs mass. The only systematic uncertainties we consider are those which could lead to disagreements between the data and the background and signal models, thereby hiding a true signal or creating a false signal.

We can further restrict the types of systematic errors of concern to us by considering the effect of our background normalization procedure. If the effect of an error is to increase or reduce the event yield of all MC samples equally, the normalization procedure compensates for it. After normalization, the event yields are insensitive to this type of error. Therefore, we consider only those systematic uncertainties with the potential to change the shape of the modeled RF output distribution. Most notably, we ignore the uncertainty in the integrated luminosity measurement. Our application of the measured luminosity to all MC samples only serves to ensure that k_ϵ and k_Z in equation 5.1 are close to unity.

It is useful to distinguish between two remaining types of systematic uncertainties of importance to this analysis. The simpler type is the *flat systematic*, which only affects the event yield of a particular background or signal model. Although a flat systematic does not alter the shape of the RF output distribution for the specific model in question, the shape of the total RF output distribution may change due to one background or signal model increasing in yield with respect to the others. The canonical example of a flat systematic is the theoretical uncertainty in the cross section for a particular background process. The other type of systematic uncertainty arises because we use the full shape of

the RF output to obtain results. So-called *shape systematics* are applied on a per-event basis and cannot be reduced to an overall scale applied to an entire background or signal model. Shape systematics arise from the identification and calibration of simulated jets, from the application of *b*-tagging TRF's, and from reweightings applied to improve the accuracy of our MC samples, to name a few sources. In this section, we describe the sources and our handling of all flat and shape systematics considered for this analysis.

5.10.1 Flat Systematics

While our normalization procedure allows us to ignore some systematic uncertainties, the resultant event yields remain somewhat uncertain, and those uncertainties have several sources. We have attempted to divide the uncertainties into components corresponding to specific sources, as the source of an error controls how the error should be handled in our limit-setting procedure. In the following list, I describe the uncertainties (all flat systematics) arising from our background normalization procedure. All numbers quoted below are fractional errors indicating how much the actual value might vary from the nominal value.

- The post-tag multijet background yield is uncertain because there are very few directly tagged events in the multijet sample. Also, our application of the light-jet TRF has some effect on the shape of the multijet sample. This effect may be erroneous, considering the multijet sample does in fact contain heavy-flavor jets. Finally, our method of selecting the multijet sample may result in exaggerated heavy-flavor content in comparison to the actual multijet background. To cover all of these potential sources of error, a 30% uncertainty on the multijet background is required in RunIIa, and 20% is required in RunIIb.
- The background normalization fit gives us a statistical uncertainty in k_Z^2 , due to limited statistics in pre-tag data. This uncertainty amounts to 1.14% in RunIIa and 0.9% in RunIIb, applied only to our Z +jets background models.
- When the normalization procedure is performed using only the $\mu\mu$ leptonic channel, the obtained value of k_Z^2 is different by 2%. Because k_Z^2 should not depend on lepton

Final State	Uncertainty
$Z/\gamma^* + bb/cc$	20%
Diboson	7%
$t\bar{t}$	10%
ZH	6%

Table 5.8: Theoretical cross-section uncertainties.

flavor, we interpret this as an indication that k_ϵ may be in error. Thus, we assign a 2% uncertainty to all MC samples.

- The values of k_Z^2 obtained in the separate RunIIa and RunIIb normalization fits differ by 9%. Since k_Z^2 should not depend on data epoch, we interpret this as a sign that k_Z^2 is compensating for an error in simulated jet identification. Thus, it may be that k_Z^2 should be applied to all MC samples and not only to Z +jets. To reflect this ambiguity in our method, we assign a 9% uncertainty to all non- Z +jets MC samples.
- Our normalization procedure assumes that the inclusive Z production cross section is known with infinite precision, but of course this is not the case. We assign a 6% uncertainty to the inclusive Z cross section, and apply this uncertainty to all non- Z +jets MC samples.

The remaining flat systematics are the theoretical uncertainties in the background and signal cross sections. Table 5.8 shows the fractional uncertainties we assign to each simulated background and signal process. Associated production of Z bosons with light-flavor jets is absent from the table because it is constrained by the background normalization procedure.

5.10.2 Shape Systematics

The following list describes all shape systematics considered in this analysis. It is difficult to assign a number to quantify a shape systematic, but tables 5.9 and 5.10 provide a sense

for the impact of the various uncertainties with weighted averages of the unsigned shifts in event yield over all RF output bins.

- There are 40 independent uncertainties associated with the PDF's used in generating our MC samples, but most have very little effect on our final RF output distributions. We choose the two uncertainties with the largest effects on our event yields, and add them in quadrature. The resultant PDF systematic is applied to all MC samples.
- The inclusive trigger correction is uncertain, as it is derived from a sample of data with limited statistics. Because the overall trigger efficiency is constrained by our normalization procedure, we consider only the differences in shape that could result from an erroneous trigger correction. Without changing the total event yield, we exaggerate and diminish the shape of the trigger correction according to the statistical uncertainty in the ratio of inclusive to single-muon triggered data.
- Since jet identification is more efficient in MC than in data, simulated jets are randomly rejected according to jet ID scale factors, as described in section 4.3.3. To assess the effect of a potential error in the jet ID SF's, we shift the SF's down by their uncertainty and redo the random rejection of identified jets. Using these new events with substantially different jet content, we repeat our selection and normalization procedures, and we reevaluate the kinematic fit and the RF. This yields a shifted RF output distribution, and the difference from the nominal RF output distribution is the shape systematic from the uncertainty in the jet ID SF's.
- Since the jet energy resolution is better in MC than in data, simulated jets are randomly smeared to match the energy resolution as measured in data (see section 4.3.3). However, the amount by which jets are smeared is uncertain, and we assess the impact of this uncertainty by under- and over-smearing the jets in every MC sample. This results in samples of events with different jet content than the nominal sample, so we must repeat our jet selection and normalization using these new events. We thus

obtain shifted RF output distributions that describe the shape systematic from the uncertainty in the jet energy resolution correction.

- Similar to the treatment of the jet ID and jet resolution systematics, shifting the jet energy scale up and down by its uncertainty results in new events with different jet content. Again, jet selection and normalization are redone to yield shifted RF output distributions that describe the shape systematic from the uncertainty in the energy scale for simulated jets.
- In RunIIb, the requirement that jets be vertex-confirmed is effectively an addition to the jet definition, so the uncertainty in the corresponding vertex confirmation scale factors is handled in much the same way as the uncertainty in the jet ID scale factors. A shifted RF distribution is obtained that describes the effect of varying the vertex confirmation SF's down by their uncertainty, and varying the SF's up is assumed to have the opposite effect on the RF output distribution.
- Taggability SF's derived using events with only one jet are substantially different from those derived using two-jet events. Since the taggability SF's should not depend on jet multiplicity, we define their uncertainties to be large enough to cover this difference. Because we apply taggability scale factors as event weights, we adjust the weights up and down according to their uncertainties in order to assess the shape systematic from the uncertainty in our taggability SF's.
- It is known that the b -tagging TRF's for b and c jets are highly correlated, while the light jet fake rate is independent of the other TRF's. To approximate these correlations, we correlate all TRF's in MC samples enriched with heavy flavor, shifting all TRF's up and down by their uncertainties to obtain shifted weights for each event. The resultant, shifted RF output distributions describe the shape systematic from the uncertainty in the heavy-flavor TRF's. Likewise, we shift all TRF's up and down in a correlated fashion in all MC samples devoid of heavy-flavor jets to obtain the shape

systematic from the uncertainty in the rate at which light jets are misidentified as b jets.

- Uncertainty in the $Z p_T$ reweighting comes entirely from limited statistics in the $Z \rightarrow ee$ data used to derive the reweighting. We adjust Z +jets background event weights up and down according to this uncertainty, which yields shifted RF output distributions that describe the shape systematic from $Z p_T$ reweighting.
- Similar to $Z p_T$ reweighting, uncertainty in the jet angle reweighting comes entirely from limited statistics in the data used to derive the reweighting. Again, we adjust Z +jets background event weights up and down to assess the shape systematic from jet angle reweighting.

RunIIa							
Systematic	ZH (115)	Zjj	Zbb	Zcc	Diboson	$t\bar{t}$	Multijet
PDF	0.9	0.6	1.7	0.9	0.2	0.7	
Inclusive Trigger	1.9	1.9	2.0	2.1	2.2	1.9	
Jet ID	1.0	0.6	0.5	0.4	0.6	0.7	
Jet Resolution	3.3	3.9	8.9	7.8	5.4	5.5	
JES	2.4	1.0	3.5	1.7	3.5	4.0	
Taggability SF	2.4	2.2	2.5	2.5	2.4	2.5	
TRF (heavy)	1.0		2.3	4.0		0.4	
TRF (light)		16			3.8		1.6
$Z p_T$ RW		2.2	2.3	2.1			
Jet Angle RW		0.9	0.9	0.9			

RunIIb							
Systematic	ZH (115)	Zjj	Zbb	Zcc	Diboson	$t\bar{t}$	Multijet
PDF	1.0	0.7	1.6	1.0	0.2	0.3	
Inclusive Trigger	1.8	1.7	1.8	1.8	1.8	1.7	
Jet ID	1.5	0.1	0.4	0.3	1.5	1.0	
Jet Resolution	1.3	2.2	4.9	3.3	1.5	2.4	
JES	1.9	0.9	2.4	2.8	3.1	3.8	
VCJ SF	3.4	0.4	0.7	0.7	3.6	3.4	
Taggability SF	0.6	0.4	0.4	0.5	0.7	0.7	
TRF (heavy)	0.03		2.3	3.6		0.4	
TRF (light)		19			4.0		4.2
$Z p_T$ RW		2.4	2.6	2.3			
Jet Angle RW		0.9	0.9	0.9			

Table 5.9: Impact of shape systematics on 1T-selected events. Each number is a weighted average of the (unsigned) percent change in event yield over all RF output bins for a particular signal or background model, upon variation of a particular systematic uncertainty by one standard deviation.

RunIIa							
Systematic	ZH (115)	Zjj	Zbb	Zcc	Diboson	$t\bar{t}$	Multijet
PDF	0.7	0.6	1.5	0.6	0.3	0.7	
Inclusive Trigger	2.0	2.1	2.1	2.1	2.2	2.0	
Jet ID	1.1	0.7	0.7	0.8	0.6	0.8	
Jet Resolution	3.5	4.9	12	10	3.8	3.9	
JES	3.1	1.0	3.0	3.2	3.7	3.4	
Taggability SF	4.9	3.9	4.6	4.1	4.5	5.5	
TRF (heavy)	5.6		4.6	5.3		6.5	
TRF (light)		4.9			4.5		0.4
Z p_T RW		2.2	2.1	1.9			
Jet Angle RW		0.8	0.9	0.9			

RunIIb							
Systematic	ZH (115)	Zjj	Zbb	Zcc	Diboson	$t\bar{t}$	Multijet
PDF	0.7	0.7	1.5	1.3	0.1	0.4	
Inclusive Trigger	1.8	2.1	1.9	1.8	2.2	1.8	
Jet ID	1.5	0.1	0.4	0.3	1.5	1.7	
Jet Resolution	2.0	1.3	7.9	3.7	4.2	0.9	
JES	2.7	0.4	2.5	3.0	5.1	3.1	
VCJ SF	4.2	0.5	0.7	0.7	3.6	3.4	
Taggability SF	0.9	0.8	0.9	0.8	1.0	1.2	
TRF (heavy)	5.0		4.1	4.5		6.7	
TRF (light)		4.8			4.4		1.0
Z p_T RW		2.5	2.3	2.1			
Jet Angle RW		0.9	1.0	0.9			

Table 5.10: Impact of shape systematics on 2L-selected events. Each number is a weighted average of the (unsigned) percent change in event yield over all RF output bins for a particular signal or background model, upon variation of a particular systematic uncertainty by one standard deviation.

5.11 Limit Setting

If an event is unlikely to occur but has a large number of opportunities to occur, the number of occurrences should follow a Poisson distribution. With 1.5 million bunch crossings but only 100 events recorded per second, the data collected by DØ are well-described by a Poisson distribution. Furthermore, the likelihood that a given bunch crossing will provide an event that passes the selection cuts in this analysis is exceedingly small. Therefore, we describe the probability of d observed events in our data, given p predicted events in our background and signal models, using equation 5.5. If we consider multiple independent bins of data and predicted events, then each bin i is governed by the Poisson distribution, and the total probability is just the product of all individual bin probabilities. In this analysis, we consider several RF output bins for each data epoch and b -tag selection. By keeping the 1T and 2L RF output distributions separate, we avoid diluting the high signal significance resulting from the 2L selection. By separating RunIIa from RunIIb, we retain potentially useful information as to how S/B changes from one data epoch to another.

$$P(d|p) \approx \prod_i \frac{p_i^{d_i} e^{-p_i}}{d_i!} \quad (5.5)$$

When attempting to determine whether a Higgs signal is present in the data, the background-only hypothesis H_B must be compared to the signal+background hypothesis H_{S+B} . Each hypothesis yields a predicted number of events, and the Poisson distribution provides a measure of compatibility with the data. We construct the log-likelihood-ratio (LLR) \mathcal{L} , defined in equation 5.6, to quantify which hypothesis is more compatible with the data. If $\mathcal{L} > 0$, then the data prefers H_B ; if $\mathcal{L} < 0$, then the H_{S+B} is preferred.

$$\mathcal{L}(d|H_B, H_{S+B}) = -2 \ln \left(\frac{P(d|p(H_{S+B}))}{P(d|p(H_B))} \right) \quad (5.6)$$

Complications arise from the fact that the predicted number of events for a given hypothesis is uncertain. Systematic errors could lead to false predictions, so we need a method of taking systematic uncertainties into account. For each hypothesis, we allow the predicted number of events to vary according to the systematic uncertainties, and we maximize compatibility with the data. In other words, we minimize the quantity λ , defined in equation

5.7 [48]. Here, \vec{s} is a vector whose components represent the variations of the prediction in the “directions” of the various systematic uncertainties. For example, if $\vec{s} = (1, 0, -2)$, then there are three systematic uncertainties, and we have varied the prediction by shifting the first systematic up by one standard deviation and the third systematic down by two standard deviations. We minimize λ by varying \vec{s} , which in turn varies the predicted numbers of events $p(H, \vec{s})$. In this step, we have approximated $\ln(d_i!) \approx d_i \ln(d_i) - d_i$, which is valid for moderately high values (~ 100 or greater) of d_i .

$$\begin{aligned} \lambda(d|H, \vec{s}) &= 2 \sum_i \left(p_i(H, \vec{s}) - d_i - d_i \ln \left(\frac{p_i(H, \vec{s})}{d_i} \right) \right) + |\vec{s}|^2 \\ &\approx -2 \ln(P(d|p(H, \vec{s}))) + |\vec{s}|^2 \end{aligned} \quad (5.7)$$

Let $\vec{s}_B(d)$ be the systematic shifts that minimize λ , given the background-only hypothesis H_B , and let $\vec{s}_{S+B}(d)$ be the systematic shifts that minimize λ , given the signal+background hypothesis H_{S+B} . We redefine the LLR in terms of the minimal values of λ , using equation 5.8. Taking into account the systematic uncertainties in predicted event yields, the quantity \mathcal{L}' indicates the hypothesis most compatible with the data.

$$\mathcal{L}'(d|H_B, H_{S+B}) = \lambda(d|H_{S+B}, \vec{s}_{S+B}(d)) - \lambda(d|H_B, \vec{s}_B(d)) \quad (5.8)$$

To determine the statistical significance of the observed value of \mathcal{L}' , we generate a large number of pseudoexperiments (PE’s) for both hypotheses. To generate a PE, we randomly sample an n -dimensional Gaussian distribution with unit variance and mean at the origin, where n is the number of systematic uncertainties. Thus, we obtain a systematic shift vector \vec{s} and an associated prediction $p(H, \vec{s})$. This prediction is randomly fluctuated according to Poisson statistics to produce pseudodata \tilde{d} . For each PE, we evaluate $\mathcal{L}'(\tilde{d}|H_B, H_{S+B})$. The collection of \mathcal{L}' values from all background-only PE’s yields a probability distribution $f_B(\mathcal{L}')$. Likewise, the $S+B$ PE’s yield a probability distribution $f_{S+B}(\mathcal{L}')$. Minimal overlap between f_B and f_{S+B} would indicate high sensitivity to the signal.

The value of \mathcal{L}' at the peak of the f_B distribution is the expected LLR ($\mathcal{L}'_{\text{exp}}$) in the absence of any signal, while the peak of the f_{S+B} distribution gives the expected LLR if the

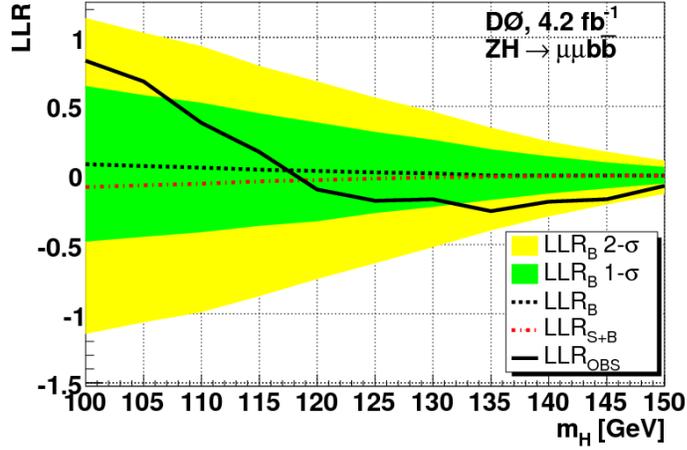


Figure 5.16: Observed and expected LLR values versus hypothetical Higgs mass, using only the $\mu\mu$ leptonic channel. The solid black line is the observed LLR, the dashed black line is the background-only expected LLR, and the dashed red line is the signal+background expected LLR. The green and yellow shaded areas are the 1σ and 2σ intervals for the background-only LLR distribution.

$S + B$ hypothesis is correct. If the observed LLR ($\mathcal{L}'_{\text{obs}} = \mathcal{L}'(d)$) were to differ from $\mathcal{L}'_{\text{exp}}$ by more than two standard deviations of f_B , we would begin to doubt either the modeling of the background or the background-only hypothesis itself. As shown in figure 5.16, this is not the case, so we conclude that the background-only hypothesis is consistent with the data, and we proceed to set upper limits on Higgs production.

Because we aim to exclude the $S + B$ hypothesis, we are interested in the integrals of f_B and f_{S+B} from the observed LLR to infinity in the background-like (positive) direction. Shown explicitly in equations 5.9 and 5.10, each integral yields a confidence level (CL) for the one-sided confidence interval $[\mathcal{L}'(d), \infty)$. To take into account potential mismodeling of the background, we only reject the $S+B$ hypothesis if the background-only hypothesis shows significantly better agreement with the data. Thus, we use the quantity $\text{CL}_S = \text{CL}_{S+B}/\text{CL}_B$ to set limits on Higgs production [49]. If $\text{CL}_S = 0.05$, this means a background-only PE is 20 times more likely than a $S + B$ PE to be at least as background-like as the data, so we would exclude the signal with 95% confidence.

$$\text{CL}_B(d) = \int_{\mathcal{L}'(d)}^{\infty} f_B(\mathcal{L}') d\mathcal{L}' \quad (5.9)$$

$$\text{CL}_{S+B}(d) = \int_{\mathcal{L}'(d)}^{\infty} f_{S+B}(\mathcal{L}') d\mathcal{L}' \quad (5.10)$$

With so much overlap between the f_B and f_{S+B} distributions, CL_S is actually quite close to one, so the confidence with which we can exclude the expected, standard-model Higgs signal is very small. To place upper limits on Higgs production, we exaggerate the signal expectation until $\text{CL}_S = 0.05$, so the exaggerated signal is excluded with 95% confidence. We report the scale factor by which the signal was exaggerated as our observed limit, and we are 95% confident that the actual Higgs production rate must be lower than this factor times the SM expectation. Additionally, we report expected limits obtained from the pseudodata that maximize f_B . As such, the expected limits assume that no Higgs signal exists, so the first hint of an actual Higgs signal would be an observed limit significantly higher than the corresponding expected limit.

In table 5.11 and figure 5.17, the observed and expected limits on ZH production are presented, using only the $\mu\mu$ leptonic channel. This entire limit-setting procedure was repeated 11 times for different values of the Higgs mass, ranging between 100 and 150 GeV in 5-GeV steps. The limits are smallest at low values of m_H because that is where the expected rate of $ZH \rightarrow \ell\ell bb$ production is largest. One might be concerned that there are sizable differences between the expected and observed limits at certain values of the Higgs mass, but it is important to note that the data is everywhere compatible with the background-only hypothesis (see figure 5.16), so these differences are not significant.

As the $ZH \rightarrow \mu\mu bb$ analysis is part of a larger $ZH \rightarrow \ell\ell bb$ effort, limits were also determined from the combination of the $\mu\mu$ channel with all other leptonic channels (ee , ee_{ICR} , and $\mu\mu_{\text{TRK}}$). The expected and observed LLR values obtained from this combination are plotted in figure 5.18, while expected and observed limits for all 11 hypothetical values of the Higgs mass are shown in figure 5.19 and table 5.12.

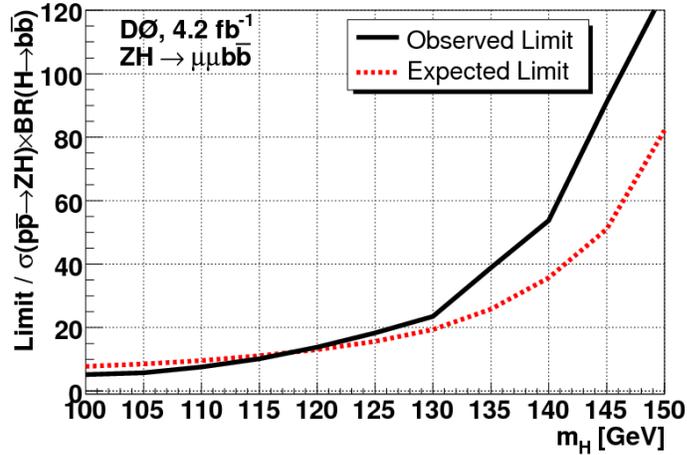


Figure 5.17: Expected and observed upper limits (95% CL) on ZH production versus hypothetical Higgs mass, using only the $\mu\mu$ leptonic channel. The solid black line is the observed limit, and the dashed red line is the background-only expected limit.

m_H (GeV)	Expected Limit	Observed Limit
100	7.7	5.2
105	8.5	5.7
110	9.5	7.6
115	11.1	10.1
120	13.0	13.8
125	15.6	18.4
130	19.4	23.5
135	25.8	38.8
140	35.7	53.8
145	51.0	91.0
150	82.1	126.0

Table 5.11: Expected and observed upper limits on ZH production, using only the $\mu\mu$ leptonic channel, for 11 hypothetical values of the Higgs mass between 100 and 150 GeV. The limits are expressed in units of the expected SM production rate.

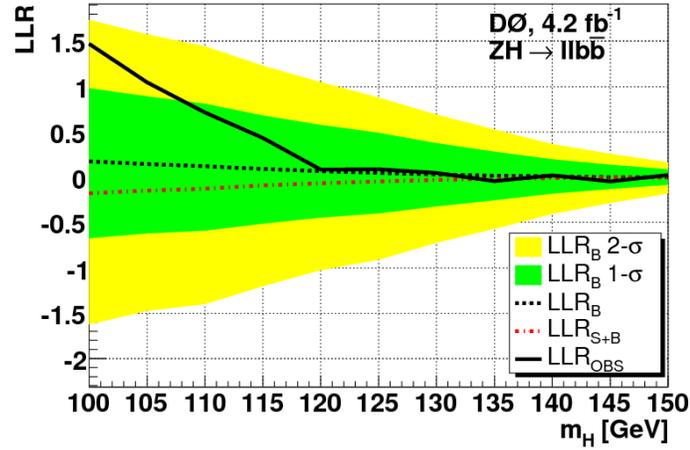


Figure 5.18: Observed and expected LLR values versus hypothetical Higgs mass, combining all $ZH \rightarrow \ell\ell b\bar{b}$ leptonic channels. The solid black line is the observed LLR, the dashed black line is the background-only expected LLR, and the dashed red line is the signal+background expected LLR. The green and yellow shaded areas are the 1σ and 2σ intervals for the background-only LLR distribution.

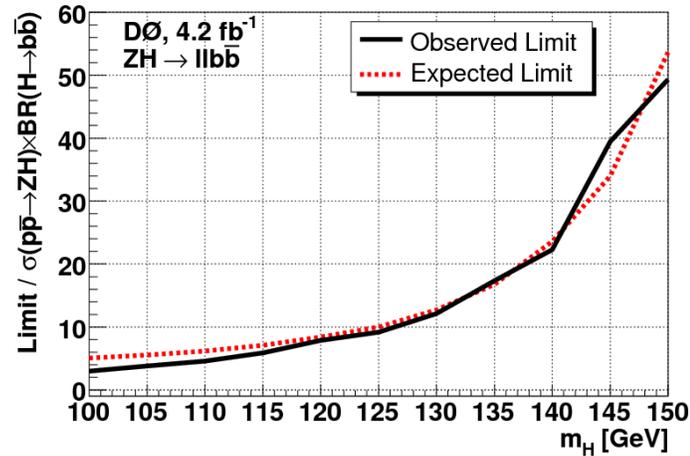


Figure 5.19: Expected and observed upper limits (95% CL) on ZH production versus hypothetical Higgs mass, combining all $ZH \rightarrow \ell\ell b\bar{b}$ leptonic channels. The solid black line is the observed limit, and the dashed red line is the background-only expected limit.

m_H (GeV)	Expected Limit	Observed Limit
100	5.1	3.0
105	5.6	3.4
110	6.3	4.3
115	7.2	5.4
120	8.4	8.0
125	9.8	9.6
130	12.4	13.3
135	16.4	17.0
140	23.4	22.7
145	33.4	42.8
150	53.0	50.6

Table 5.12: Expected and observed upper limits on ZH production, combining all $ZH \rightarrow \ell\ell b\bar{b}$ leptonic channels, for 11 hypothetical values of the Higgs mass between 100 and 150 GeV. The limits are expressed in units of the expected SM production rate.

Chapter 6

CONCLUSION

We have presented the results of a search for associated production of Z and Higgs bosons in the $\mu\mu bb$ final state, using 4.2 fb^{-1} of $\bar{p}p$ collisions collected with the DØ detector. We employed several sophisticated analysis techniques, such as the kinematic fit and the multivariate random forest discriminant. To reduce the impact of potential systematic errors, we used the data to constrain many uncertainties in our signal and background models. As no significant excess over background was observed in the data, we evaluated 95% C.L. upper limits on ZH production by comparing the data with large collections of background-only and signal+background pseudoexperiments. Assuming a Higgs boson of mass 115 GeV, we exclude ZH production above 10.1 times the standard-model expectation, using only the $\mu\mu bb$ leptonic channel. Combining all $ZH \rightarrow \ell\ell bb$ leptonic channels, we exclude ZH production above 5.4 times the standard-model expectation. These results were recently accepted for publication by Physical Review Letters [50].

No single Higgs search at the Tevatron is expected to be sensitive to the Higgs boson. Sensitivity is achieved by combining channels and production modes, and even by combining results from DØ and CDF. The latest results from the Tevatron-wide combination [4], using up to 6.7 fb^{-1} of data, were able to exclude Higgs bosons of mass between 158 and 175 GeV using the dominant decay for high-mass Higgs bosons, $H \rightarrow W^+W^-$. In the lower mass range where $ZH \rightarrow \ell\ell bb$ contributes substantial sensitivity, the combined 95% C.L. upper limits are all less than ~ 2.5 times the SM expectation. At $m_H = 115 \text{ GeV}$, Higgs production above $1.6 \times \text{SM}$ is excluded.

Already, more than 8.7 fb^{-1} of $\bar{p}p$ collisions have been recorded by DØ, with data accumulating at over 2 fb^{-1} per year. By the end of RunII, analyses at the Tevatron will be searching for the Higgs in at least 10 fb^{-1} of data, which will bring them close to the sensitivity that would be required to exclude the Higgs boson across the entire range

of allowed masses. Expected improvements to each analysis will further boost sensitivity, perhaps allowing the combined Tevatron analysis to report evidence for Higgs production. In the $ZH \rightarrow \ell\ell b\bar{b}$ analysis, efforts are currently underway to improve the kinematic fit by treating initial- and final-state radiation properly, to optimize b tagging by providing variables sensitive to b jets as inputs to the multivariate classifier, and to increase Higgs sensitivity with the use of additional multivariate discriminants.

Within the next few years, from observations at the Tevatron or at the Large Hadron Collider, we can expect to have either excluded or found evidence for the Higgs boson, the remaining undiscovered particle in the standard model. In either eventuality, this will be an exciting milestone for particle physics. If the Higgs boson is excluded, this will be the first result from collider experiments that is clearly in conflict with the standard model. If the Higgs boson is discovered, then the study its properties will provide many new paths of research and a rich variety of measurements potentially sensitive to particle physics beyond the standard model.

BIBLIOGRAPHY

- [1] M. Steinhauser and T. Seidensticker, Phys. Lett. **B467**, 271 (1999).
- [2] K. Nakamura *et al.*, “The Review of Particle Physics,” J. Phys. G **37**, 075021 (2010).
- [3] ALEPH Collaboration, DELPHI Collaboration, L3 Collaboration, OPAL Collaboration, and the LEP Working Group for Higgs Boson Searches, “Search for the Standard Model Higgs boson at LEP,” Phys. Lett. B **565**, 61 (2003).
- [4] CDF Collaboration, DØ Collaboration, and the Tevatron New Phenomena and Higgs Working Group, “Combined CDF and DØ Upper Limits on Standard Model Higgs-Boson Production with up to 6.7 fb^{-1} of Data,” arXiv:1007.4587 [hep-ex] (2010).
- [5] The LEP Electroweak Working Group, website.
<http://lepewwg.web.cern.ch/LEPEWWG/plots/summer2010/>
- [6] The Tev4LHC Higgs Working Group, website.
<http://maltoni.home.cern.ch/maltoni/TeV4LHC/SM.html>
- [7] A. Djouadi, J. Kalinowski, and M. Spira, “HDECAY: a Program for Higgs Boson Decays in the Standard Model and its Supersymmetric Extension,” Comput. Phys. Commun. **108**, 56 (1998).
- [8] Fermilab Visual Media Services, website.
http://www-visualmedia.fnal.gov/VMS_Site_2/index.shtml
- [9] B. Worthel, “Linac Rookie Book.”
http://www-bdnew.fnal.gov/operations/rookie.books/LINAC_v2.pdf (2006).
- [10] W. Chou, M. Kostin, and Z. Tang, “Stripping Efficiency and Lifetime of Carbon Foils,” Nucl. Instrum. Methods Phys. Res. Sec. A **590**, 1 (2008).
- [11] S. Abachi *et al.*, “The DØ Detector,” Nucl. Instrum. Methods Phys. Res. Sect. A **325**, 393 (1993).
- [12] V. M. Abazov *et al.*, “The Upgraded DØ Detector,” Nucl. Instrum. Methods Phys. Res. Sect. A **565**, 463 (2006).
- [13] S. N. Ahmed *et al.*, “The DØ Silicon Microstrip Tracker,” arXiv:1005.0801 [physics.ins-det] (2010).

- [14] R. Angstadt *et al.*, “The Layer 0 Inner Silicon Detector of the DØ Experiment,” Nucl. Instrum. Methods Phys. Res. Sect. A **622**, 298 (2010).
- [15] A. Bross *et al.*, “Characterization and performance of visible light photon counters (VLPCs) for the upgraded DO detector at the Fermilab Tevatron,” Nucl. Instrum. Methods Phys. Res. Sect. A **477**, 172 (2002).
- [16] B. Casey *et al.*, “The DØ Run 2b Luminosity Constant,” DØ Note 5945 (2009).
- [17] A. Khanov, “HTF: histogramming method for finding tracks. The algorithm description,” DØ Note 3778 (2000).
- [18] G. Borissov, “Ordering a Chaos or... Technical Details of AA Tracking,” All DØ Meeting, February 28, 2003. http://www-d0.fnal.gov/global_tracking/talks/20030228/talk-adm-030228.ps
- [19] H. Greenlee, “The DØ Kalman Track Fit,” DØ Note 4303 (2003).
- [20] A. Schwartzman and C. Tully, “Primary Vertex Reconstruction by Means of Adaptive Vertex Fitting,” DØ Note 4918 (2005).
- [21] A. Schwartzman and M. Narain, “Probabilistic Primary Vertex Selection,” DØ Note 4042 (2002).
- [22] P. Calfayan *et al.*, “Muon Identification Certification for p17 data,” DØ Note 5157 (2007).
- [23] G. Bernardi and S. Trincaz-Duvoid, “Improvement of the NADA Algorithm: Hot Cell Killing in DØ Run II Data,” DØ Note 4057 (2002).
- [24] U. Bassler and G. Bernardi, “Towards a Coherent Treatment of Calorimetric Energies: Missing Transverse Energy, Jets, E.M. Objects and the T42 Algorithm,” DØ Note 4124 (2002).
- [25] G. C. Blazey *et al.*, “Run II Jet Physics: Proceedings of the Run II QCD and Weak Boson Physics Workshop,” arXiv:hep-ex/0005012 (2000); B. Andrieu and E. Busato, “Jet Algorithms in the DØ Run II Software: Description and User’s Guide,” DØ Note 4457 (2004).
- [26] A. Harel, “Jet ID Optimization,” DØ Note 4919 (2005).
- [27] V. M. Abazov *et al.*, “Measurement of the Inclusive Jet Cross Section in $\bar{p}p$ Collisions at $\sqrt{s} = 1.96$ TeV,” Phys. Rev. Lett. **101**, 062001 (2008).

- [28] K. DeVaughan *et al.*, “Jet Energy Scale Determination for DØ Run IIb,” DØ Note 5801 (2009).
- [29] V. M. Abazov *et al.*, “b-Jet Identification in the DØ Experiment,” Nucl. Instrum. Methods Phys. Res. Sect. A **620**, 490 (2010).
- [30] E. Berger (Ed.), “Research directions for the decade. Proceedings, 1990 Summer Study on High-Energy Physics, Snowmass, USA, June 25 – July 13, 1990,” Singapore: World Scientific (1992).
- [31] J. Pumplin *et al.*, “New generation of parton distributions with uncertainties from global QCD analysis,” J. High Energy Phys. **07**, 012 (2002).
- [32] T. Sjöstrand, S. Mrenna, and P. Skands, “PYTHIA 6.4 Physics and Manual,” J. High Energy Phys. **05**, 026 (2006).
- [33] M. L. Mangano *et al.*, “ALPGEN, a generator for hard multiparton processes in hadronic collisions,” J. High Energy Phys. **07**, 001 (2003).
- [34] S. Hoeche *et al.*, “Matching Parton Showers and Matrix Elements,” arXiv:hep-ph/0602031 (2006).
- [35] R. Brun and F. Carminati, “GEANT: Detector Description and Simulation Tool,” CERN Program Library Long Writeup W5013 (1993).
- [36] P. Calfayan, “Combination of Single Muon Triggers in p17 data by using an inclusive OR,” DØ Note 5329 (2008).
- [37] J.-F. Grivaz and N. Makovec, “Shifting, Smearing and Removing Simulated Jets,” DØ Note 4914 (2005).
- [38] T. Gadfort *et al.*, “Performance of the DØ NN b-tagging Tool on p20 Data,” DØ Note 5554 (2008).
- [39] L. Ancu *et al.*, “Search for $ZH \rightarrow \ell^+ \ell^- b\bar{b}$ in 4.2 fb⁻¹ of Data Using a Random Forest in Dimuon and Dielectron Events,” DØ Note 6071 (2010).
- [40] R. Hamberg, W.L. van Neerven, and T. Matsuura, Nucl. Phys. B **359**, 343 (1991), [*Erratum ibid.* B **644**, 403 (2002)].
- [41] N. Kidonakis and R. Vogt, “Theoretical top quark cross section at the Fermilab Tevatron and the CERN LHC,” Phys. Rev. D **78**, 074005 (2008).

- [42] J. M. Campbell and R. K. Ellis, “Vector boson pair production at the Tevatron, including all spin correlations of the boson decay products,” *Phys. Rev. D* **60**, 113006 (1999).
- [43] K. A. Assamagan *et al.*, “The Higgs Working Group: Summary Report 2003,” arXiv:hep-ph/0406152 (2004); O. Brein, A. Djouadi, and R. Harlander, “NNLO QCD corrections to the Higgs-strahlung processes at hadron colliders,” *Phys. Lett. B* **579**, 149 (2004); M. L. Ciccolini, S. Dittmaier, and M. Kramer, “Electroweak radiative corrections to associated WH and ZH production at hadron colliders,” *Phys. Rev. D* **68**, 073003 (2003); J. Baglio and A. Djouadi, “Predictions for Higgs production at the Tevatron and the associated uncertainties,” *J. High Energy Phys.* **10**, 064 (2010).
- [44] V. M. Abazov *et al.*, “Measurement of the Shape of the Boson-Transverse Momentum Distribution in $\bar{p}p \rightarrow Z/\gamma^* \rightarrow e^+e^- + X$ Events Produced at $\sqrt{s} = 1.96$ TeV,” *Phys. Rev. Lett.* **100**, 102002 (2008).
- [45] W. Fisher *et al.*, “Studies of Alpgen parameters, corrections and associated uncertainties,” DØ Note 5966 (2009).
- [46] E. A. Strauss, “From ZZ to ZH: How Low Can These Cross Sections Go or Everybody, Let’s Cross Section Limbo!”, FERMILAB-THESIS-2009-29 (2009).
- [47] A. Hoecker *et al.*, “TMVA: Toolkit for Multivariate Data Analysis,” arXiv:physics/0703039 [physics.data-an] (2007).
- [48] W. Fisher, “Systematics and limit calculations,” FERMILAB-TM-2386-E (2007).
- [49] A. L. Read, “Presentation of search results: the CLs technique,” *J. Phys. G* **28**, 2693 (2002).
- [50] V. Abazov *et al.*, “Search for ZH to llbb production in 4.2 fb^{-1} of ppbar collisions at $\sqrt{s} = 1.96$ TeV,” accepted for publication in *Phys. Rev. Lett.*, arXiv:1008.3564 [hep-ex] (2010).

VITA

John BackusMayes was born in Worcester, Massachusetts, in 1984. He graduated from Lakeridge High School in Lake Oswego, Oregon, in 2001. In 2005, he earned a Bachelor of Arts degree with honors in physics from Williams College, Williamstown, Massachusetts. After completing two years of graduate-level courses and passing the qualifying exam in physics at the University of Washington in Seattle, in 2007 he relocated to Chicago, Illinois, to pursue research with the $D\bar{D}$ experiment on the Tevatron accelerator at Fermilab. Based on his analysis of the data collected by $D\bar{D}$, he has graduated with a Ph. D. in physics from the University of Washington in 2010. He looks forward to continuing research in the field of high-energy particle physics as a member of the ATLAS experiment on the Large Hadron Collider at CERN.