UNIVERSITY OF CALIFORNIA
Santa Barbara

The Top Quark Mass, Systematic Limitations, and my Tracker-Driven Measurements

A Dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy

in

Physics

by

Ford Garberson

Committee in Charge:

Professor Joe Incandela, Chair

Professor Claudio Campagnari

Professor David Berenstein

August 2008

The Dissertation of
Ford Garberson is approved:

Professor Claudio Campagnari

Professor David Berenstein

Professor Joe Incandela, Committee Chairperson

The Top Quark Mass, Systematic Limitations, and my Tracker-Driven Measurements

by

Ford Garberson

Acknowledgements

It is standard in a thesis dissertation to acknowledge the people who have supported you in your graduate career. This is an impossible task on a modern collider experiment. Without the army of engineers and scientists building, testing, commissioning, monitoring, and repairing hardware and electronics systems there would be no collider, no experiment, and no data to analyze. Similarly, without experts to calibrate, correct, and validate the data quality, triggers, analysis objects, physics tools, and standard software, any data analysis would be a vastly (if not impossibly) more complicated project to complete. Many of these vital tasks are tedious and go unrecognized on the job market, but without the dedicated professionals who carry them out this field would collapse. There are probably thousands of people who have been instrumental to my graduation who's names I do not even know. There is no way I can express my gratitude to them.

On a more personal note, I feel honored to have been a part of UC Santa Barbara's high energy physics group. It has been a pleasure to have worked some of the most energetic, creative, and thorough scientists in the field. There are a number of people in particular who I must thank: Joe Incandela, for turning me loose and supporting me on wide variety of physics and hardware projects in three different parts of the world and on two different experiments; Chris Hill, for

pushing me through my first analysis trial by fire and showing me how to stand on my own two feet; Tony Affolder, for teaching me everything I know about silicon and coming to my rescue when my advancement presentation almost fell flat on its face; Antonio Boveia, Ben Brau, and Adam Scott for giving me a glimpse into their unique approaches to breaking down and understanding complex data, their patience in teaching me how to code under cdfsoft, and their stimulating lunchtime conversations. Of course, there are also many people from outside my group to whom I am indebted. I especially want to thank Chris Neu for his patience in supervising me and teaching me about b-tagging, Tom Wright for demonstrating how an honest and thorough scientist deals with hard problems, and Charles Plager for providing me with so many tools that have made my code less painful to write.

Abstract

The Top Quark Mass, Systematic Limitations, and my Tracker-Driven Measurements

Ford Garberson

Top quark mass measurements have achieved an unexpected level of accuracy in the last several years. This accuracy is only possible because of a new procedure that calibrates away the dominant jet energy uncertainty of past measurements. In this thesis I present some studies illustrating my suspicions that this procedure is leading them to claim overly optimistic results. Additionally, I present three measurements of the top quark mass that will be almost entirely independent of jet energies, and will thus serve as important cross checks of the standard measurements once enough statistics have been collected.

I perform my measurements of the top quark mass in the lepton plus jets channel with approximately 1.9 fb^{-1} of integrated luminosity collected with the CDF detector using quantities with minimal dependence on the jet energies. One measurement exploits the transverse decay length of b-tagged jets to determine a top quark mass of $166.9^{+9.5}_{-8.5}$ (stat)$\pm$2.9 (syst) GeV/c^2, and another the transverse momentum of electrons and muons from Wdecays to determine a top quark mass of $173.5^{+8.8}_{-8.9}$ (stat) $\pm$ 3.8 (syst) GeV/c^2. I combine these quantities in a

third, simultaneous mass measurement to determine a top quark mass of 170.7 ± 6.3 (stat) $\pm$ 2.6 (syst) GeV/c^2.

Contents

Chapter 1

The Big Picture

My thesis project is a measurement of the top quark mass. This is a very specialized topic, so before I jump into it, I want begin with a look at the broader context. I this chapter I will start in a very non-technical way with a discussion of the history of the field of particle physics, and talk about what this science is seeking to accomplish. I will go on to explain the current state of our understanding of particle physics and present the Standard Model. As an experimentalist, I almost never find myself thinking about the details of any theory. Nevertheless, since I was required to spend almost a year studying them in my physics classes, I feel that I should try my best to give a rough overview of the mathematical framework of the Standard Model. This part of the chapter will be highly technical. I will go on to discuss the unanswered questions that

will motivate future research, and explain how my thesis analysis fits into the overall picture.

1.1 What is particle physics?

Throughout the history of civilization, humans have stared up at the stars and at the world around them and wondered what everything is made of and how it all works. I would say that such questions were the domain of philosophy until the chemists of the 1700s began making real advancements in understanding the properties of materials. The first experimental evidence that matter is made up of tiny particles did not come until 1785 when Jan Ingenhousz looked at coal dust floating on a surface of alcohol under a microscope and observed that they moved around in a very jittery way. It then took well over a century (and the rediscovery of the phenomenon by Robert Brown) for the physics community to recognize that this "Brownian Motion" is caused by discrete molecules bouncing off of the dust.

Of course, water molecules are not "elementary" particles, as they can be broken into smaller pieces. The first truly elementary particle (as far as we know) to be discovered was the electron in 1897 by J. J. Thompson. He produced a current in a cathode ray tube and passed it through external magnetic and elec-

tric fields. By observing the deflection of the current, he was able to determine that it was composed of negatively charged particles, and calculate their velocity and charge to mass ratio. It was soon determined that the particles was about a thousand times lighter than the lightest known atom. Thompson proposed that atoms were made of these electrons encapsulated in a soup of positive charge to make a full atom charge neutral (the "plum pudding" model).

This idea was dramatically overturned in 1909 by Rutherford's scattering experiment. In this experiment Rutherford's group shot (positively charged) alpha particles at gold foil that was only a few atoms thick and observed that a tiny fraction of them were bounced off of the foil at more than a ninety degree angle. Since it was inconceivable that a diffuse "pudding" of charge could cause such a dramatic recoil, this proved that there were tiny clumps of positive charge distributed within the atoms, and so the atomic nucleus was discovered. Future experiments determined that alpha particles could shatter these nuclei into smaller charged pieces that could potentially have a charge as small as that of a hydrogen nucleus. In this way the proton was discovered. But a simple model of atoms as being composed of an equal number of negatively charged electrons and positively charged protons was defeated by the fact that atomic masses were heavier than this model would permit for a given number of electrons. This mystery was given partial explanation with the discovery in the 1930s of highly

penetrating neutral radiation being produced from these collisions that became attributed to neutrons.

These discoveries of new particles constitute only the tip of the iceberg, but it would serve no purpose to continue listing them here. What I want to emphasize is that particle physics is about more than the cataloging of particles. It is really about understanding how the universe works on a fundamental level, and studying particles is one of the most effective ways to do this. For example, by the year 1900 there were two fundamental forces that were thought to govern all interactions in the universe: gravity and electromagnetism (although a true understanding of the link between electricity and magnetism would require the development of Special Relativity by Einstein). This simple picture is helpless to explain the handful of experimental results mentioned above. How could electrons form a negatively charged cloud outside of a positively charged nucleus? What prevents them from collapsing into the nucleus? A full theory of quantum mechanics was needed to explain this mystery. And what keeps dozens of positively charged protons clustered together in a tiny nucleus? You would think that their electric repulsion would cause them to fly apart. This counterintuitive fact pointed to the existence of a third fundamental force that overpowered the electrical repulsion of the protons on the short ranges of the nucleus. This force was (not so creatively) called the strong force. Due to the short range nature

of this force, however, large atoms are quite difficult to hold together, which is why very large nuclei undergo fission and are radioactive, and why no elements heavier than uranium have survived in nature.

But even quantum mechanics and the strong force were not enough to explain some of the bizarre observations that physicists were making in those days. In the 1930s physicists were studying the decays of heavy nuclei. Sometimes a neutron in one of these nuclei will decay into a proton and an electron, leaving behind a positively charged nucleus rebounding against an electron in a process known as beta decay. But the amazing thing about this decay was that the nucleus and the electron did not have the proper velocities: momentum was not conserved! Some physicists were so baffled that they proposed that conservation of energy and momentum should be thrown out, but the true explanation was that a new neutral particle was being created called the neutrino that was escaping completely undetected. In 1956 this neutrino was directly discovered when it was demonstrated that beta decay can work in reverse: you can shoot a beam of atoms that beta decay to create a beam of neutrinos, and some of these neutrinos may interact with protons in target atoms producing neutrons and positrons (the antimatter twin of electrons). These neutrinos interact through a fourth type of force that was (even less creatively) called the weak force.

As if our understanding of the universe had not already been challenged enough, in the 1940s and early 1950s a wide variety of new particles were observed in the debris from cosmic rays impacting the atmosphere. These unexpected particles were quite heavy and had a bafflingly long lifetime, so they were called "strange". Much like the enormous periodic table of chemicals with predictable properties pointed to a simpler underlying structure, so too the suspicion arose that all these new particles had some underlying structure as well. The quark theory was proposed: protons and neutrons are composed of two kinds of quarks called up and down with electrical charge equal to some fraction of the electron, and the new particles from cosmic rays had their bizarre properties due to the presence of a new, "strange" quark. Scattering of electrons off of protons gave tantalizing suggestions of the presence of these quarks, but it required the development of modern particle colliders to make it possible to study these new particles in a controlled environment to make much headway. In 1974 the J/Ψ particle was discovered at these machines, which was the first particle to contain a new, fourth flavor of quark, the charm quark. When this and a variety of other new particles turned out to have exactly the properties predicted by the quark model we finally had confirmation of this theory. Unfortunately, this was the last truly dramatic and unexpected discovery that has ever been made in particle physics.

Before long, the theorists had turned the wealth of counterintuitive experimental data into a well packaged model called the Standard Model that has successfully predicted all further observations at particle colliders to date. It predicts the existence of the 17 fundamental particles with the properties described in Table 1.1. Other properties such as the typical mass and the lifetimes of these particles are shown graphically in Figure 1.1. The lifetimes depend on a number of physical principles and are explained by the Standard Model and the particle masses. But the masses themselves are completely arbitrary as far as we know. The arbitrariness of the masses and the patterns that emerge from this table suggest that there may be still further underlying physics that have not yet been discovered, just as the properties of the periodic table are explained by atomic theory, the properties of atomic nuclei are explained by protons and neutrons, and the properties of protons and neutrons are explained by the Standard Model particles. Many new theories have been developed aimed at explaining these properties.

To put it another way, from the 1930s through the 1970s the particle physics community faced the problem of a wealth of confusing data with no solid theories to explain it. Today we face the reverse problem. We have many theories that are consistent with all observable data, and we can only wait for future, highly ambitious experiments to find out which ones are correct.

Table 1.1: The particles of the Standard Model. The first category are the quarks which combine in twos and threes to make up composite hadrons such as protons and neutrons. Electrons and their heavier cousins, the muons and taus form the leptons, along with their weak force partners, the neutrinos. Each force except for gravity has an associated boson force carrier that is well described by the Standard Model. Finally, the only particle predicted by the Standard Model that has not yet been discovered is the Higgs boson, which is a consequence of Electroweak Symmetry breaking (Section 1.3.3).

Type	Mass Up →	Charge (e^- mult)	Spin
Quarks	$u \quad c \quad t$	$2/3$	$1/2$
	$d \quad s \quad b$	$-1/3$	$1/2$
Leptons	$e \quad \mu \quad \tau$	-1	$1/2$
	$\nu_e \quad \nu_\mu \quad \nu_\tau$	0	$1/2$
Force Carriers	γ (E&M)	0	1
	g (Strong)	0	1
	$W^\pm$ (Weak)	± 1	1
	Z (Weak)	0	1
Electroweak Unification	Higgs	0	0

1.2 The Evolution of Particle Physics Models

Before tackling the Standard Model, I must first explain in more detail the theories that have developed from the last century of discoveries. Perhaps the most dramatic development in this time was the theory of Quantum Mechanics. According to previous physical understanding, the particles making up the universe are like billiard balls. If you know the position and velocity of every particle at some time, you can in principle predict the position and velocity of every particle at a later or earlier time with an accuracy that is only limited by your ability to grind through the equations to get a solution.

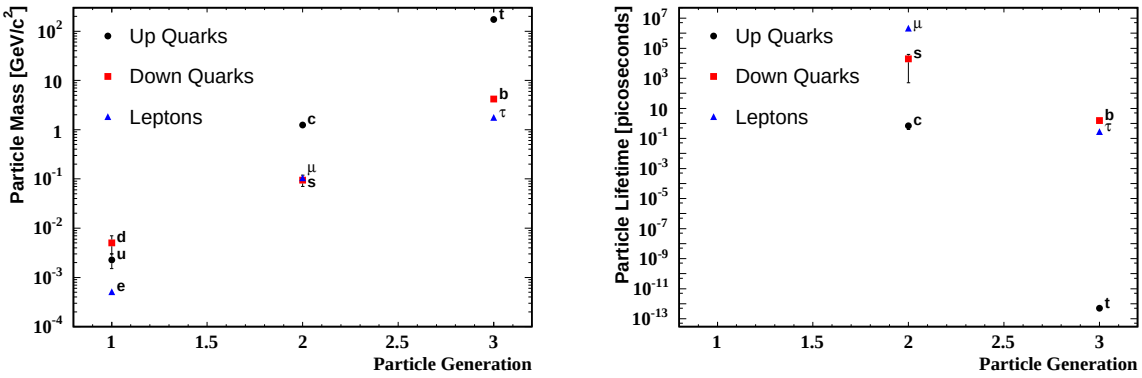

(a) (b)

Figure 1.1: Masses (left) and typical lifetimes (right) for the quarks and leptons of the Standard Model. Neutrino masses are ignored, and lifetimes are in the particle's center of mass frame. For s, c, and b quarks and average of their typical hadron lifetimes are used to determine a rough central value and range for a typical quark. Current experimental techniques are grossly inadequate to directly measure the lifetime of a top quark, so this result is taken from theory.

This neat picture was overthrown with the development of Quantum Mechanics in the early twentieth century. For very small objects, one cannot regard a particle as having a definite position and a definite velocity. All you can do is determine probabilities that if you look for a particle with a given position/velocity you will find one. The picture of a swarm of billiard balls interacting in well understood ways with one another is replaced by a picture of probability clouds representing the likelihoods for particles to exist in specific arrangements deforming one another through their interactions.

Of course, this picture is still too simplified. A proper theory must account for special relativity. One consequence of relativity is that mass, energy, and momentum are directly linked, $E^2 = p^2c^2 + m^2c^4$. It is quite possible for certain particles to annihilate one another, reducing the mass of a system while keeping the energy and momentum constant. Similarly, massless particles can annihilate one another, producing new particles that do have mass, again keeping the energy and momentum of the system constant. Thus, one cannot really think of reality as being composed of interactions between a fixed number of probability clouds each of which represent the possible locations of a given particle. Particles have to be allowed to disappear entirely or to appear as the results of some other interaction.

The mathematical framework for handling such a scenario is called Quantum Field Theory. Instead of dealing with probability amplitudes for individual particles that evolve in time, this formalism deals with time evolving objects that in mathematics are called fields. There is one field for each possible particle, and these fields can interact with the vacuum to produce particles of the appropriate type, or with a particle to destroy it. This mathematical framework is at the heart of the Standard Model of particle physics, which is the most effective theory of reality that we have so far.

The details of how Quantum Field Theory is used to predict interactions under a theory are quite complicated, and would distract from the thrust of this thesis. The short explanation is that for a given theory, a Lagrangian density is hypothesized in terms of the quantum fields. The Lagrangian is required to be invariant under local gauge transformations of the fields, and this invariance is what produces the interactions of the theory. The Lagrangian can then be used to calculate the likelihoods for any physical process to occur, usually by developing Feynman Rules to be applied under a given order of approximation. For interested readers, I have provided a highly simplified, qualitative example of how such a calculation can be performed In Appendix A. Like many other theories, the Standard Model has a particular Lagrangian, which I will introduce in the next section.

I would recommend [42] as a solid textbook for anyone who is interested in diving into Quantum Field Theory in its full glory. For readers who are interested in a more experimental approach to understanding the Standard Model that has much less emphasis on the formalism, [28] provides an easily readable introduction, and [29] delves into the physics a bit deeper. Finally, I think that [19] provides a good balance between an experimental and a theoretical emphasis.

1.3 The Standard Model

In the following sections I will present the Lagrangians for the various interactions that occur under the Standard Model, and try to give some motivation for them. For an example of how the probability amplitude can be calculated from the Lagrangian, see Appendix A.

1.3.1 Quantum Electrodynamics

In Appendix A I present the free particle Lagrangian for a spinless, electrically charged particle. Real electrons and quarks, however, have spin 1/2, so we must modify the Lagrangian found in Equation A.5 to the following:

$$\mathcal{L}^0_{EM} = i\bar{\psi}\gamma^\mu\partial_\mu\psi - m\bar{\psi}\psi \tag{1.1}$$

Here, γ^μ are the four Dirac matrices, and ψ is the quantum field for a charged, spin 1/2 fermion which must satisfy the Dirac Equation, $(\gamma_\mu p^\mu - m)\psi = 0$, and I am suppressing the hat icons on the operators. We then apply local gauge invariance under $\psi(x) \rightarrow e^{i\alpha(x)}\psi(x)$, which requires the replacement $\partial_\mu \rightarrow \partial_\mu - ieA_\mu$, where the photon field is also invariant under $A_\mu \rightarrow A_\mu + \frac{1}{e}\partial_\mu\alpha$. This results in a the full electromagnetic Lagrangian listed below.

$$\mathcal{L}_{EM} = \bar{\psi}(i\gamma^\mu \partial_\mu - m)\psi + e\bar{\psi}\gamma^\mu A_\mu \psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} \tag{1.2}$$

One needs to bear in mind, however, that this is still an approximate Lagrangian that ignores the connection between electromagnetism and the weak force. A more correct Lagrangian will be introduced in Section 1.3.2.

1.3.2 Electroweak Unification

The weak force encompasses both electrically charged and neutral interactions. A charged boson carrying the weak force will necessarily change the particle it interacts with into a different particle in a manner that conserves electrical charge. The appropriate charged weak probability currents for electrons have the form:

$$\begin{aligned} J^+_\mu &= \bar{\nu}\gamma_\mu \frac{1}{2}(1-\gamma^5)e = \bar{\nu}_L \gamma_\mu e_L \\ J^-_\mu &= \bar{e}_L \gamma_\mu \nu_L \end{aligned} \tag{1.3}$$

Here the ν's represent the quantum fields for electron-neutrinos, and the $e's$ represent the electron quantum fields. The factor $\frac{1}{2}(1-\gamma^5)$ has the effect of picking out the left handed helicity final state (in the limit of zero electron mass), and so the equations are often written in terms with the subscript L's as

shown above. These two equations turn out to represent interactions from two components of an SU(2) group where the full interactions are described by the equations:

$$J^i_\mu = \bar{\chi}_L \gamma_\mu \frac{1}{2} \tau_i \chi_L, \qquad \chi_L = \begin{pmatrix} \nu_L \\ \\ e_L \end{pmatrix} \tag{1.4}$$

Here the τ's represent the standard Pauli spin matrices. The charged currents of Equation 1.3 are the appropriate ones associated with the raising and lowering operators in this spin space. That is, $J^\pm_\mu = \bar{\chi_L} \gamma_\mu \tau_\pm \chi_L$, where $\tau_\pm = \frac{1}{2}(\tau_1 \pm i\tau_2)$. While these first two components of the current are associated with charged interactions, the third component is written in terms of the diagonal τ_3 matrix, and so the particle charge remains unchanged. Except for the restriction to left handed states, the characteristics of this interaction are very similar to those of the electromagnetic interaction, and so it is natural to combine them into a neutral electroweak current:

$$j^Y_\mu = 2(j^{em}_\mu - J^3_\mu), \tag{1.5}$$

$$j^{em}_\mu = -\bar{e}\gamma_\mu e = -\bar{e}_R \gamma_\mu e_R - \bar{e}_L \gamma_\mu e_L \tag{1.6}$$

When developing the full theory, the basic electromagnetic interaction, $-iej^\mu_{EM}$ (associated with Equation A.8), is replaced with:

$$- i\frac{e}{\sin\theta_W} J^{i\mu} W^i_\mu - i\frac{e}{2\cos\theta_W} j^{Y\mu} B_\mu \tag{1.7}$$

Where θ_W is the electroweak mixing angle (a fundamental parameter of the Standard Model). The W-bosons correspond to the fields $W^\pm_\mu = \sqrt{\frac{1}{2}}(W^1_\mu \mp i W^2_\mu)$. Similarly, the photon and Z-boson correspond to the fields:

$$\begin{aligned} A_\mu &= B_\mu \cos\theta_W + W^3_\mu \sin\theta_W \\ Z_\mu &= -B_\mu \sin\theta_W + W^3_\mu \cos\theta_W \end{aligned} \tag{1.8}$$

The corresponding Lagrangian (still neglecting all fermions except for the electron-neutrino pair) is then:

$$\begin{aligned} \mathcal{L}_{EW} &= \bar{\chi}_L \gamma^\mu [i\partial_\mu - \frac{e}{2\sin\theta_W}\tau^i W^i_\mu + \frac{e}{2\cos\theta_W} B_\mu]\chi_L \\ &+ \bar{e}_R \gamma^\mu [i\partial_\mu + \frac{e}{\cos\theta_W} B_\mu] e_R - \frac{1}{4} W^i_{\mu\nu} W^{i\mu\nu} - \frac{1}{4} B_{\mu\nu} B^{\mu\nu} \end{aligned} \tag{1.9}$$

Here the kinetic terms for the bosons are given by $B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu$, and $W^i_{\mu\nu} = \partial_\mu W^i_\nu - \partial_\nu W^i_\mu - \frac{e}{\sin\theta_W}(W_\mu \times W_\nu)^i$. This Lagrangian is, however, missing mass terms for the W and Z bosons. The difficulty is that mass terms, e.g., $M^2_W W_\mu W^\mu$, break local gauge invariance, and lead to a non-renormalizeable theory if they are added in directly by hand. The solution to this dilemma comes from introducing the Higgs Field.

1.3.3 The Higgs Mechanism

The Higgs Mechanism is introduced in the form of four scalar fields:

$$\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi_1 + i\phi_2 \\ \\ \phi_3 + i\phi_4 \end{pmatrix} \tag{1.10}$$

These fields show up in the electroweak Lagrangian in the following form:

$$\mathcal{L}_{EW}^{Higgs} = |(i\partial_\mu - i\frac{e}{\sin\theta_W}\tau^i W_\mu^i - i\frac{e}{2\cos\theta_W}B_\mu)\phi|^2 - V(\phi) \tag{1.11}$$

Here, the vacuum potential is $V(\phi) = \mu^2\phi^\dagger\phi + \lambda(\phi^\dagger\phi)^2$ at lowest order. The electroweak gauge boson fields were needed to make this Lagrangian locally gauge invariant under the SU(2) phase rotations $\phi \rightarrow e^{i\theta_a(x)\tau_a/2}\phi$. The three degrees of gauge freedom, $\theta_a(x)$ allow us to rotate away three out of the four Higgs fields in Equation 1.10. The Higgs SU(2) doublet becomes:

$$\phi(x) \rightarrow \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ \\ v + h(x) \end{pmatrix} \tag{1.12}$$

Only one Higgs Field, $h(x)$ remains, and a constant v. In oder to determine what v should be, we note that for the perturbative Feynman rules to be effective, the fields need to represent fluctuations about a minimum in the vacuum

potential. For the potential of Equation 1.11, the minimum of this potential will occur when $v = \sqrt{-\frac{\mu^2}{\lambda}}$ in Equation 1.12.

With this definition of ϕ, the expansion of Equation 1.11 has the term $(\frac{ev}{2\sin\theta_W})^2 W_\mu^+ W^{-\mu}$ in it, predicting a leading order W boson mass of $m_W = \frac{ev}{2\sin\theta_W}$. A Z mass can similarly be evaluated, but the coefficient of the $A_\mu A^\mu$ term evaluates to zero, indicating that photons are massless as expected. An almost identical procedure can also be applied to create masses for the fermions. For example, the necessary Lagrangian term to produce a mass for the electron is given by:

$$\mathcal{L}_{ele}^{Higgs} = -G_e[\bar{\chi}_L^e \phi e_R + \bar{e}_R \phi^\dagger \chi_L^e] \tag{1.13}$$

The electron mass will be given by $m_e = (1/\sqrt{2})G_e v$, and is determined by the arbitrary choice of the coupling constant, G_e. The terms for the quark masses are more complicated due to the fact that they contain both left and right handed helicity states, and the fact that their flavor eigenstates are not the same as their mass eigenstates (leading to the topic of the CKM matrix which I will not cover here). But neither of these complications are showstoppers, and the same Higgs field can be designed to be responsible for the masses of all elementary particles in a gauge invariant manner.

A dramatic side effect of these Lagrangians is the introduction of a scalar Higgs field with its own mass term given by $m_H = \sqrt{2v^2\lambda}$ at lowest order. The devilish part about this physical particle, however, is that it is very challenging to observe experimentally. As shown above, it couples to other particles in proportion to their mass. So particles such as electrons that are easy to create have an essentially zero probability of producing a Higgs, while massive particles such as the top quark have a high probability (relatively speaking) of producing a Higgs, but are very challenging to produce themselves. Probably because of this complication, the Higgs boson is the final predicted Standard Model particle which has not yet been experimentally discovered.

1.3.4 Quantum Chromodynamics

Quantum Chromodynamics refers to interactions between quarks and gluons involving the strong force. In addition to electric charge, quarks carry a strong force quantum number that is called color and that is allowed to take on three unique values. Quarks can never be found as lone particles. They must either be bound together with an antiquark of the same color, or with two other quarks each of different colors into a final compound hadronic state that is "color neutral". The interactions that create these color neutral hadrons are known as fragmentation, and occur whenever a quark or a gluon is produced in an interac-

tion. At a collider, the hadrons created in this process are usually highly boosted, and generally traverse the detector as a narrow stream of particles called a jet. The free quark in color state j will have an associated Lagrangian:

$$\mathcal{L}^0_{QCD} = \bar{q}_j(i\gamma^\mu\partial_\mu - m)q_j \tag{1.14}$$

We must generalize our local gauge invariance requirement to account for this color:

$$q(x) \rightarrow e^{i\alpha_a(x)T_a}q(x) \tag{1.15}$$

Here, T_a can be any set of eight linearly independent traceless 3x3 matrices. As matrices, these terms clearly do not commute, and will obey a particular set of commutation relations, $[T_a, T_b] = if_{abc}T_c$. This invariance requires the replacement of the derivative by

$$\partial_\mu \rightarrow \partial_\mu + igT_aG^a_\mu \tag{1.16}$$

Eight new quantum fields have been created which represent the gluons in all their possible color neutral forms. Like photons, gluons will have a (more complicated) kinetic energy term involving their stress tensor, $G^a_{\mu\nu} = \partial_\mu G^a_\nu - \partial_\nu G^a_\mu - gf_{abc}G^b_\mu G^c_\nu$. The final QCD Lagrangian is then:

$$\mathcal{L}_{QCD} = \bar{q}_j(i\gamma^\mu\partial_\mu - m)q_j - g(\bar{q}_j\gamma^\mu T_a q_j)G^a_\mu - \frac{1}{4}G^a_{\mu\nu}G^{\mu\nu}_a \tag{1.17}$$

Where a sum over the three color indices, j, is implied. In addition to the standard fermion-fermion-boson interactions from QED, the last term in this Lagrangian also predicts the self coupling of gluons. It is interesting to note that these couplings are expected to allow the existence of stable three and four boson final states called glueballs. These glueballs have yet to be discovered experimentally, though there have been some tantalizing hints.

1.3.5 Putting it all together

Combining the QCD, Electroweak, and Higgs Lagrangians above, and summing over all quark flavors, lepton flavors, and QCD colors, we arrive at the full Standard Model Lagrangian:

$$\mathcal{L}_{SM} = \sum_{j}^{N_{LH}} \bar{L}_j \gamma^\mu (i\partial_\mu - \frac{e}{2\sin\theta_W}\vec{\tau}\cdot\vec{W}_\mu - \frac{eY_j}{2\cos\theta_W}B_\mu)L_j +$$
$$\sum_{k}^{N_{RH}} \bar{R}_k \gamma^\mu (i\partial_\mu - Y_k\frac{e}{2\cos\theta_W}B_\mu)R_k - \frac{1}{4}\vec{W}_{\mu\nu}\cdot\vec{W}^{\mu\nu} - \frac{1}{4}B_{\mu\nu}B^{\mu\nu} -$$
$$\sum_{l}^{N_{LH}} (G_l^{lepton}\bar{L}_l\phi R_l + G_l^{quark}\bar{L}_l(-i\tau_2\phi^*)R_l + h.c.) +$$
$$|(i\partial_\mu - \frac{e}{2\sin\theta_W}\vec{\tau}\cdot\vec{W}_\mu - \frac{e}{2cos\theta_W}B_\mu)\phi|^2 - V(\phi) +$$
$$\sum_{mna}^{N_c,N_q,N_g} [\bar{q}_{mn}(i\gamma^\mu\partial_\mu - m_n)q_{mn} - g(\bar{q}_{mn}\gamma^\mu T_a q_{mn})G_\mu^a] - \frac{1}{4}G_{\mu\nu}^a G_a^{\mu\nu} \quad (1.18)$$

Here, the first line represents the couplings between the electroweak bosons and the $N_{LH} = 6$ left handed fermion states with weak hypercharge Y. The second line covers the neutral electroweak couplings with the right handed fermion states and the kinetic energies of the electroweak bosons. The third line provides the couplings between the fermions and the Higgs field, along with the fermionic mass terms. The fourth line gives us the same for the massive electroweak bosons and the Higgs potential, which in the leading order approximation is $V(\phi) = \mu^2\phi^\dagger\phi + \lambda(\phi^\dagger\phi)^2$. Finally, the fifth line gives us the QCD contributions: the gluon self couplings and kinetic energy, and the couplings between quarks and gluons summed over the number of color states, $N_c = 3$ for quarks and $N_g = 8$ for gluons, and the number of quark states, $N_q = 6$.

1.3.6 Standard Model Limitations

Although the Standard Model has successfully explained all collider data to date, the theory is incomplete, and has logical inconsistencies. If it were not then there would be much less interest in studying particle physics! Some of the limitations of the Standard Model are listed below.

- Gravity. The Standard Model incorporates only the electroweak and strong forces, but it is very challenging to develop a quantum theory of gravity. This deficiency has not posed a problem so far because gravity is so weak compared to the other forces that it can only play an appreciable role on *very* short distance scales. The collider energies needed to probe such tiny scales are on the order of the Planck Mass, about 10^{14} times larger than those that will be achieved at the LHC. This number, however, assumes that gravity scales in strength following the same trends that we observe in the macroscopic world around us. For all we know gravity may become powerful much faster than expected at small distance scales and start to play a dominant role. Some theories involving extra dimensions, for example, make such a prediction.

- Cosmological inconsitencies. There are a number of cosmological mysteries that contradict the Standard Model. For example, if there were no CP vi-

olation, all matter would have been expected to annihilate with antimatter in the early universe after the Big Bang. The fact that there is matter and we exist provides evidence of a sizable CP violation in the fundamental laws of physics. The Standard Model allows for CP violation, however experimental measurements at colliders have restricted the magnitude of this violation to a value that is too small to explain the required imbalance. As another example, the Standard Model cannot explain the observed dark matter in the universe. The most likely explanation for dark matter is the presence of a new type of a stable particle that does not interact through the electromagnetic or strong forces, and makes up on order of 80% of the mass of the universe. There are high hopes that such a particle may be discovered at the LHC.

- The "Hierarchy Problem". As described in Section 1.3.3 , the Higgs mass can be calculated at lowest order to be $m_H = \sqrt{2v^2\lambda}$. But when the higher order corrections involving loops from other Standard Model particles are included, unless the non-radiative Higgs mass has exactly the right value down to incredible precision, then the calculated Higgs mass blows up to an infinite value. Unless the non-radiative mass is incredibly fine tuned, new physics must come into play to reduce the corrected mass. Indirect

electroweak limits [1] suggest that the Standard Model Higgs mass is less than that of the top quark, suggesting that the new physics is likely to appear at energies accessible by the LHC.

- The "Arbitrariness Problem". The Standard Model may make successful predictions, but there is no motivation for its structure. Why should we expect three generations of particles? Why should they be arranged in the $SU(3) \times SU(2) \times U(1)$ group structure? In its simplest form there are 19 free parameters in the Standard Model, all of which must be determined by experiment. Even small changes in these parameters would have dramatic consequences, as creatively explained in [24]. For example, if the electron mass were not so tiny then muons would convert all protons to neutrons and neutrinos. The universe would degenerate into a lightless assortment of neutron stars, black holes, and empty voids.

1.4 The Top Mass and the Higgs

Now that I have presented the Standard Model and explained its successes and failures, it is reasonable to ask how my thesis project will help to advance our understanding of particle physics. As mentioned above, the top mass is only one of 19 free parameters in the Standard Model. Measuring it precisely does

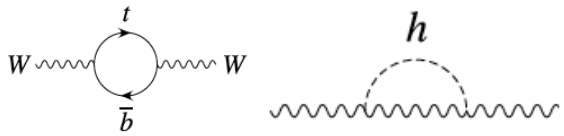

(a) (b)

Figure 1.2: Feynman diagrams demonstrating some of the couplings that alter the mass of the Wboson. a: a correction according to the top mass. b: a correction according to the Higgs mass.

not directly help to solve any of the mysteries laid out in Section 1.3.6. However, the top quark mass can be used to predict the mass of the Higgs, assuming the particle exists. Such a prediction can be used to guide searches for the Higgs, and if the Higgs is discovered, it will provide a direct test of the validity of the Standard Model.

The Wmass determined in Section 1.3.3 is, of course, only a leading order approximation of the value that is measured in experiments. Corrections are needed to account for loops from other particles within the W propagator such as those shown in Figure 1.2. A more complete formula for the Wmass is [32]:

$$m_W^2 = \frac{\pi \alpha_{em}}{\sqrt{2} G_F \sin^2 \theta_W (1 - \Delta r)} \tag{1.19}$$

Here, Δr represents the correction from higher order loops, and will contain mass terms for all particles that contribute to such loops. Of these particles, the most important ones are the top quark and the Higgs (if it exists). After doing

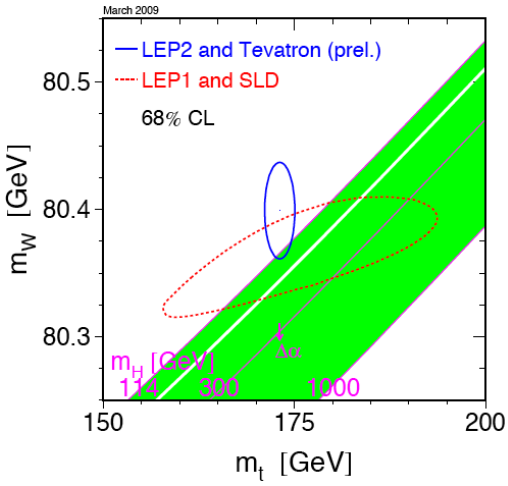

Figure 1.3: Standard Model Higgs mass constraints based upon the most recent top and W mass measuremnts

the appropriate calculations, one can solve for the Higgs mass. The uncertainties on the Higgs mass will then be driven by uncertainties on the top and W mass measurements as well as uncertainties on the measurements that determine $\sin\theta_W$ and the couplings. Periodic fits are performed to the measurements of various Standard Model parameters such as these. When the latest results [1] are projected onto the m_t vs m_W plane, the one-standard deviation error ellipse is shown in Figure 1.3, it is evident that a light Higgs is preferred. Taking all uncertainties into account, the Higgs mass is constrained to be $< 163\ GeV/c^2$ at a 95% confidence level.

But this motivation for my thesis is deceptive. As should be obvious from Figure 1.3, the uncertainty on the Higgs mass is limited far more by the W mass

uncertainty than by the top mass uncertainty. The published top mass results are now so precise that there may no longer be much benefit in improving them. However, as I will explain in Section 5.2, this resolution depends entirely on the validity of one new analysis technique, and I suspect that the uncertainties on this technique may be significantly underestimated, for reasons I will explain in Section 5.4. Since my measurements are performed in a completely different manner than other analyses, I consider my thesis to be more interesting in its capacity to provide a cross check to the conventional measurements, than in its more limited capacity to improve the world average top mass results.

Chapter 2

Particle Colliders

Physicists have had centuries to a solid understanding of the everyday world around us. The aim of particle colliders has been to push this understanding deeper by studying the behavior of matter at higher and higher energies, at which unique interactions may occur. As explained in Section 1.1, powerful accelerators have been responsible for discovering countless new particles, and for illuminating the true nature of the strong and weak forces.

2.1 Accelerators

The original accelerator design was linear. In these machines, charged particles travel through a vacuum in a straight line, precisely boosted at periodic

intervals (usually with radio waves these days) to incrementally higher and higher energies. The energies attainable at such accelerators are limited by the tuning and power of the boosters, and the length of the accelerator you were willing to build.

Competing circular accelerator designs are also in wide use. The earliest designs were known as cyclotrons. Charged particles travel around these machines, constrained to move in a spiral by a magnetic field. A high frequency alternating current gives the particle a kick each time it passes the same angle, boosting the energies higher each time while the radius of revolution increases to stay in phase with the boosts. Newer circular accelerators (synchrotrons) vary the magnetic fields and the frequencies of the boosts, allowing the particles to travel in a closed circle up to the highest possible energies. Circular colliders allow you to accelerate particles around as many revolutions as you like, with the maximum available energies limited by the maximum strength of the magnetic field (or sometimes in the case of electrons, by their energy loss rate through synchrotron radiation). Both limitations can be offset by building the colliders at a larger size so that the particles need not be forced into such tight curves.

Under a simple design the accelerated beams can be directed at a fixed collision target. Most modern accelerator experiments, however, have two accelerated beams which are aimed to collide with each other. One major advantage of these

designs is that it allows one to achieve much higher effective energies. For example, a 1 GeV beam of electrons directed at a fixed electron target corresponds to about 32 MeV in the rest frame of the collision. Having the beam collide with a second 1 GeV beam of electrons going in the opposite direction does not increase the center of mass energy by a factor of 2, it increases it by a factor of more than 30. In fact, the energy boost gained from head on collisions goes as the square of the energy of each individual beam. Thus, despite the formidable engineering challenges, all modern colliders on the energy frontier have been built to smash two beams into each other.

As shown in Figure 2.1, collider energy reaches have been increasing at an exponential rate over most of the last fifty years. Progress has slowed recently, however. This is primarily because of the limitations described above. Linear colliders must be built larger to provide greater energy reach, and circular colliders must be built larger for the magnets to be capable of constraining the particle trajectories and to reduce the synchrotron radiation losses. Larger colliders have simply become prohibitively expensive to build. It may well be that the LHC will mark the last generation of colliders. The only way that the energy frontier could ever become cheap again would be from technological breakthroughs. In fact, some promising research is being directed at developing acceleration tech-

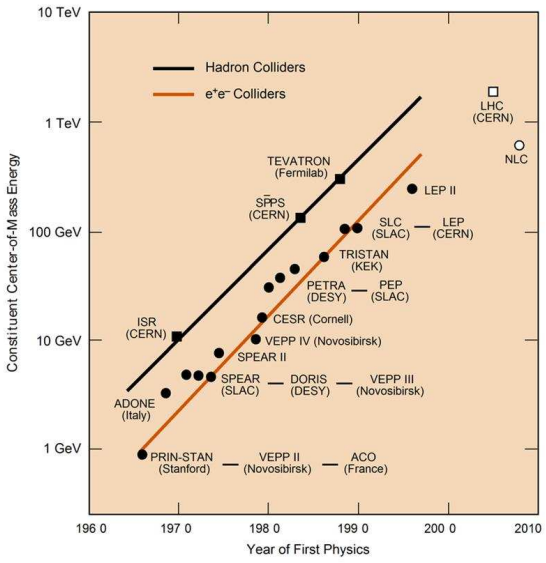

Figure 2.1: Historical evolution of the energies of colliding partons in accelerators. The predicted turn on dates of the LHC and the next linear collider are, of course, far too optimistic in this plot.

niques that may be orders of magnitude more efficient [43] [23]. Such research deserves more attention, but these technologies are still in their infancy.

2.2 The Tevatron

I carried out my thesis on the CDF experiment on the Fermilab Tevatron accelerator, shown in Figure 2.2. The Tevatron consists of a series of accelerators and storage rings, culminating in the collision of protons with antiprotons at a center of mass energy of almost 2 TeV. Though the average collision energies of the individual quarks and gluons are considerably lower than this (see Fig-

ure 2.1), the Tevatron will remain be the world's most powerful collider until the LHC turns on, and is the only collider that has ever been capable of studying top quarks.

The acceleration procedure is as follows. Protons are taken from a hydrogen bottle and bonded with two electrons to form a negatively charged atom. They are then accelerated through a crockoft-walton and linear accelerator (LINAC) to a kinetic energy of 400 MeV. These atoms are then inserted into the booster where they immediately have their electrons stripped from them by a carbon foil. Additional packets of H^- atoms can be easily added because they will not be electrically repelled. Inside the booster the proton energies are ramped to 8 GeV before they get passed on to the main injector.

Once in the injector the protons are ramped to 120 GeV and a beam is shot at a nickel target where antiprotons have a tiny chance of being created and captured. The antiprotons are then passed through a debuncher to regularize their energies before they get stored in the accumulator ring, and are periodically passed on to the larger Recycler ring. The main injector then ramps the proton and antiproton energies to 150 GeV, before they are injected to the Tevatron, where each beam is accelerated to the world's highest energies of 980 GeV. There are six locations where the proton and antiproton beams can be made to cross. The multipurpose CDF and D0 experiments have been set up at two of them.

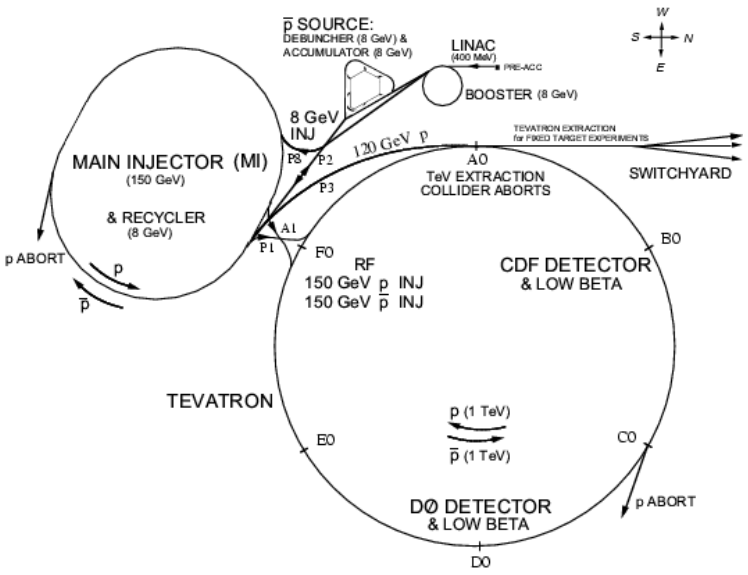

Figure 2.2: The Tevatron accelerator complex.

Chapter 3

Detection of Particles

3.1 How Particles Interact in Detectors

Before discussing the design of particle detectors it is helpful to first review the physics of how particles interact with matter. The reasons for the design will then be much clearer. I will give a brief overview of some of the common types of interactions below, and what behaviors should be expected for different classes of particles.

An electrically charged particle will interact in matter in very different ways than neutral particles. For the sorts of high energy particles produced at colliders, the sign of the charge will have almost no impact, all that matters is the charge magnitude. The particle will undergo constant collisions as its elec-

tromagnetic field interacts with the electromagnetic field of the electrons in the atoms of the material through which it passes. This will cause the particle to gradually lose energy. In the case of a heavy incoming charged particle, the average rate at which it looses energy through these interactions, dE/dx, will be a function of its mass, charge, relativistic boost, and the properties of the material through which it passes. Holding all other quantities fixed, this rate of energy loss is shown in Figure 3.1 below, taken from the Particle Data Group (PDG) [27]. The plot is complicated, but I will explain its essential features below.

Heavy charged particles in the usual energy range found in particle colliders are fairly well described by the Beth-Bloch formula, shown in Equation 3.1. This equation is a lowest order approximation[1] of the rate of energy loss the particle will suffer from ionization and atomic excitation interactions. Here, N_A represents Avagadro's number, r_e the classical radius of the electron, z the charge of the incident particle, Z the atomic number of the absorber, A its atomic mass, T_{max} the maximum kinetic energy that may be imparted to a free electron in a collision, I the average exitation energy for the absorber electrons, and δ is a correction to account for long range collisions due to relativistic stretching of the incident particle's electric field. Despite being a lowest order approximation, and despite ignoring nuclear interactions, electronic motion, and photonic radiation

[1]Strangely, the Bloch correction is a higher order term that is generally not written in the "Bethe-Bloch" equation

off of the incident particle, this equation accurately describes the data in the central area of Figure 3.1, between $0.05 < \beta\gamma < 500$. For lower energy particles, the orbital properties of electrons within the atoms of the material become important, and the sign of the charge of the incident particle starts to play a role. The higher order Barkas term [40] is needed to model these interactions. Above $\beta\gamma = 0.05$, electronic ionizations start to become more rare, and the particle loses energy more slowly until the β^2/β^2 term becomes the most significant, at which point the particle becomes "minimum ionizing". For a *very* boosted particle, the relativistically enhanced electromagnetic fields of electrons within the bulk cause the particle to accelerate enough to emit significant Bremsstrahlung radiation. This is the dramatic rise in the rate of energy loss seen at very high boosts.

$$-\frac{dE}{dx} = \frac{4\pi N_A r_e^2 m_e c^2 z^2 Z}{A\beta^2}\{\frac{1}{2}ln(\frac{2m_e c^2\beta^2\gamma^2 T_{max}}{I^2}) - \beta^2 - \frac{\delta(\beta\gamma)}{2}\} \tag{3.1}$$

High energy electrons and photons both lose their energy over a characteristic distance scale known as a radiation length (the distance over which the electron loses all but 1/e of its energy). For electrons, this can easily be understood by their low mass. They behave similarly to the muons shown in Figure 3.1, but since they are about 200 times lighter than the muons they will have correspond-

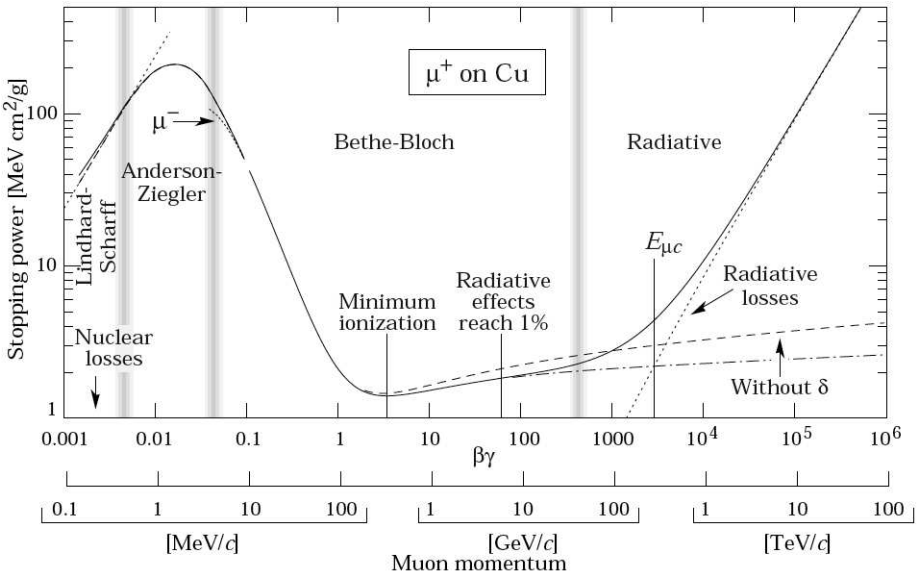

Figure 3.1: A characteristic $\frac{1}{\rho}\frac{dE}{dx}$ distribution for charged, heavy particles passing through matter (in this case antimuons passing through copper).

ingly larger boosts. Thus, they will radiate away their energy very rapidly. For high energy photons, the radiation length represents the distance over which 7/9 of them will be absorbed in the production of an electron positron pair. It should be noted that incident high energy electrons mostly lose energy by producing photons, while in turn pair produce into electrons, and so on in a characteristic cascade. Similarly, incident high energy photons quickly pair produce into electrons, and so on, in an almost identical looking cascade. The interactions of electrons and photons are thus very different from those of heavy charged particles, but are difficult to distinguish from each other.

There is also a much more complicated class of hadronic interactions. To lowest order, neutral hadrons will have no interactions with the electron clouds of the material. Charged hadrons will continuously deposit energy in the electrons according to Figure 3.1. However, since most hadrons at colliders will be minimum ionizing, they will lose energy very slowly. The primary manner in which a hadron stops in matter is when it interats with the nucleii of the atoms in the bulk through the strong force. This interaction will often be very destructive, and can cause the hadron and nucleus with which it collides to shatter in a spray of other particles. These reactions are much more complicated and are very difficult to understand.

The final class of particles are those which interact neither through the electromagnetic force nor the strong force. As far as we know, the only long-lived particles that behave this way are neutrinos. The weak force through which neutrinos interact is so feeble that neutrinos are virtually guaranteed to pass through any collider detector without depositing a detectable amount of energy.

3.2 Tracking detectors

Tracking detectors all share some common characteristics. A good tracking detector will extract enough energy from a charged particle passing through it

to reliably read out a trajectory, but not so much that it deflects the particle or significantly reduces its overall energy. In addition to the trajectories, many experiments also set up a magnetic field throughout the tracking system. The direction in which the particles bend under the influence of the field determines their charge, and the amount of curvature determines the particle's momentum. If the field is constant and parallel to the beam axis, the relationship is quite simple: $p_T = KsB$, where s is the radius of curvature, B is the magnitude of the field, and K is the appropriate constant depending upon your units. This is the only reliable way to determine the momentum of a muon in your detector.

The cloud chamber is perhaps the oldest type of tracker. Invented in 1911 by Charles Wilson, it consisted of a container of a supersaturated gas. The vapor was expanded adiabatically by a diaphragm until it reached a very low temperature, but the chamber contents would remain gaseous. However, when a charged particle passed through the tank it would ionize the vapor. The ions would act as nuclei around which condensation could form. The trail of condensation would then clearly describe the trajectory of a track.

The next major advancement in tracking technology occurred in 1952 with the invention of the bubble chamber by Donald Glaser. Bubble chambers are similar to cloud chambers except that they are filled with a liquid instead of a

gas[2], and instead of being cooled, the liquid is kept at a very high temperature. When a charged particle passed through the liquid the resulting ions would send the nearby liquid over the boiling point resulting in microscopic bubbles. At regular intervals a piston would be released allowing the liquid to expand until the bubbles became visible and the track could be seen. The bubble chamber could be made larger than the cloud chambers, allowing for a better momentum resolution of charged particles in the presence of a magnetic field, and the liquid could be made much denser, allowing for the detection of more energetic particles which might otherwise pass through undetected if they are in the minimum ionization range of Figure 3.1.

Both cloud chambers and bubble chambers share a common drawback, however. They both operate visually, so the only way to save records of track trajectories was to constantly snap photographs and store the pictures to be scanned later by bored graduate students. For high luminosities it quickly becomes impossible to have a person manually scan the millions or trillions of images that would be needed. In the modern era all collider tracking detectors operate by extracting the charge from the ions the incoming particles produce in the material. The charge is then amplified to a detectable level, and saved to disk so that

[2]Supposedly Dr. Glaser used beer as the liquid in his early prototypes

computers can reconstruct the tracks with high accuracy, and run automated algorithms to decide which events are interesting.

One modern tracker design uses the principles of scintillation. Briefly, when a charged particle passes through a scintillator it leaves behind a wake of excited molecules. The material of the detector is chosen so that when the molecules decay to their ground state they will emit a photon that can easily be detected (usually with a wavelength near the optical range). The CDF outer tracker operates on an entirely different principle, using what is known as a wire drift chamber design. It consists of eight concentric cylinders of wires running through a gas mixture which ionizes when a charged particle passes through it. A steady electric field throughout the drift chamber causes the ions to collect in the wires and the charge is then read out. One might expect the hit resolution to be limited by the spacing between the wires. However, by measuring the time at which the peak of the signal pulse is read out, one can indirectly determine how long it took the charge to drift through the gas mixture to each wire, and thus how far from the wire the charged particle passed. In the end, a hit resolution on the order of 150 microns is achieved.

3.2.1 Silicon tracking detectors

Most of the newest generation of tracking detectors are silicon based. The whole of CDF's inner tracking is made of silicon detectors. For my first project with the Santa Barbara group I spent time running electronic tests on individual components of CMS's silicon detector called modules. For this reason I will describe silicon trackers in a bit more detail and explain my involvement. Far more complete details on the properties of silicon detectors can be found in [30].

Semiconductor detectors, such as silicon, have a number of key advantages compared to the detectors discussed previously. In particular, the required energy deposit to produce a detectable electrical signal is much smaller. In silicon, for example, the energy required to ionize an electron is only 3.6 eV (after accounting for the required phonon emission). In contrast, ionization in a gaseous detector generally requires on the order of 30 eV. As a result one can generally achieve a much cleaner signal in a silicon detector. Another advantage of silicon is that the material is cheap because of its high abundance. Silicon is also easy to shape, allowing very high resolution detectors to be constructed. Finally, these detectors have a read out time that is limited only by the speed of the front end electronics, make silicon an ideal choice for use in the very high luminosity environments of modern colliders that would quickly saturate most other types of tracking detectors.

In order to understand how a silicon detector operates one must understand how to control its noise. The disadvantage of having such a small band gap in silicon is that it is quite easy for thermal fluctuations to excite electrons to the conduction band, leading to an observed current. In fact, at room temperature the signal for a minimum ionizing particle is expected to be about 50,000 times smaller than this thermal noise. One could solve this noise problem by cryogenically cooling the entire detector to very cold temperatures. However, there is a different trick that can be applied much more easily. This trick requires us to modify the silicon by inserting impurities .

The insertion of these impurities (called "doping") is generally designed to reduce the size of the band gap in semiconductors. Inserting impurities with single, loosely bound electrons in their outer shell makes the silicon "n-type" doped, since it then becomes very easy to excite negatively charged electrons into the conduction band. On the other hand, one can insert impurities the have only a single electron missing from their outer electron level which makes the silicon "p-type" doped. Such impurities are very likely to capture electrons from the valence band of the silicon, creating a hole in the bulk which will also induce a charge flow. But the most remarkable aspect of this doping is that If one manufactures n-type and p-type doped silicon bulks and brings them into contact with one another it creates what is known as a pn-junction. Since the fermi level

must be constant in any material with a high conductivity, this connection will cause holes to flow from the p-bulk into the n-bulk, and electrons to flow from the n-bulk into the p-bulk, until the energy levels shift enough to make the fermi energy constant again. This creates a small depleted region where there are no charge carriers available to induce a current.

The electrical potential in the pn-junction I have described is naturally higher on the n-side than the p-side due to the charge imbalance. As a result, if one applies an external voltage to the system that is positive on the p-side an negative on the n-side, the potential difference will flatten out and the depleted region will shrink. until charge can flow freely. If, on the other hand, one applies the voltage the other way around, this potential difference will be enhanced and the depleted region will grow until it covers the entire silicon bulk. In this scenario there will still be a small leakage current as thermal electron-hole pairs are created from deeper in the bulk of the silicon and swept aside, but this current will be many orders of magnitude smaller unless the voltage is increased so high that it can initiate electrical avalanches.

With the noise well under control a functioning detector can be designed. The silicon modules for the CMS tracker, were manufactured as thin, n-doped silicon wafers, in which hundreds of narrow strips of p-doped silicon were implanted. When powered, a 400 V external voltage is applied, where the backplane of the

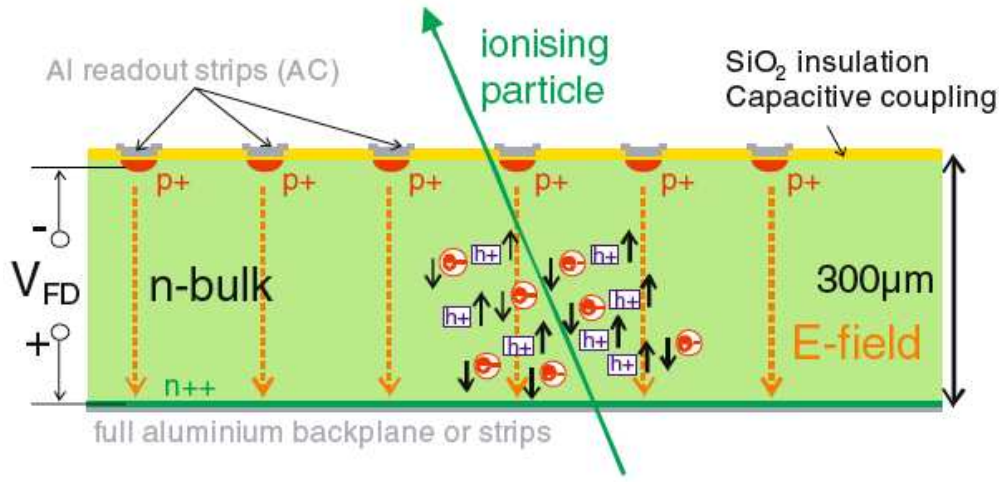

Figure 3.2: A cross section of a CMS silicon sensor. Charges are ionized in the depleted n-bulk, and flow under influence of the reverse biased voltage to either the p-implants or the backplane. Capacitive coupling induces currents in the aluminum strips, which are then amplified, shaped, and read out.

n-bulk is held at a positive potential and the strips are kept negative. This potential is more than enough to enforce complete depletion of the bulk. Then, when a charged particle enters the bulk, it will excite electrons to the conduction band that will be swept to the backplane, while the holes will be drawn to the p-doped implants. In order to protect the sensitive silicon from damage it is essential to coat the surface of the silicon with a non-conductive, non-reactive SiO_2 layer. This layer functions as a dielectric. So capacitors can be designed by layering aluminum strips on top of the SiO_2, oriented above and parallel to the p-implants. Current in the bulk will therefore lead to charge flow through capacitive coupling. The charge flows out to amplifiers which enhance the signal

to an easily detectable level. The full cross section of a sensor is shown in Figure 3.2 (image taken from [30]). At CMS a signal to noise ratio of about 40 to 1 is achieved in this manner, along with a hit resolution on the order of 10 microns.

I became involved in the CMS silicon effort by studying the completed modules for defects. I will briefly summarize my work. I increased the reverse bias voltage and studied the leakage current through the bulk to test for breakdown. I ran tests of the "pipelines" that stored data in the electronic readout systems in the chips. In addition, since this hardware will be ran in a cooled state, I thermal cycled hundreds of modules several times between temperature extremes far faster and over a greater range than they will experience under operating conditions. These cycles had the inadvertent side effect of demonstrating that the modules would not be seriously damaged by condensation, which occurred when outside air leaked into the environmental chamber (two engineers from the company we rented it from assured me that this could not happen, but they were very wrong). I ran electronic tests scanning the modules for defects before and after these cycles, and found that they survived them well.

The electronic tests I ran searched for unusual noise and current signatures that would be indicative of physical damage in the sensors. Perhaps the most common defect in the silicon module would occur as the result of a broken bond

in a wire connecting the aluminum strips to the readout. In the absence of any signal the amount of noise that is read out from a healthy channel is directly proportional to the applied voltage, and inversely proportional to the capacitance between the aluminum and p-implant strips (I am ignoring electronic sources of noise, but that does not change the validity of these tests). When a wire bond is broken, all sensors upstream of the bond cease to read out results, and the capacitance drops proportionally. This leads to an easily identified decrease in the noise of the strip. The observed noise can be used to identify which strips have broken bonds, and even to identify whether the bond is broken between two sensors (more noise), or just before the readout (less noise). A related defect can result from damage that causes a conductive connection between two aluminum strips. In this case the two strips act as one strip with double the capacitance, and thus a greatly increased noise. Two consecutive channels with similar, significantly increased noise levels are indicative of such a short.

A more serious type of electrical short might occur due to a puncture that causes a connection between the aluminum and the p-implant. When this happens the zero resistance connection causes the amount of charge read out to increase many orders of magnitude. In fact, a couple of these defects (called pinholes) can be enough to saturate the entire electronic readout chip, causing all 128 readout channels to shut down. Thus, these are very serious problems.

In this high current environment the afflicted channel will usually have an even lower noise than a broken wirebond, and thus can be directly identified through the noise test. I also applied a more direct approach, which was to shine laser light upon the biased sensor and see how the channels react. The band gap in silicon is small enough for even visible light to create a current (it is a very sorry technician that forgets to turn off the bias voltage before opening the testing chamber to remove a module). In the presence of low intensity laser light, a current will flow in the healthy channels. In saturated channels with pinholes, however, the readout has already been shutdown, so no signal will be seen.

The software and electronic test stands I used to run these tests were set up by others. I would estimate that I tested about 1,000 modules (in many cases multiple times), thermal cycled about 100 of them, photographed many of the defects I discovered under a microscope, and wrote software to upload and retrieve the test results for several thousand of these modules from an SQL database.

3.3 Calorimetry

I have been significantly less involved in calorimeters over the course of my graduate student career, so I will have less to say about them. The word

"calorimeter" can refer to different things in different fields, but in high energy physics, calorimeters refer to hardware that is used to absorb particles and measure the energy that they deposit.

Calorimeters can be either homogeneous or sampling. Homogeneous calorimeters are built so that signal is read out from the entire volume at the same time as it is absorbed. Building the entire device out of a heavy scintillating crystal is a common approach. Sampling calorimeters are built in layers of absorber and scintillator. In sampling calorimeters only a small fraction of the deposited energy is measured, but more effective absorbers can be chosen, and are generally sandwiched in layers between scintillators which read out the electrical cascades that are produced. These are the type of calorimeters that are used at CDF.

In collider experiments there are generally two further categories of calorimeters that are built. First an inner calorimeter is built that is intended to absorb and measure the energies of electromagnetic showers from photons and electrons. Outside of this is a calorimeter that is intended to measure the energies of hadrons. The electromagnetic calorimeter can generally be built thinner, using a lighter absorber, because electromagnetic showers are more localized as explained in Section 3.1. This also leads to a better energy resolution for electromagnetic particles. Hadronic particles, on the other hand, will easily puncture

the thin layers of electromagnetic calorimetery, and must be absorbed in the (usually coarser) hadronic calorimeter.

3.4 A Needle in a Cosmic Sized Haystack (Triggers)

When high energy protons and antiprotons collide, by far the most likely outcome is a soft scatter or the production of a handful of jets. At Tevatron energies, if your interest is in b-physics, you must contend with the fact that $b\bar{b}$ events will be produced in less than 0.1% of collisions, while only about one in a million collisions produce W's. When it comes to top physics, $t\bar{t}$ events are only produced about once in every 10 billion events, of which only about 20% decay into the decay channel I am studying for my thesis. Searches for new physics or Higgs production must isolate collision events that are even rarer than this. The only way to have a hope of studying interesting physics at the energy frontier is to ramp up luminosities so that extremely rare events will sometimes be produced. At the Tevatron, the proton and antiproton bunches are made to cross within each detector once every 396 ns, and most of these crossings will produce at least one collision.

Of course, this creates a daunting data handling problem. The idea of reading out the entire detector two and a half million times per second is absurd. In practice, sophisticated systems of electronic devices and computers are set up to rapidly study each event as it is produced, and only to read out the most interesting ones. At CDF, Level 1, the lowest, crudest level of triggering, vetoes about 99% of the crossing events, passing data on to level 2 at a rate of about 20 kHz. Level 2, in turn, performs slightly more sophisticated rejection algorithms, and passes collisions on to the level 3 computing farm at a rate of several hundred Hz. In the end, about 100 bunch crossing events are read out to disk per second for use in analyses. These events are vastly more likely to contain interesting physics than the average collision event.

For my analysis, I am interested in selecting $t\bar{t}$ events in the semileptonic decay channel (see Section 4.1). Thus, I use events that pass a series of trigger selections that require a muon (or electron) with transverse momentum (or energy) greater than 18 GeV/c (or GeV). Further details of the CDF triggering system are provided in [17].

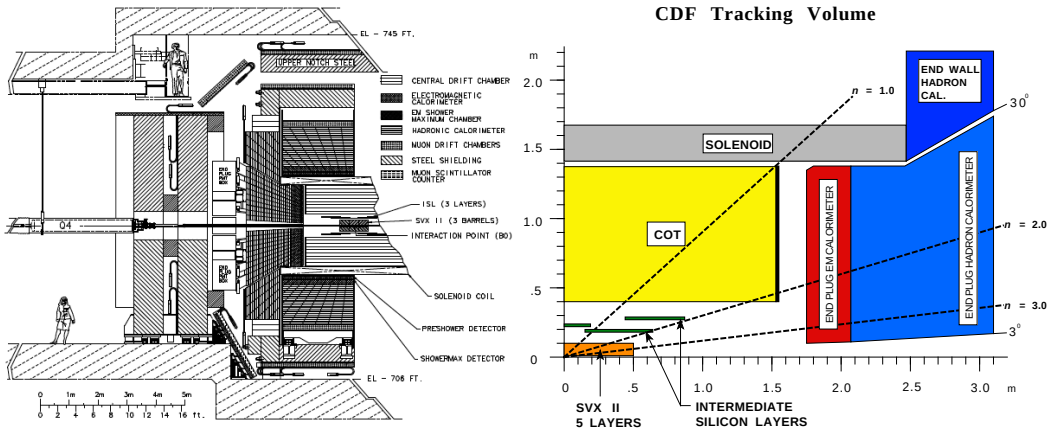

(a) (b)

Figure 3.3: The CDF detector components are arranged in cylindrical layers as shown in these figures. The left picture shows a cross section of the left half of the detector and gives a sense of scale. The right picture zooms in on the inner tracking system.

3.5 The CDF Detector

The CDF detector (see Figure 3.3) incorporates all of the important detector components described above as well as other specialized hardware that I will not go into. I will give some brief details of the important parts of the detector for my thesis. A more complete description can be found elsewhere [17].

As described above, the CDF tracking system consists of an inner silicon tracker, inside of an outer wire drift chamber (the COT), immersed in a 1.4 T magnetic field. The COT consists of eight alternating layers of wire clusters, covering the geometric region between 0.40 and 1.37 m from the beam axis, and less than 1.0 in pseudorapidity in the area at the greatest radial distance

from the beamline. Each successive wire layer is rotated at an angle of two degrees with respect to the previous layer in order to provide longitudinal as well as axial direction information. The inner tracker consists of eight layers of double sided silicon, covering the radial range of 1.35 to 29 cm, and less than 2.0 in pseudorapidity. Again, some implants are rotated relative to the others to provide longitudinal position information. The silicon tracker is vital for the vertexing in my analysis, providing a two dimensional impact parameter resolution of about 70 micro meters for tracks with transverse momentum greater than 0.5 Gev/c, while the wire tracker is vital for the lepton measurements in my analysis, providing a transverse momentum resolution of about 5% for 50 GeV/c muons. Leptons and b-jets are reliably identified out to a pseudorapidity of about 1.0, after which the efficiency drops rapidly due to tracks falling outside of the range of full tracker coverage. Muons are identified from tracks in the inner tracking, which match with track stubs in outer muon drift chambers. Electrons are identified from a track pointing at an isolated cluster in the electromagnetic calorimeter. This calorimeter consists of alternating layers of lead-scintillator, and is located just outside of the solenoid. It provides an energy resolution of about 3% on electrons of 50 GeV transverse energy.

Chapter 4

Top Physics

As shown in Figure 1.1, top quarks are the heaviest particles in the Standard Model, and they decay very quickly. They decay so quickly, in fact, that they do not have time to hadronize. Unlike for any other quark, this allows the properties of the bare top quark, including its mass, to be studied with no fragmentation contamination. In this chapter I will explain what $t\bar{t}$ events look like in the detector, and talk about how they are selected. A key aspect of this event selection is the identification ("tagging") of the b-jets in the decay products. I will put a special emphasis on how b-tagging performance is understood in data, since I have been deeply involved in this work throughout my graduate career, and a solid understanding is essential for my thesis measurement.

4.1 The Top Signature

The top decays in a very characteristic manner, an example of which is shown in Figure 4.1. Electroweak mixing is minimal for the top quark, so under the Standard Model almost 100% of its decays will be to Wb. The Wcan in turn decay either leptonically, or into two jets. So $t\bar{t}$ analyses will involve the selection of jets, and sometimes lepton(s) and missing energy from neutrino(s) depending on which decay channel you are interested in. The most likely decay channel (BR=44%) is the all hadronic channel where both W's decay into jets. This is the most challenging channel to work in due to the overwhelming QCD background. In the semileptonic decay channel one Winstead decays into an electron or muon plus neutrino (BR=29%). This channel still has a relatively high cross section, but the backgrounds are much more manageable. Finally, in the dilepton channel both W's decay to electrons or muons plus neutrinos (BR=5%). The backgrounds are very small in this channel, however analyses requiring large statistics may have trouble finding enough events. Due to the difficulty in identifying τ's, analyses that search for τ's in the finally state are rarely pursued.

For my thesis I measured the mass of the top quark in the lepton plus jets channel. Thus, the signature I was searching for involved four or more jets, a

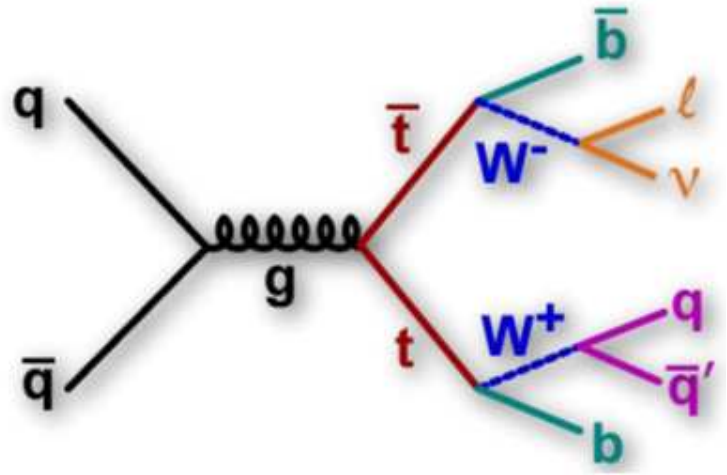

Figure 4.1: Production and decay of a $t\bar{t}$ event in the lepton plus jets channel.

lepton, and missing transverse energy from the escaping neutrino. The base standard event selection used for these analyses at CDF require 4 jets with transverse energy above 20 GeV. This selection leads to a sample that is predominantly W+jets, with a $t\bar{t}$ signal purity of only about 35%. An analysis on an event sample with such a large background would be quite challenging. The most effective way to reduce this background is to tag one or more of the jets as originating from b-quarks. Since most of the W+jet and QCD backgrounds come from light flavor jets, a reliable identification algorithm will dramatically increase the $t\bar{t}$ purity. This brings me into the topic of b-tagging and my first analysis project.

4.2 b-Tagging

The identification ("tagging") of b-jets is important not only for top physics, but also for many other types of analyses, such as in searches for a low mass Higgs. There are a number of characteristic properties of b-jets that one can exploit when searching for them. For example, b-jets have a fairly good chance of containing a lepton in their decay products, so there are simple b-taggers at CDF that simply tag the jet if it contains a muon or an electron (though electrons are much more difficult to reliably identify within a jet). The branching fractions to these final states are small enough, however, that such taggers are quite inefficient. There are many other properties of b-jets that one can utilize, the most effective of which is the long lifetime of b-hadrons.

As shown in Figure 1.1, b-hadrons tend to have lifetimes in excess of a picosecond, meaning that when they are given the sort of boost that they will often have at the Tevatron, they are likely to travel on the order of half a centimeter before they decay. Unlike long lived strange hadrons and muons, their decay products will usually be completely contained inside of the tracking system. Unlike unlike top quarks and $W$and Zbosons, b-hadrons will usually travel far enough before they decay that their tracks can be reliably identified as being displaced by CDF's precision silicon tracking system. The false identification

("mistag") of a gluon, u, d, or s jet will be quite rare. However, lifetime based taggers will have a fairly large rate of accepting charm jets since charm hadrons have only a slightly smaller characteristic decay length than b-hadrons, so one must be willing to accept this contamination.

The most widely used lifetime based CDF b-tagger attempts to reconstruct the displaced decay vertex from tracks within the jet. The algorithm starts by selecting tracks that have impact parameters that are well displaced from the primary collision vertex. These tracks are quite likely to originate from a b-hadron decay. Attempts are made to veto tracks that appear to be consistent with the characteristic decays of light flavor particles. Nevertheless a sizeable background of tracks from material interactions and light decays will populate this sample. A fit is then performed to hopefully locate a displaced using these tracks as shown in Figure 4.2. Two track vertices will not be considered unless they pass extra stringent quality requirements to cut down on mistags. If the significance of the vertex's transverse impact parameter projected onto the jet axis (Lxy/σ_{Lxy}) exceeds some cut then the jet is tagged as a b-jet. This is known as the SecVtx algorithm [13]. There are three variants on the input parameters to the algorithm, know as the loose, tight, and ultratight operating points, which have the b-tagging and mistagging efficiencies shown in Figure 4.3. I will show the results of some of my validation studies for all three taggers below. The most

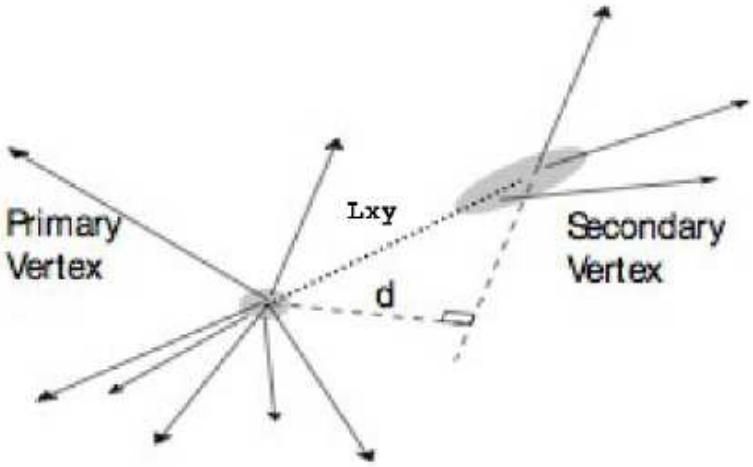

Figure 4.2: A simple illustration of a jet that has been identified ("tagged") with the SecVtx algorithm.

widely used operating point is the tight one, which has about a 40% b-tagging efficiency for a 1% light flavor (udsg) mistag rate for mid energy, central jets.

4.3 b-Tagging Efficiency

Tagging algorithms such as SecVtx lead to dramatic improvements in the sensitivity of top analyses. The real challenge is to understand the performance of the taggers well enough to have confidence in measurements that use them. I have helped develop procedures to determine both the b-tagging and the mistagging efficiencies for the CMS experiment at the LHC, and I have twice performed b-tagging efficiency measurements at CDF. Mistagging efficiency measurements, while an important and challenging topic, play only a small role in my thesis,

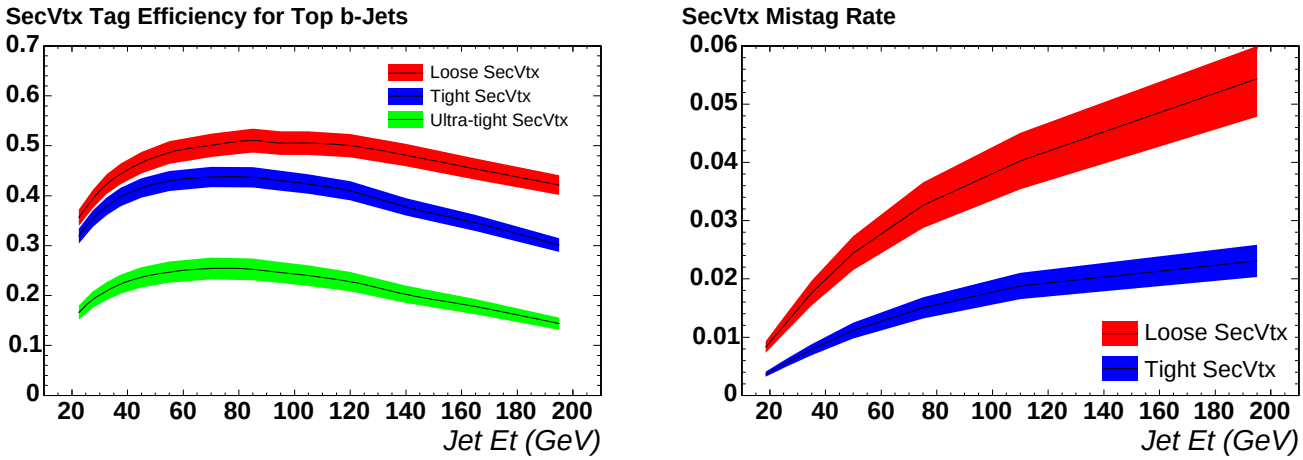

(a) (b)

Figure 4.3: This plot shows the tagging efficiencies and one sigma uncertainty bands for each of the three SecVtx operating points as a function of jet energy. a: corrected efficiencies for b-jets from $t\bar{t}$ events. b: corrected light flavor (udsg) mistag rates as measured in multijet data.

and so I will not discuss them in detail. The b-tagging efficiency measurement, however, is central to the understanding of my thesis, so I will provide a brief explanation of the procedure and my involvement here.

There are two widely used approaches to determine the b-tagging efficiency at CDF, which produce results that are consistent with one another. The one that I have worked on is slightly more sensitive. It is based on selecting jets containing muons and using the properties of the muons to determine the number of b-jets in the sample before and after tagging, and therefore, the tagging efficiency. I will explain this technique in detail below. I have performed this analysis twice,

once on a dataset of 690 pb^{-1}, and again on a dataset of 1.2 fb^{-1}. The numbers and plots that I show will be for the more recent, 1.2 fb^{-1} dataset.

4.3.1 The b-Tagging Efficiency from Muon $p_{T,rel}$

In the Monte Carlo, the tagging efficiency can be trivially determined by counting the number of b-jets before and after tagging at the generator level. in order to evaluate the accuracy of the Monte Carlo, one must determine the same information for the data. The core idea of the muon $p_{T,rel}$ technique is that since b-hadrons are more massive than charm or light flavor hadrons, the muons they produce in their decays will tend to have more transverse boost relative to the jet direction. Thus, by fitting the transverse momentum of the muon relative to the measured jet axis (the muon $p_{T,rel}$) to templates for the different jet flavors, one can determine the b-flavor fractions before and after tagging. The challenge is in determining flavor templates that you can trust, and in determining how tagging efficiencies might change when you look in different samples. As will be seen, the Monte Carlo gets the tagging efficiency close to correct. So the approach used to determine tagging efficiencies for b-jets in data is to trust the Monte Carlo to model the rates accurately to lowest order, and then apply a "scale factor" correction to arrive at a final rate which is derived from studies such as mine, $SF = \epsilon^b_{data}/\epsilon^b_{MC}$, within appropriate uncertainties.

To determine this scale factor, the first step is to select a sample that is dominated by b-jets that you can study. We selected this sample from multijet events where two jets were found back-to-back in ϕ (greater than two radians), one of which contained a muon, and the other of which was b-tagged. The two jets are quite likely to be $b\bar{b}$ from this simple selection. Both jets are required to pass some minimal transverse energy requirements, and the muon is required to be contained within a cone size of 0.4 of its associated jet in $\eta\phi$ space, to have $p_T > 9\ GeV/c$, to be non-isolated and central, to pass through all the silicon layers, and to have inner tracks which match with appropriate stubs in the outer muon tracking chambers. In cases where more than one muon was found within the jet cone, only the highest p_T muon was considered. For this analysis I also required the invariant mass of the tracks used in the SecVtx tagged vertex on the away jet to be $M_{vtx} > 1.5\ GeV/c^2$, though I later convinced myself that this cut was unnecessary because it does not significantly improve the b-purity.

4.3.2 The $p_{T,rel}$ fitting templates

The data for this analysis was taken on a trigger path that requires a muon with $p_T > 8\ GeV/c$. It should be noted that this trigger is prescaled (the trigger is designed to only fire with a small random probability to conserve bandwidth). Since this prescaling becomes more significant at high luminosity, I reweighted

the events to counteract any biases this might produce. To determine templates for the $p_{T,rel}$ distributions of b- and charm- jets, dijet Pythia Monte Carlo was generated and filtered to require the presence of a muon with 8 GeV/c of transverse momentum. We applied the event selection listed above and recorded the muon $p_{T,rel}$. When fitting for the case where the muon jet is tagged in the data, we built our b-template from tagged jets, to account for the small tagging bias to higher $p_{T,rel}$. We did no apply this requirement for the charm and light flavor templates, however, as there were insufficient statistics to do so.

We had three different approaches to determine light flavor templates, all of which produced fairly consistent results. The first was to generate Pythia dijet events and filter out all events with charm or b-jets. We also had two data driven light flavor templates. One of these we took from data where we required that there be no displaced tracks in the event. For another light flavor template we inverted the matching requirement between tracks in the inner tracking and in the muon chambers to select muons that were most likely fakes. All of these templates are shown in Figure 4.4.

Since the non-b templates have such similar shapes it is very difficult to perform a three template fit to the data. Instead, we performed two way fits using each of the four non-b templates in turn to determine the b-tagging efficiencies.

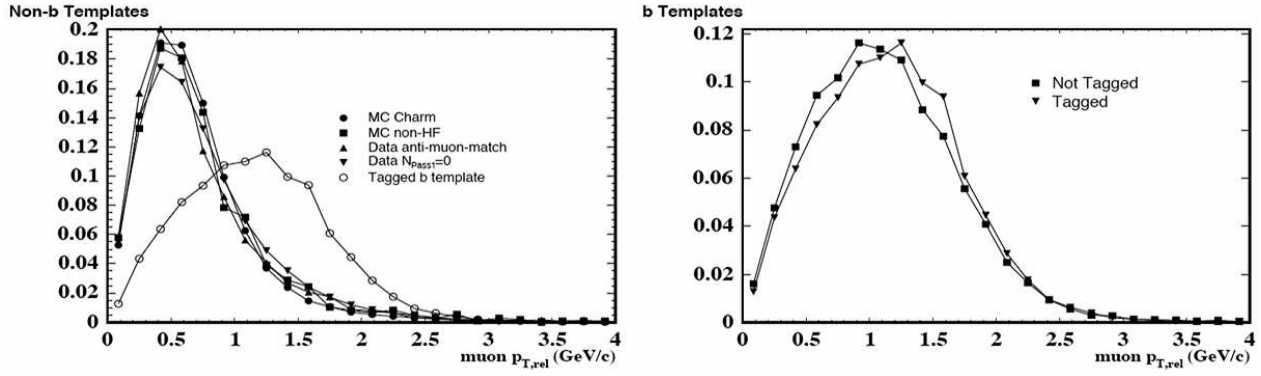

(a) (b)

Figure 4.4: Left: The two data driven light flavor templates, the light flavor Monte Carlo template, and the charm template all show good separation from the tagged b-template. Right: tagging does produce a small bias in the b-template shape.

We averaged the results using each of these non-b templates to measure the tagging efficiency, and took their standard deviation as a systematic uncertainty.

4.3.3 The $p_{T,rel}$ fitting results and trends

Final fit results before and after tagging are shown in Figure 4.5 using charm as the non-b template. We used the ROOT package TFractionFitter to perform the fits. This package performs a maximum likelihood fit to find the fractions of each template within the data assuming Poisson statistics for each bin. In addition, it correctly accounts for statistical uncertainties in the flavor templates by allowing the shapes of the templates to analytically vary within Poisson statis-

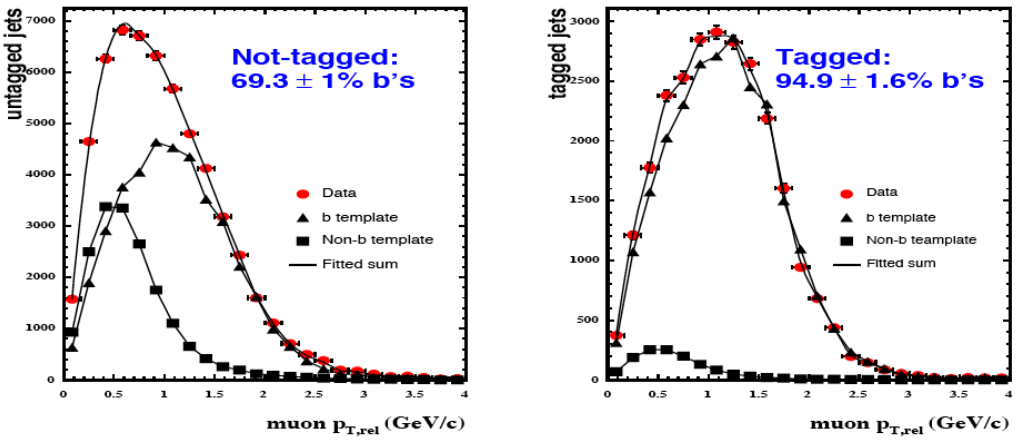

Figure 4.5: Fits to determine the b-fraction for muon jets that are not tagged (left) and that are tagged (right), using the charm template for the non-$b$contribution.

tics. The full fitting algorithm is described in [21]. The final b-fractions and measured scale factors are shown in Figure 4.6.

It is also important to verify that the b-tagging scale factor that is derived using the jets in a $b\bar{b}$ sample is appropriate for jets in other samples. As part of these analyses I checked how the tagging efficiencies and scale factors vary over a variety of kinematic quantities. Examples of three of the variables I studied are shown in Figure 4.7. These plots show the tagging efficiencies in black, the ratio of which determine the scale factor in red. The fitting functions demonstrate that the scale factor is roughly consistent under a variety of conditions. The jet transverse energy and eta plots are important to study because b-jets from $t\bar{t}$ decays will have very different kinematics than in $b\bar{b}$ events. Other variables

Quantity	Loose	Tight	UltraTight
Data efficiency, charm	0.449 ± 0.006	0.382 ± 0.006	0.163 ± 0.004
Data efficiency, nonhf	0.462 ± 0.008	0.393 ± 0.007	0.169 ± 0.005
Data efficiency, $N_{pass1} = 0$	0.465 ± 0.006	0.396 ± 0.006	0.169 ± 0.004
Data efficiency, anti-matched	0.449 ± 0.006	0.382 ± 0.006	0.164 ± 0.004
Average data efficiency	$0.456 \pm 0.006 \pm 0.007$	$0.388 \pm 0.006 \pm 0.006$	$0.166 \pm 0.004 \pm 0.003$
MC efficiency	0.483 ± 0.003	0.416 ± 0.003	0.183 ± 0.002
Scale Factor	$0.944 \pm 0.015 \pm 0.015$	$0.932 \pm 0.016 \pm 0.015$	$0.908 \pm 0.026 \pm 0.015$

Figure 4.6: Tagging efficiencies and scale factor results for each of the three variants of the SecVtx tagger. The data efficiency and scale factor results show two uncertainties: the first is the statistical uncertainty from the fits, and the second is the systematic uncertainty due to the standard deviation of the results using each of the four non-$b$templates in turn.

such as the phi of the jet allow us to study the detector to make sure there are no problems developing in some part of the silicon that is not modeled in the Monte Carlo.

4.3.4 The $p_{T,rel}$ systematic uncertainties

We studied a wide variety of systematic uncertainties for this analysis. Since three of the four non-b templates used in the averaging represent light flavor, this procedure assumes a 25% charm fraction. Reweighting the charm fraction to more realistic values produces a very small systematic shift. Another uncertainty arises from higher luminosity in the data compared to the Monte Carlo. Extra jets from multiple interactions per beam crossing can smear the jet direction as measured in the calorimeter relative to the muon. The tracks of the jet should

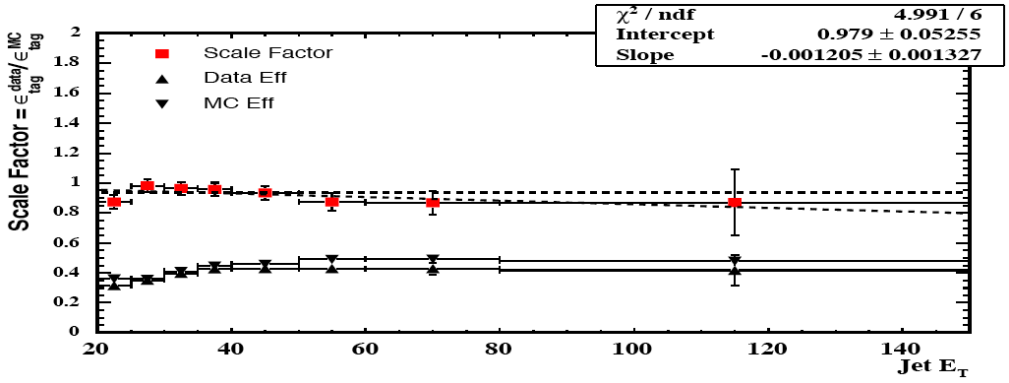

(a)

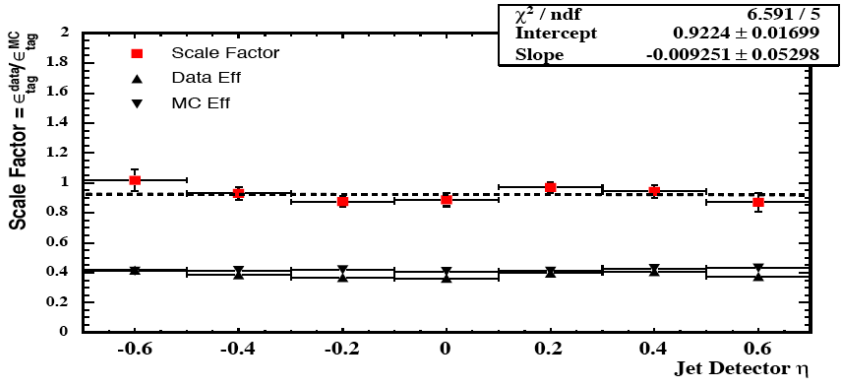

(b)

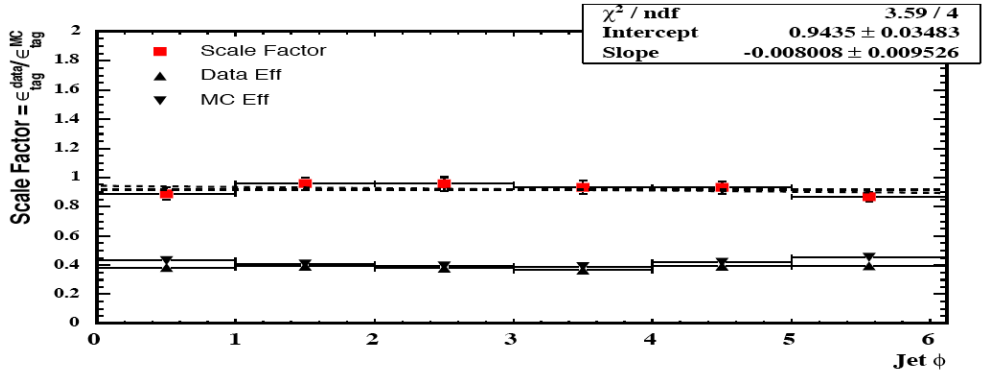

(c)

Figure 4.7: Tagging efficiency trends (black), scale factor trends (red), and their fits as a function of corrected jet transverse energy (top), detector eta (middle, assuming the jet comes from the origin), and phi (bottom).

not be effected, however, so we estimated this systematic by determining jet directions with tracking, and taking the shift in the measured efficiency results as a systematic. An uncertainty due to possible Monte Carlo mismodelling of the tagging bias on the $p_{T,rel}$ distributions was determined by swapping in the wrong templates (tagged instead of untagged and vice versa), and taking 20% of the resulting shift in efficiency as an uncertainty (under the assumption that the Monte Carlo gets this bias at least 80% right). Inaccuracies in the Monte Carlo modeling of b and c-hadron fragmentation and production ratios, which were estimated for a previous version of the same analysis and carried over into mine, are less than 1%.

As expected, the largest systematic uncertainties on this method arise due to determining the tagging efficiency scale factor in the muon jets from the $b\bar{b}$ sample and then applying it to jets in other samples (effects that might show up in trends like Figure 4.7). Differences in jet eta and the track multiplicity of the hadron decays could potentially produce sizeable shifts. Though the systematic shifts that I found for these effects are consistent with zero within statistics, the statical uncertainty is large enough that we took it as a systematic just to be on the safe side. The largest uncertainty, however, comes from the differences in jet energy between the samples. Similarly, the appropriate scale factor for jets in the $t\bar{t}$ sample may be different than for jets from the $b\bar{b}$ sample due to

the much larger transverse energy of the jets. This is a very important problem for my thesis, and it is what drives a substantial fraction of the work I have done as a graduate student. When you convolute the jet transverse energy scale factor trends in Figure 4.3.3 with the spectrum from $t\bar{t}$ jets you get a scale factor that is shifted downwards, though again, this may only be because of a lack of statistics. The systematic associated with this shift is the largest uncertainty on the final scale factor measurement. The complete scale factor uncertainties are shown in Table 4.1.

Table 4.1: The relative (%) systematic uncertainties on the scale factor measurements for each of the tagging operating points.

Systematic	Loose	Tight	Ultratight
Charm Fraction	0.2	0.2	0.2
Multiple Interactions	0.3	0.7	4.4
Tagging Bias	1.1	1.4	1.4
HF Modeling	0.8	0.8	0.8
Decay Multiplicity	0.6	1.7	2.4
Jet η	3.4	1.5	1.2
Jet E_T	3.6	3.9	4.7
Totals	5.2	4.9	7.2

4.4 Tagging Performance in Data

If your goal is to understand the characteristics of a sample that you have b-tagged, you will need to know how the b-tagging has changed things. How many jets of each flavor have passed the tagging selection, and how has the tagging

sculpted their properties? My $p_{T,rel}$ fits show that the simulation does a fairly good job of modeling the characteristics of b-tagging. Under the assumption that it is not drastically worse for charm jets, the procedure used at CDF is to take the simulation results as a lowest order model for the rate and type of b- and c-jets that are tagged. The only correction that is applied is the scale factor, which is about 0.95 ± 0.05 (after combining my results with those of a parallel method) for the tight SecVtx tagger. This correction is usually applied by either deweighting all tagged simulated jets by five percent, or by throwing five percent of them out at random. A five percent uncertainty is then taken on the normalization of the results. This uncertainty is generally increased to ten or fifteen percent for charm jets to account for the fact that the simulation could be significantly worse for them without us knowing. These normalization uncertainties will not, however, cover any biases in the properties of the jets that pass tagging in simulation. Such uncertainties are rarely considered to be significant, but there are certain special cases where it is important. For example, top mass measurements will be biased if the simulation gets the jet energy tagging dependence significantly wrong, so we published the jet energy dependence of our scale factor results (Figure 4.3.3). Analyses were encouraged to reweight their tagged jet energies according to these trends to determine potential biases in the analysis results.

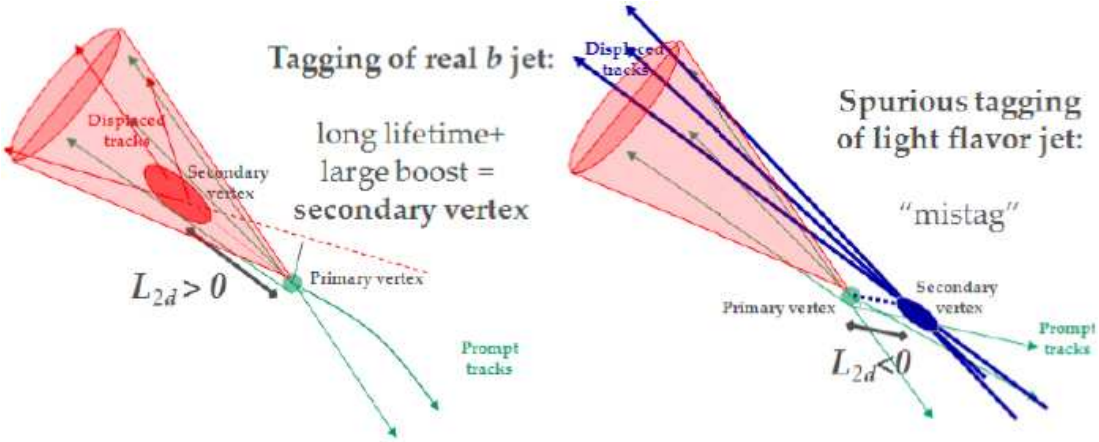

Figure 4.8: "Positive" and "negative" SecVtx tags. Mistagged light flavor jets are roughly equally likely to be positively or negatively tagged.

The simulation has been seen to be much less accurate for light flavor tags (called "mistags"), however, so no one would really trust the simulated results even after a scale factor correction. There are three primary sources of mistags: mistaken tags due to errors in the tracking or vertexing, tags from the displaced decays of real long lived light flavor particles, and tags from vertexed material interactions. Tags from the first category are generally expected to be the most significant. A useful property of this category of mistags is that they are expected to be symmetric about the collision vertex. What I mean by this is that the tagging vertex is equally likely to be displaced from the vertex in the same direction as the jet points as it is to be displaced in the opposite direction. The former type of tag is called a "positive" tag, and the latter type is called a "negative tag", see Figure 4.8.

Thus, to lowest order, negatively tagged jets in data are taken as a good model of the rate and types of negative tags. This rate of negative tags is measured in multijet data and is parameterized in a "negative tag matrix" with five kinematic variables: the jet transverse energy, the jet rapidity, the number of tracks in the jet, the number of extra reconstructed collision vertices, and the sum of the transverse energy of all jets in the event.

This negative tag rate, however, will not account for mistakenly tagged light flavor jets that result from material interactions or real long lived particles, and it will include undesired small contributions from heavy flavor jets, which are mistakenly negatively tagged. All of these effects are compensated for by applying a "mistag asymmetry" correction factor, which is derived (with uncertainties) from vertex mass and lifetime fits to determine the flavor composition of the positively and negatively tagged jets. This asymmetry is applied to convert the negative tag matrix into a mistag matrix. The rate and types of negatively tagged jets are then determined by running over all jets in a simulated sample. But instead of keeping jets that are tagged in the simulation, all jets are kept, weighted by the probability that they will be tagged, determined from the mistag matrix entry that is appropriate for their kinematics, with associated uncertainties.

4.5 b-tagging at CMS

I have also spent some time working on b-tagging at the CMS experiment on the LHC. Many more tagging algorithms algorithms are being pursued at CMS than are in use at CDF, and more tagging validation algorithms are in partial stages of development. One prominent new approach is to use b-jets from $t\bar{t}$ events to predict the tagging rate. This was an option that was not available at CDF due to the much smaller $t\bar{t}$ production cross section there. This approach will probably be very successful, however since it has never been tried before it should be validated against the well-understood algorithms developed at past experiments. Additionally, anyone doing measurements in the top sector will need to account for complicated correlations due to the fact that they are applying the efficiency measurement results to a very similar event sample to what it was derived from.

My involvement has been to supervise two students in implementing certain CDF tagging validation algorithms. We have prepared the CDF algorithms for determining the tagging efficiency using muon $p_{T,rel}$, and for determining the mistag asymmetry (described in Section 4.4). The higher energy range at the LHC poses significant challenges to the $p_{T,rel}$ technique due to the fact that the distinction between the b and non-b $p_{T,rel}$ templates begins to break down. This is

especially dramatic for charm jets. My suspicion is that the mass of the b-hadron begins to play a much smaller role compared fragmentation effects (similar for both b- and c-jets) in determining the template shapes, but I have not had time to study this. Most likely the muon $p_{T,rel}$ algorithm will be a technique that only works reliably for jets with transverse energy below 150 GeV/c^2 or so, and will end up serving as a cross check in this energy range for other techniques that are more appropriate for highly energetic jets.

Chapter 5

Top Quark Mass Measurements

There are both practical and theoretical reasons why an accurate knowledge of the top quark mass is desirable. As explained in Section 1.4, the top mass is one of the most important parameters of the Standard Model when it comes to calculating the expected Higgs mass. The top quark will also be produced in copious quantities at the LHC. The better its properties are understood, the easier it will be to isolate and remove it as a background in searches for new physics. Finally, top quark decays will provide an excellent source of high energy jets at the LHC of all different flavors. Measurements of the jet energies under the constraint that the top quark mass must agree between data and simulation can be used to calibrate the energy measurements of highly boosted jets, and will be the best way to reliably calibrate the energy measurements of b-jets.

However, as explained in Section 1.4, the top mass is already known very accurately. The latest world average top mass combination claims a result of $m_t = 173.1 \pm 0.6(\text{stat.}) \pm 1.1(\text{syst.})$ [2]. Further improvements in the precision of this measurement will be very challenging at a hadron collider, and are unlikely to provide substantial benefits for any of the three applications I have discussed above. Indeed, taking these results at face value, it appears that the top quark mass has been measured with all the accuracy we are likely to need for the foreseeable future, and there is little left of interest to do.

However, as I discuss in Appendix B, measurements sometimes turn out to be badly wrong because they underestimate their uncertainties. We should not just assume that the world average top mass measurements are correct. In fact, the precision of the published measurements are far better than anyone would have predicted ten years ago. Before we dismiss the top mass as a solved problem, it is important to understand what has changed. I will begin by explaining the most important differences between how Run I and Run II top mass measurements are performed, and how the systematic uncertainties are determined. I am quite skeptical of the new results, so I will go on to present some studies I have done that I think illustrate some very significant potential problems with the world average top mass results. These discussions will motivate my thesis, which is

vulnerable to none of these problems, and will serve as an important cross check once enough statistics have been accumulated.

5.1 Run I Top Mass Measurements

As explained in Section 4.1, there are three important types of final states for $t\bar{t}$ production. By far the most sensitive of these channels is the lepton plus jets channel, in which a highly pure sample of $t\bar{t}$ events can be selected with reasonably good statistics. For this reason, measurements in the lepton plus jets channel have always carried a dominant weight in the world average mass combination.

Let us begin by looking at an example Run I top mass measurement in the lepton plus jets channel [11] performed at CDF with 109 pb^{-1} of integrated luminosity. In this measurement the invariant masses of the top and antitop quarks were reconstructed from the jets, lepton, and neutrino decay products, with proper parton assignments chosen according to a χ^2 minimization technique, constrained based upon which jets are b-tagged. A HERWIG Monte Carlo simulation was used to determine invariant mass distributions for a variety of top mass hypotheses, and a maximum likelihood fit was performed to determine

a measured top mass of $m_t = 175.9 \pm 4.8(\text{stat}) \pm 4.9(\text{syst})$. The systematic uncertainty is broken down into their components in Table 5.1.

Table 5.1: Sources of systematic uncertainty for a sample Run I top mass measurement

Systematic	Value (GeV/c^2)
Jet energy measurement	4.4
Initial and final state radiation	1.8
Shape of background spectrum	1.3
b-tag bias	0.4
Parton distribution functions	0.3

Clearly, the systematic uncertainty is dominated by uncertainties in simulations of the measurement of jet energies. Studies had shown that the Monte Carlo modeled these jet energy measurements down to an accuracy of about 4-5% per jet. With two top quarks decaying into six partons, four of which form jets, a 2.5% uncertainty on the mass results due to jet energy measurements is not unexpectedly high.

5.2 Run II Top Mass Measurements

Let us now consider a modern CDF top mass measurements in the lepton plus jets channel. I will present the results of the one that is input to the world average mass combination [3]. It is performed with 3.2 fb^{-1} of integrated luminosity, and measures a top mass of $172.1 \pm 0.9(\text{stat.}) \pm 0.7(\text{JES}) \pm 1.1(\text{syst.})$.

Understanding top mass measurements in Run II is not trivial. Almost all of them are performed with much more sophisticated algorithms. There are some for which the data is fit to signal and background templates to determine the signal fraction and the top mass simultaneously in a manner that bears some resemblance to the Run I approach. Other measurements (such as the one discussed here) use a very sophisticated and computationally intensive approach to determine an event by event mass likelihood based upon the event's kinematic consistency with theoretical predictions, and accounting for jet resolution inaccuracies and the momentum probability distributions for the colliding partons given modern understanding of the proton structure. I do not claim to fully understand these measurements, but the full details of them are not relevant to my concerns. What is important is that the systematic uncertainty is now much smaller, and we need to understand why. The full breakdown of the systematic uncertainty is shown in Table 5.2. There are a lot of new categories of uncertainty listed here, but Similar improvements are seen in other modern top mass analyses.

The most important change is that the huge jet energy based systematic is no longer present. This dramatic improvement is not unique to the complicated algorithm used in this analysis, it is the same for all measurements in the all-hadronic and lepton plus jets channels. The cause of the improvement is a new

Table 5.2: Sources of systematic uncertainty for a sample Run II top mass measurement

Systematic	Value (GeV/c^2)
Calibration	0.2
MC Generator	0.5
ISR and FSR	0.3
Residual JES	0.5
b-JES	0.4
Lepton P_T	0.2
Multiple hadron interactions	0.1
PDFs	0.2
Background	0.5
Color reconnection	0.4

trick that is being used. In the Run I measurements, the jet energies in simulation were determined as accurately as possible and the top mass was measured. Systematic uncertainties were determined by shifting the energies of all jets in the event up or down together within the jet energy uncertainty bounds, and the top mass was remeasured, with the shifts taken as a systematic mass uncertainty.

In Run II measurements, the jet energies are not fixed, rather they are determined simultaneously along with the top mass in the fit. More specifically, the assumption is made that the bias in the simulation of the jet energies is determined by one parameter called Δ_{JES} (JES = Jet Energy Scale), which is the same for all jets in the top decays. Reconstructed top mass distributions as a function of Δ_{JES} are shown in Figure 5.1(a) for a representative analysis in the all-hadronic channel [4]. Regardless of how the fit is done, there is only one

independent way to determine the value of this parameter, which is to constrain the reconstructed Wmass from the hadronic Wdecays to match between data and simulation. This constraint fixes the Δ_{JES} between the data and the Monte Carlo. Most of the jet energy uncertainty then gets wrapped into the fit, which determines Δ_{JES} and m_t simultaneously as shown in Figure 5.1(b) from [3]. As a fit parameter, the JES uncertainty has become a statistical uncertainty, and will approach zero in the high statistics limit.

This really is quite a clever trick. My concerns are not with this calibration procedure itself, but rather with how its uncertainties are determined. In Table 5.2 we see two remaining non-statistical jet energy uncertainties, the "Residual JES" and the "b-JES" uncertainties, which sum in quadrature to 0.6 GeV/c^2. I will explain how these uncertainties are determined in the next couple of sections, but note that the new total jet energy systematic uncertainty has dropped by a factor of more than seven.

5.3 Jet Energy Corrections and Uncertainties

Regardless of whether a jet is simulated or a real data jet, a number of correction factors are applied to the raw energy measurement in the calorimeter to increase its accuracy. A large body of work has gone into determining these

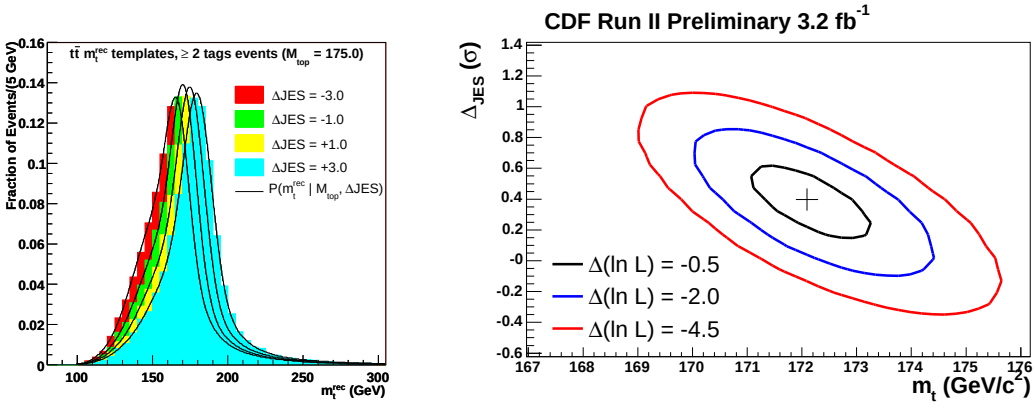

(a) (b)

Figure 5.1: a: Distributions of reconstructed top mass for a fixed generated top mass but where the jet energies are biased up or down by up to three sigma ($-3 < \Delta_{\rm JES} < 3$) from an all hadronic analysis [4]. b: Results of the simultaneous fit to the top mass and the jet energy bias for a lepton plus jets analysis [3].

correction factors, as well as in determining the inaccuracies in the simulation which translate into systematic jet energy uncertainties. Full details on this work are given in [22]. There are five correction factors that are measured. Any or all of these can be applied depending on the tastes of the group doing the analysis, each of which will bring the measured jet energy closer to the true energy of the underlying parton, p_T^{parton}:

$$p_T^{parton} = C_{abs}(C_\eta p_T^{jet} - C_{MI}) - C_{UE} + C_{OOC} \tag{5.1}$$

Many of the correction factors listed here depend upon the kinematics of the jet. C_η corrects the measured jet energies to account for differences in how

different parts of the detector respond to the jets. C_{MI} removes the splash energy associated with extra collisions ("Multiple Interactions") that artificially raise the jet energy. One can then apply an "absolute" correction, C_{abs}, which attempts to correct to the true energy of the particles in the jet, given biases in how the calorimeter responds to the particle interactions. The C_{UE} correction subtracts out the "Underlying Event" energy deposited into the jet cone due to initial state radiation and collisions of other partons within the same $p\bar{p}$ interaction. Finally, one can recover the energy that is deposited "Out-of-Cone" (C_{OOC}), beyond the calorimeter area that is used to reconstruct the jet. In practice, most top mass analyses choose to apply only the C_{η}, C_{MI}, and C_{abs} corrections. The full effective correction factor that is applied to p_T^{jet} in your analysis is called the jet energy scale (JES).

Systematic differences between the data and simulation associated with each of these five categories must be considered, however, regardless of which correction factors are applied. These uncertainties are what modern top mass measurements attempt to eliminate by directly determining the bias in the simulated JES compared to that of the data, Δ_{JES}, as explained in Section 5.2. Next I will explain how C_{abs}, C_{OOC}, and their uncertainties are determined. These will be the two most important factors to understand for top mass analyses.

5.3.1 Absolute Jet Energy Corrections

To determine the absolute jet energy correction factor, C_{abs}, CDF relies upon simulation, and determines uncertainties based upon comparisons with data. CDF finds its corrections by clustering particles in simulated dijet events into generated jets, and comparing the resulting transverse momentum to that of the fully reconstructed jet energies in the calorimeter. The generated particle jet generally has a higher energy than the reconstructed jet as shown in Figure 5.2(a). The correction factor that is needed to convert from the calorimeter jet energy to the particle jet energy is shown in Figure 5.2(b). Regardless of whether C_{abs} is applied, any inaccuracies in the simulation of particle energy deposits in the calorimeter must be taken into account as a systematic uncertainty. There are several sources of uncertainty that must be considered.

One category of uncertainty comes from studying the response of the detector to the collisions of single particles and comparing between data and simulation. The response of photons and electrons are determined from $Z \rightarrow e^+e^-$, $W \rightarrow e\nu_e$, and $J/\psi \rightarrow e^+e^-$ decays. Hadronic particle responses are studied by comparing calorimeter deposits with associated measured track momenta (E/p) for isolated tracks and pions in collision and minimum bias data. For the highest energy particles there are not enough statistics to understand the response under running conditions, and so it is necessary to use the results of test beam

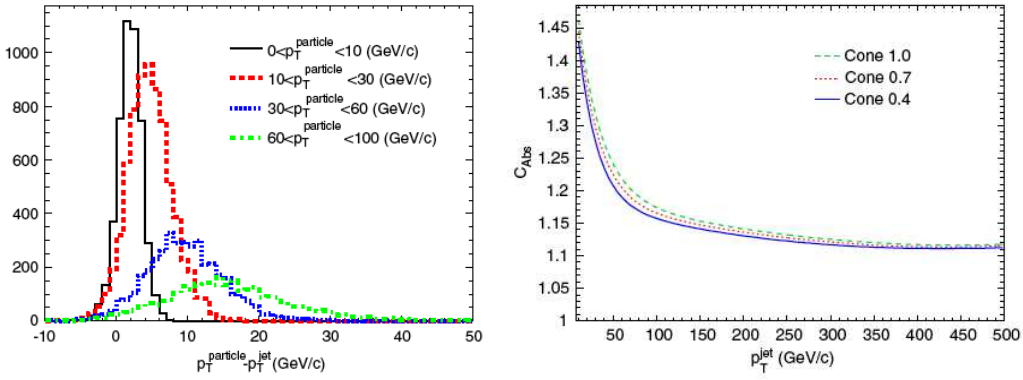

(a) (b)

Figure 5.2: a: Differences between generated and reconstructed jet transverse momenta as a function of true jet transverse momenta for jets clustered with cone size 0.4. b: A larger correction (C_{abs}) is needed for low energy jets.

.

studies performed by shooting high energy pions at the detector. Reasonably good agreement between data and simulation is seen provided that the particles are restricted to hit the center of a calorimeter cell, while disagreements of up to 10% are observed as collisions approach the cell edges. When particles are restricted to impact within the central 90% x 90% of the calorimeter cells, the comparison of data and Monte Carlo performance is shown in Figure 5.3(a). The systematic uncertainty that is chosen to represent possible single particle response inaccuracies is drawn on as dashed lines. For high momentum particles, this uncertainty is driven by uncertainties in the test beam (absolute energy and changes in readout of the modern detector electronics), while for low momentum

particles the smaller uncertainty due to the inaccuracies of the simulation of non-central particles plays the largest role.

Additional uncertainties are needed due to inaccuracies in the simulation of the calorimeter response to multiple particles in a small area, as well as calorimeter response drift during running conditions. The former is determined by comparing the energies of jets clustered from tracks with those clustered from the calorimeter in simulation and data, while the latter is restricted to a small range of uncertainty by constantly tuning and calibrating the calorimeter during running conditions. In the end, the final uncertainty due to calorimeter response, C_{abs}, is given as a function of jet energy in Figure 5.3(b).

One important thing to note about this uncertainty is that it is a composition of several sources of uncertainty (test beam energy, behavior of electronics, behavior of overlapping particles, etc.). Each of the uncertainty curves in Figure 5.3(b) is drawn in a somewhat arbitrary but conservative manner to bracket how inaccurate the simulation might be. Bear in mind that the shapes of these curves were chosen in a somewhat arbitrary manner. Had different shapes been chosen, the impact on the top mass systematic uncertainty could have been significant as I will explain later.

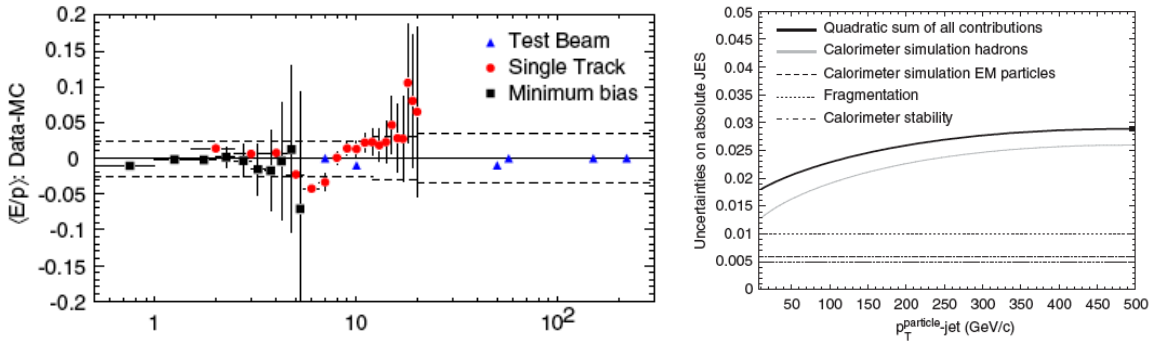

(a) (b)

Figure 5.3: a: Comparison of calorimeter response to charged hadronic particles as a function of particle momentum from a variety of sources under favorable conditions. The one sigma systematic uncertainties are determined from other sources and drawn in as dashed lines. b: All sources of uncertainty on calorimeter response ("Absolute JES") as a function of the true jet energy.

.

5.3.2 Out of Cone Jet Energy Corrections

Most top mass analyses at CDF use jets clustered in a cone size of 0.4 in $\eta\phi$-space. Since low energy jets are quite broad, a significant fraction of their energy can flow outside of this cone and be lost to the reconstruction. The necessary correction, C_{abs}, is determined through simulation in essentially the same manner as for the particle response correction, however since most top mass analyses do not apply this correction, I will not explain the details. What is more important is how the uncertainty is determined.

The consistency between out-of-cone energy in data and simulation is tested in γ+jet and Z+jet samples. In events with no extra jets, the boson transverse momentum must equal the true transverse momentum of the jet by simple conservation principles. The fraction of the boson's energy that is found in calorimeter towers outside of the jet cone is then taken as a measure of the fraction of the jet's energy that is lost due to out-of-cone effects. The mean out of cone fractions for data, and for Pythia and Herwig Monte Carlo simulation are compared. If the means were the same then there one could conclude that the simulation was modeling the data accurately. It turns out, however, that there is significantly more energy out-of-cone in the data than is modeled in either Pythia or Herwig. The difference between the mean out-of-cone energy fractions in data and the two simulations is shown in Figure 5.4(a). Systematic uncertainties are drawn in as the dashed curves to bracket the disagreement between data and simulation, scaled by the absolute calorimeter corrections described in Section 5.3.1. The total jet energy scale uncertainty, accounting for all five categories, is shown as a function of jet transverse energy in Figure 5.4(b).

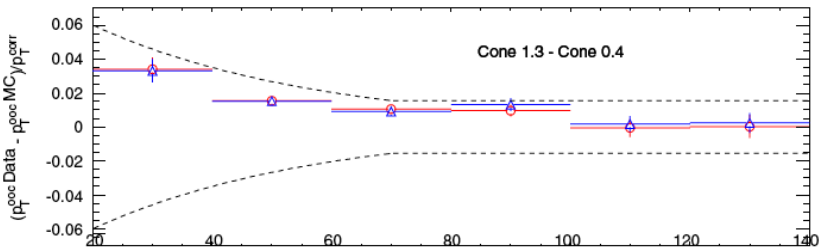

(a)

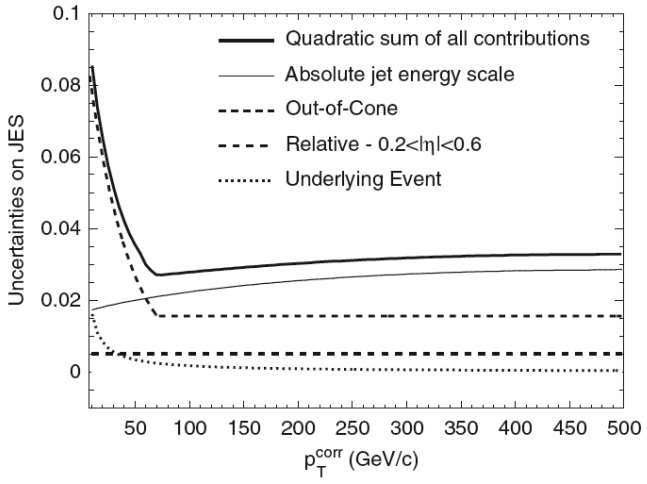

(b)

Figure 5.4: a: The difference in the mean fraction of a jet's energy that flows out of the jet cone in data compared to Pythia and Herwig simulation. The systematic uncertainties in the simulation are then drawn in as the dashed lines, scaled by C_{abs} corrections. b: The total jet energy uncertainty as a function of jet transverse energy from all sources.

5.3.3 Kinematic Jet Energy Uncertainties and the Top Mass

While the W-mass calibration procedure is expected to calibrate away a large part of the jet energy uncertainties in the high statistics limit, systematic uncertainties will necessarily remain due to the fact that there is no calibration information for the b-jets. Official CDF (and D0) prescriptions call for two categories of systematic uncertainty to be evaluated to to cover this effect. One is a "residual" systematic due to kinematic differences between the b-jets and the W-jets which will be described in this section. A second category is due to how differences in flavor between the b-jets and the light flavor jets can bias the results, which I will discuss in Section 5.3.4.

As explained in Section 5.3, there are five categories of corrections that can be applied at CDF to make a jet's measured energy as close as possible to the true parton energy. Regardless of which corrections are applied, each of these corrections accounts for a real physics effect for which there is an associated systematic uncertainty in the simulation[1]. The top mass uncertainties are evaluated for each of these jet energy components individually. Specifically, for a

[1]Actually, there are six categories of residual uncertainty that are considered at CDF. For some reason the out-of-cone uncertainty is split into two parts, one called "out-of-cone" which covers jet energy which is lost out to a ΔR of 1.3, and another called "splash-out", which covers the small amount of jet energy that is lost even farther away than this.

given type of uncertainty, the energies of all simulated jets are shifted up or down by one-sigma together (given, for example, by the solid line in Figure 5.3(b) for the absolute jet energy uncertainty), and top mass is evaluated. One half of the difference between the one sigma up and one sigma down systematic shifts in the top mass are taken as a systematic uncertainty. Under the assumption that correlations between these systematic effects are negligible, the systematics for each component are then added in quadrature to determine a final jet energy systematic uncertainty for the top mass. The CDF lepton plus jet top mass analysis that I have described [3] claims the following systematic uncertainties, shown in Table 5.3.

Table 5.3: Residual jet energy systematic uncertainty components for the CDF lepton plus jets top mass analysis that is used in the world average combination.

(GeV/c^2)	Eta	Multiple Interactions	Absolute	Underlying Event	Out-of-Cone and Splash Out
Systematic	0.06	0.02	0.40	0.10	0.23

5.3.4 Flavor Jet Energy Uncertainties and the Top Mass

The five categories of jet energy based systematic uncertainties discussed above are intended to account for all kinematic differences between the b-jets

and the light flavor jets descending from the W-boson decay. Additional systematics, however, will be needed to account for the fact that b-jets may behave in a very different way than the light flavor jets that were studied to determine the standard uncertainties. For example, b-jets have a slightly higher average energy per parton in their decay products. This will lead to a very small uncertainty that is similar to the single particle response uncertainty discussed in Section 5.3.1. An additional uncertainty is needed due to the fact that a significant fraction of b-jets decay semileptonically, and will lose energy through a neutrino and (sometimes) a muon. The small uncertainty on the branching fraction of semileptonic b-decays translates directly into a small uncertainty on the simulated energy of b-jets.

A third category of "fragmentation" flavor uncertainties that is considered is due to biases in the simulation of the fraction of a b-jet's energy that is carried by the b-hadron. The D0-collaboration has performed studies [5] to determine the Pythia parameters that are needed to reproduce LEP and SLD fragmentation results. When new Monte Carlo samples were simulated using the appropriate tuning, it was seen that the the default Pythia samples produce hadrons that are about 2% too soft compared to the fragmentation products, see Figure 5.5. While there is no reason to expect this difference to significantly effect the measured jet energies, the simulated events were reweighted to reproduce the SLD and

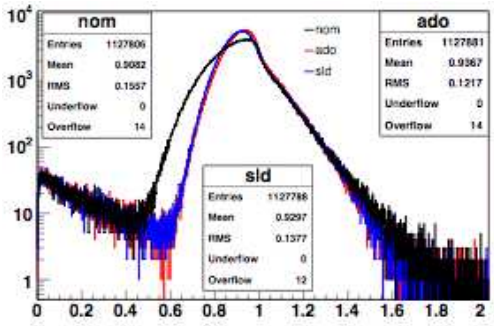

Figure 5.5: This plot shows the distribution of the b-hadron energy within b-jets in Pythia CDF simulation. The quantity plotted is the Bowler parameter, $z = \frac{E_{had}+p_{||,had}}{E_{jet}+p_{||,jet}}$, a quantity that is very similar to the fraction of the jet energy that is carried by the hadron. As can be seen, the default Pythia parameters significantly underestimate z. The rise in the tail at low z is unexpected. I discovered that this is due to a bug in the CDF code that calculates z. I corrected it for my own analysis, but it effects only a small fraction of events.

LEP fragmentation schemes, and a small systematic shift in the measured top quark mass was observed. There are, of course, much more obvious reasons why fragmentation inaccuracies could bias the results of my decay length top mass measurement, so I will spend some time discussing how I handle this later. In the end, these three flavor based jet energy systematic uncertainties have a relatively small impact on most top quark mass measurements. These uncertainties for the same CDF lepton plus jets top mass analysis [3] are shown in Table 5.4.

Table 5.4: Flavor based jet energy systematic uncertainty components for the CDF lepton plus jets top mass analysis that is used in the world average combination.

(GeV/c^2)	Semileptonic BR	Fragmentation	Single Particle Response
Systematic	0.10	0.35	0.12

5.4 My Concerns

To summarize the top mass results and uncertainties so far, the top quark mass used to be measured with a precision that was limited primarily by the jet energy systematic uncertainty. Then, as discussed in Section 5.2, a clever trick was developed to calibrate the jet energy uncertainties by constraining the W-mass to match between data and simulation. This calibration turns all of the previously claimed jet energy systematic uncertainties into statistical uncertainties, and new systematic uncertainties must be understood for the calibration. As I have explained in Sections 5.3.3 and 5.3.4, there are many new categories of jet energy uncertainties that are being claimed to cover this calibration, but none of them are very large. In the end the final systematic uncertainty for all jet energy effects combines to about 0.6 GeV/c^2, about a factor of seven reduction in the uncertainty. This trick is singlehandedly responsible for the impressive precision that is claimed in modern top mass measurements, and therefore in turn is responsible for a significant part of the modern limits on the Standard Model Higgs Mass (Section 1.4). When so much rides on the validity of a single technique, I think it is important to examine and test its assumptions and limitations.

One of my concerns is that no one has ever looked at b-jet shapes to see if it is appropriate to claim the official out-of-cone jet energy systematic uncertainties for them. So I went ahead and did some studies myself which I will present in Section 5.4.1. Ideally I could apply the results of my studies to a real analyses using real data. But I cannot use my own analysis because it works in a very different way, and I was unable to find any volunteers to apply my results to their measurements. I was unwilling to spend the time to transcribe complicated algorithms or to run over large datasets. So I instead devised a simulated "toy" analysis algorithm and toy data to study, which I will introduce in Section 5.4.2. I will then use this toy analysis to estimate the how big a flavor based out-of-cone uncertainty might be, and examine some other assumptions of the W-mass calibration procedure.

5.4.1 Flavor and Jet Shape

Of all the systematic uncertainties that are assigned for flavor based jet energy differences (Section 5.3.4), none of them address differences between the jet shapes of b-jets and light flavor jets. As discussed in Section 5.3.2, there are significant uncertainties in how much of a jet's energy flows outside of the jet cone, especially for low energy jets. Figure 5.4(a) shows that this uncertainty is determined by comparing the amount of jet energy that is found outside of

jet cones in γ+jet events in data and in simulation. The simulation appears to underestimate the fraction of a jet's energy that flows outside of the cone.

These jets are, of course, almost all light flavor, and probably will have very similar shapes to those of the jets from the W-decays that are used for calibration. On the other hand, they may have very different shapes from the b-jets from the direct top decays. In particular, most of the out-of-cone energy that is found in Figure 5.4(a) probably comes from fragmentation products. Thus, inaccuracies in the fragmentation simulation is probably what drives the differences observed here. For b-jets, the decay products of the b-hadron itself are occasionally sent outside of the jet cone for low energy jets. Additionally, the muons and neutrinos in the decay products will be invisible to the calorimeter, leading to a smeared jet direction, and a potentially exacerbated out-of-cone problem.

For my b-jet shape studies I selected a sample of $b\bar{b}$ jets using the same selection as for the b-tagging efficiency scale factor measurement that I performed (Section 4.3.1). Specifically, I required two tagged jets, mostly back-to-back in ϕ, one of which contained a muon. My fits in Figure 4.5 show that the jet containing the muon is about 95% likely to be a b-jet in data. But to avoid semileptonic biases in the jet shape, I instead studied the other, "away" jet, which will also

usually be a b-jet. As a cross check I also used this tracking algorithm on the γ+jets events to see if I could successfully reproduce the official results.

The dataset files I had easy access to did not contain calorimeter tower information, so I instead studied jet shapes using tracking. Basically, I studied the ΔR distribution of tracks about the jet axis as measured in the calorimeter. To suppress contamination from multiple interactions I vetoed all events where a even a single track was found that was more than 3 cm displaced from the primary vertex in z, and to suppress contamination from overlapping jets I vetoed events where any third jet with corrected E_T ¿ 12 GeV was found. I required the tracks themselves to pass through all layers of the silicon and the central tracker, and to reach the calorimeter. I extrapolated the trajectory of the tracks through the magnetic field to the calorimeter surface and plotted the momentum distribution in ΔR of these tracks. I then compared the mean fractions of jet energy that flow outside of the jet cone between data and simulation. The results for both b and light flavor jets are shown in Figure 5.6.

I can use these results to draw my own predictions of what the out-of-cone jet energy uncertainties should be. I have plotted the difference in the fraction of jet energies found outside of the cone between the data and the simulation in Figure 5.7. The behavior of the light flavor results looks very consistent with the official out of cone results of Figure 5.4(a).

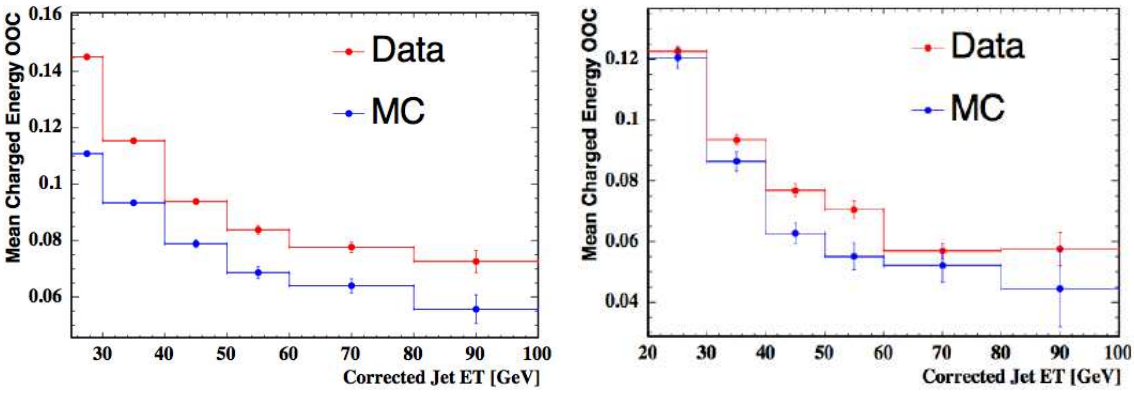

(a) (b)

Figure 5.6: Fraction of tracking energy that I found outside of the jet cone for a: light flavor jets, and b: b-jets, as a function of jet transverse energy.

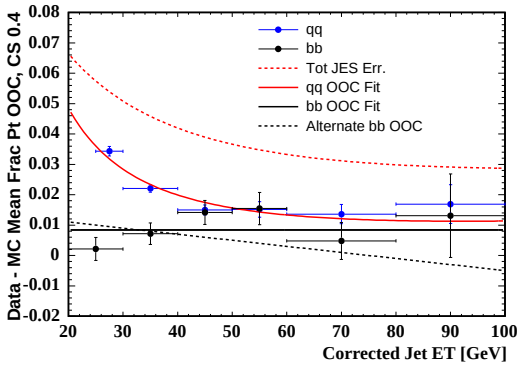

Figure 5.7: The difference in the fraction of jet energy that is found outside of the jet cone between data and simulation for b-jets and light flavor jets (qq). Fits to the results are shown in the solid red and black lines. Due to statistical limitations, however, there is a lot of uncertainty in how the b-jet distribution behaves at high E_T. I have drawn on a dotted black line to represent a very different scenario that is not ruled out by the study. Also shown, in dashed red, is an approximation of the total jet energy uncertainty from Figure 5.4(b).

5.4.2 A Toy Top Mass Analysis

I wanted to test the results of my out of cone studies on a top mass analysis. So I devised a simulated toy analysis that I think captures the spirit of a real top mass analysis while being much simpler. Basically, I make approximate invariant mass distributions for the top quark, and see how much the mean of the distribution changes under alternate systematic scenarios. This shift should be a rough estimate of the top mass systematic uncertainty.

I began by taking the standard Pythia Monte Carlo sample and considering events that pass the selection cuts that I use in my own analysis (see Section 6.1). Briefly, I require an electron or muon with at least 20 GeV/c of transverse momentum, four or more jets of at least 20 GeV of corrected transverse energy, at least one of which must be b-tagged, and at least 20 GeV of missing transverse energy. From these events I look into the Pythia generator records and pick out the real simulated lepton, jet, and neutrino four-momenta, and use them to reconstruct invariant top and W-mass distributions with no ambiguity in particle assignment, and no backgrounds. These distributions are shown in Figure 5.8.

I will take the distributions of Figure 5.8 as being "true" invariant mass distributions. I will then change various aspects of the simulated jet energies to correspond to potential reconstruction inaccuracies in simulation. I will apply a W-mass "calibration" by comparing the difference between the mean hadronic

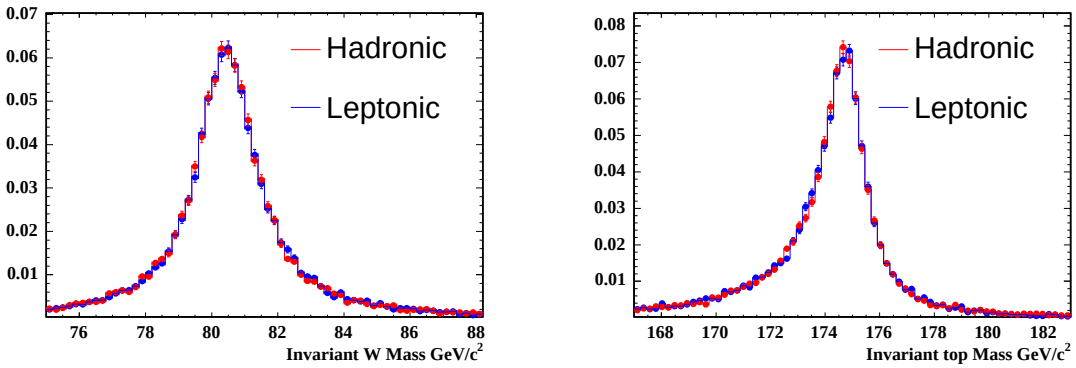

(a) (b)

Figure 5.8: Invariant top and W-mass distributions for events passing my selection. These distributions were formed by adding the simulated four-momenta of their decay products and determining the associated invariant mass.

W-mass in the shifted jets with the "true" mean. In this way I will determine a Δ_{JES}, and apply it as a correction factor to all jets in the event equally, parameterized as a function of jet transverse energy according to Figure 5.4(b). The difference between the mean of the corrected top mass distribution and the mean of the true top mass distribution will then be my approximation of the systematic uncertainty. Clearly this procedure is very simplified compared to what happens in a real analysis. But it makes the same two key assumptions: that all jets in the events can be modeled as sharing the same energy bias (Δ_{JES}), and that this bias can be determined based off of the level of agreement between the expected and observed hadronic W-mass.

5.4.3 The Residual Out-of-Cone Uncertainty for my Toy Analysis

First I will directly calculate an out-of-cone uncertainty in the official manner of Section 5.3.3 for my toy analysis. The prescription asks that I scale all jet energies down by the jet energy dependent out-of-cone uncertainty from Figure 5.4(a). As a result, the reconstructed W-mass shifts down as shown in Figure 5.9(a). This is a significant fraction (73%) of the full one-sigma jet energy uncertainty shift shown in blue in the same figure. This, my calibration procedure kicks in and scales all jet energies back up by $\Delta_{JES} = 0.73$ of the dashed red full one sigma jet energy correction curve of Figure 5.7. The result is an almost exact cancellation of the out-of-cone bias as shown in Figure 5.9(b). This is exactly what the W-mass calibration is supposed to do.

Of course, the real question is what effect this calibration has on the top mass. Since b-jets have a higher energy on average than the light flavor jets, we will be sampling slightly different parts of the out-of-cone and combined uncertainty spectra. Since a one-sigma out-of-cone shift is slightly smaller at higher energies, this leads to a slightly over-estimated correction. The results are shown in Figure 5.10. The full analysis with all jets shifted down one-sigma in out-of-cone uncertainty leads me to an upward mean top mass shift of 0.33

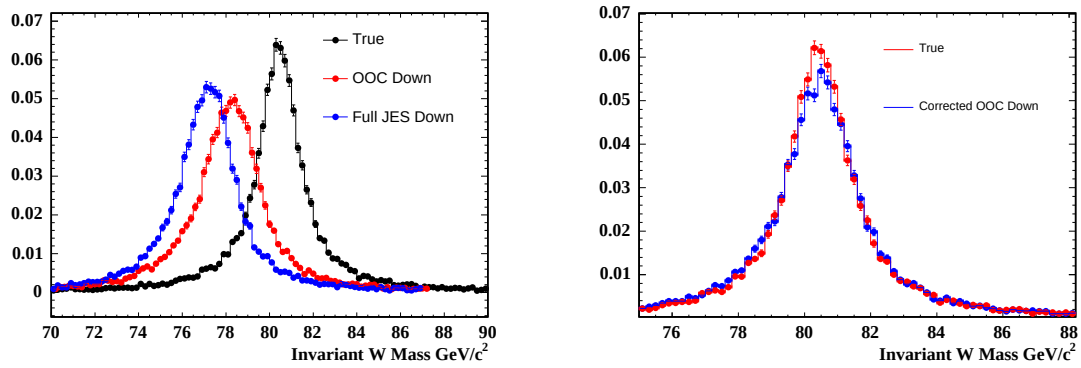

(a) (b)

Figure 5.9: Plots of the invariant W mass. a: The relative size of a one sigma out-of-cone jet energy shift compared to a full one sigma shift combining all sources of jet energy uncertainty. b: The W mass calibration procedure almost exactly cancels out the effects of the out-of-cone shift.

GeV/c^2 on the hadronic side, and 0.271 GeV/c^2 on the leptonic side. Following the official prescription[2] for determining out-of-cone uncertainties I would be led to claim a 0.3 GeV/c^2 residual systematic due to out-of-cone effects, which is very similar to what real analyses that use the W-calibration claim, see Table 5.3.

So that's what I get when I evaluate my uncertainties the way I am *supposed* to evaluate them. But what do I think the uncertainties *really* should be?

[2]Technically the prescription calls for me to take half the difference between the analysis results for out-of-cone up and out-of-cone down for my systematic. But since I am not short on statistics I am being lazy and using the full shift between out-of-cone down and the central value instead

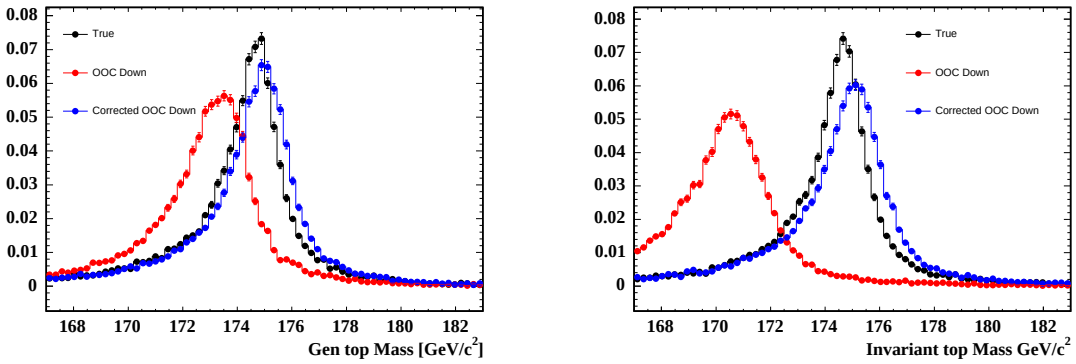

(a) (b)

Figure 5.10: The true invariant top mass distribution (black), the distribution after applying a one-sigma downwards out-of-cone shift (red), and the distribution after applying the W-calibration on top of this (blue). a: Results for the leptonically decaying top quark, b: Results for the hadronically decaying top quark. The hadronic top quark has a much larger out-of-cone shift due to the fact that three jets are being corrected instead of just one, however the calibration is no less effective.

5.4.4 Why I Doubt the Official Top Mass Results

A reasonable definition of an out-of-cone uncertainty should cover most of the possible range of fluctuation in the top mass that could result from out-of-cone mismodelling in the simulation. Well, my studies in Figure 5.6 tell me how wrong the simulation is, depending on jet flavor and energy. Let me treat the uncorrected results as the "true" jet energies. I will then construct a "data" distribution by scaling the jet energies down according to the solid black and red fits in Figure 5.6 to account for the fact that more energy is lost out-of-cone in the data than in the simulation.

The analysis sees a mean W-mass that is 54% of a full one-sigma jet energy uncertainty shift lower in the data than in the simulation, so it corrects the energies of data jets upwards by 54% of a one-sigma jet energy shift. Once again, the corrected W-mass in data perfectly matches the "true" W-mass from the simulation. However, the same cannot be said for the top mass. As can be seen in Figure 5.6, the simulation is much more accurate for b-jets than it is for light flavor jets. A 54% correction will be much too large, and leads to an overcorrected top mass (by 0.69 GeV/c^2 on the hadronic side and 0.63 GeV/c^2 on the leptonic side). These results are shown in Figure 5.11. Note that this discrepancy is more than twice as large as predicted by following the official systematic uncertainty prescription.

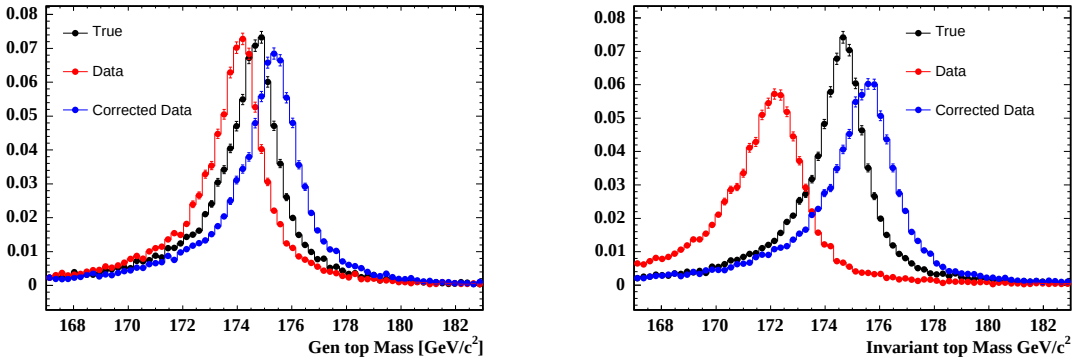

(a) (b)

Figure 5.11: These plots show the simulated invariant top mass in black (which I will pretend is the true distribution), the results with jet energies shifted according to the solid red and black fits of Figure 5.7 to make "data" distributions in red, and the results after calibration according to the W-mass procedure in blue. a: top mass on the leptonic side, b: top mass on the hadronic side. Compare this shift with the "official" systematic uncertainty of Figure 5.10.

But it should be noted that these results are derived assuming the fits of Figure 5.7, and there is a fairly significant statistical uncertainty on these fits, especially for the $b\bar{b}$ shape. Let us try the alternate hypothesis for the $b\bar{b}$ shape that is given by the dashed black curve of Figure 5.7. This shape is allowed within the statistics of our fit. On a naive glance you might think it is more optimistic: the curve is almost centered around zero for jets in the kinematic range seen in $t\bar{t}$ events. But this is in fact very bad news because the more accurate the $b\bar{b}$ simulation is, the less appropriate it is to use the same Δ_{JES} for them as for the much less accurate light flavor jets. When the analysis is rerun under this hypothesis, the over-correction is even more dramatic, leading

to systematic shifts of 1.18 GeV/c^2 on the hadronic side, and 1.15 GeV/c^2 on the leptonic side. If they were to see similar shifts in the official CDF lepton plus jets analysis [3], it would be about as large as all of their systematic uncertainties combined.

5.4.5 The Official Systematics are Truly Bizarre

Let me repeat: under my W-mass calibration procedure I will measure a less accurate top mass if I use more accurate simulated b-jet shapes. Further, the official prescription for estimating the systematic uncertainty associated with these shapes may be dramatically wrong. But the weirdness of the official procedure does not stop there.

Consider the official total jet energy uncertainty shape of Figure 5.4(b). It has a very distinctive shape where the uncertainty is much larger for low energy jets due to the inaccuracy of the out-of-cone modeling. Let me examine the consequences of this shape with my toy analysis. I will evaluate the systematic top mass uncertainty due to the absolute jet energy uncertainty (Section 5.3.1) under a variety of alternate hypotheses for the shape of the total jet energy systematic uncertainty. The possibilities that I will consider are shown in Figure 5.12. In red I have drawn a rough approximation of the total official jet energy uncertainty shape of Figure 5.12. The dashed purple line shows a different hypothetical

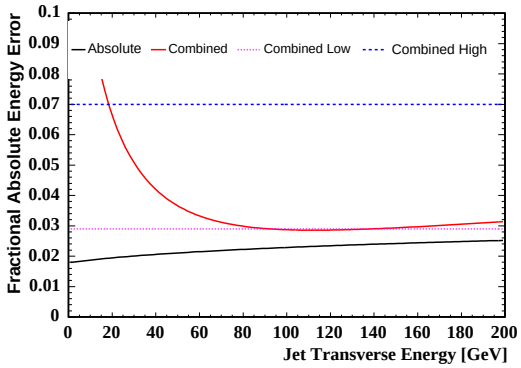

Figure 5.12: Various jet energy uncertainty shapes that I will study. Red: a rough approximation of the combined jet energy uncertainty shape from Figure 5.4(b). Dashed purple and blue: some alternate scenarios for the total uncertainty shape. Black: a rough approximation for the shape of the absolute jet energy uncertainty from Figure 5.3.

jet energy uncertainty shape, one that would roughly correspond to a situation where the out-of-cone jet energy discrepancies have been fixed. A much more pessimistic scenario is shown in the dashed blue shape. In black I show a shape that roughly corresponds to the official absolute jet energy uncertainty of Figure 5.3.

I evaluate the top mass uncertainty due to the absolute jet energy uncertainty in the official way, just like for the out-of-cone uncertainty of Section 5.4.3. I shift the energy of all jets up according to the absolute uncertainty of Figure 5.3, and evaluate the associated top mass shift. Once again, the W-calibration procedure will cause a correction to be applied to bring the top mass back close to expectations. When the official shape is used for the combined jet energy uncertainty,

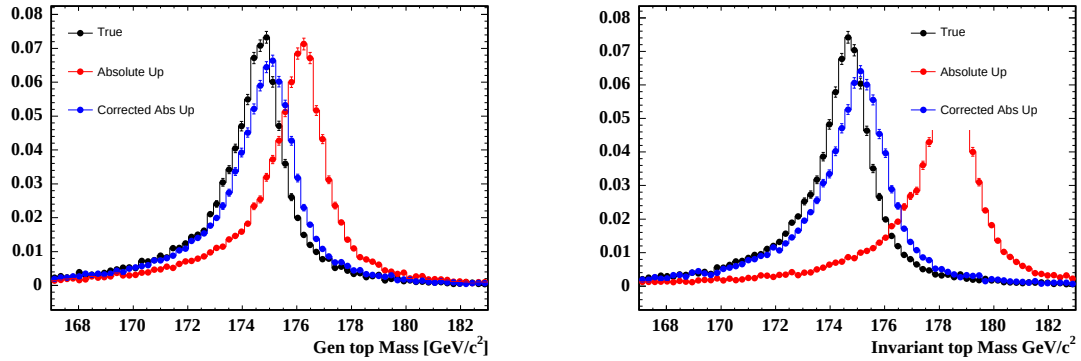

(a) (b)

Figure 5.13: Evaluation of the absolute jet energy systematic uncertainty using the "official" shape for the total jet energy uncertainty (solid red in Figure 5.12.

the final top mass is corrected according to the results of Figrure 5.13. The absolute systematic uncertainty is evaluated to be 0.28 GeV/c^2 on the leptonic side, and 0.39 GeV/c^2 on the hadronic side.

These results do seem pretty close to the results in real top mass measurements (see Table 5.3), but what happens when we evaluate these uncertainties using different shapes for the total jet energy uncertainty? One would hope that for a stable method the absolute jet energy uncertainty would only depend upon the absolute component of the uncertainty, not upon the shape of the total uncertainty curve. It turns out that the systematic uncertainties are much smaller: less than 0.1 GeV/c^2 in all cases! How did this happen?

The explanation is not trivial, but it comes down to the fact that the shapes of these jet energy uncertainties are as important as their magnitudes. CDF chose the shape of its absolute jet energy uncertainty to be mostly flat. If you also are using a flat shape for your total jet energy uncertainty then you will be shifting all jet energies up by the same about as the calibration corrects them by: there will be no differences in how jets of different kinematic ranges are treated.

Thus, if you set all jet energy uncertainties to a large flat value (over jet energy), then while this might seem conservative, you will be fixing your top mass jet energy uncertainty at zero! Another way of looking at this is to consider the fact that the only significant difference between the shapes of the default total jet energy uncertainty and that of the alternate optimistic scenario of the dashed purple line is in how the out-of-cone systematics are treated. The absolute and out-of-cone jet energy uncertainties are assumed to be completely independent of one another when the top mass uncertainties are combined. However, by changing the out-of-cone jet energy uncertainty (removing it), I have just reduced my top mass uncertainty due to absolute jet energy effects by a factor of 30! This procedure creates huge correlations between uncertainties that no one has ever attempted to understand.

Throughout these studies I am calculated a lot of systematic shifts for my toy analysis. Let me repeat that I am not claiming that these numbers I have

calculated are a real estimation of what the true uncertainties should be for real top mass analyses. My toy analysis works in a very different way and makes many approximations. But it does share what I think are the two most essential features of the most sensitive real top mass analyses: that it assumes all jet energies are biased by the same amount in the simulation, and it determines this bias by applying a W-mass calibration technique. I have shown that when the official procedures are applied to determine the top mass uncertainties, you can arrive at some truly non-sensical results, and the uncertainties might be greatly underestimated. These strange procedures that no one understands are the dominant reason why the world average top mass is claimed to be so accurate today.

Chapter 6

My Thesis: Top Mass Measurements Using Tracking

So far, I have explained why an accurate knowledge of the top mass is important, how conventional top mass measurements are performed and their uncertainties are determined, and why I am skeptical of the results. My thesis project was to develop a top mass measurement that would cross check conventional measurements. The project grew out of a technique that was developed by my adviser and his postdoc to determine the top mass using the measured transverse decay length (Lxy) of b-hadrons decaying from the top quarks [31]. Because the $t\bar{t}$ system tends to be produced with minimal boost, the transverse momenta of the top quark decay products depend approximately linearly on the top quark

mass. In turn, this means that the lifetime and decay length of the b-hadrons depend approximately linearly on the top quark mass. This was the first top quark mass measurement to be independent of calorimeter-based uncertainties, however it was also very statistically limited, and the systematic uncertainties were very large ($m_t = 180.7^{+15.5}_{-13.4}(\text{stat.}) \pm 8.6(\text{syst.})$ GeV/c^2) with 695 pb^{-1} of integrated luminosity.

For my thesis I worked to reduce both the statistical and the systematic limitations. In addition to using almost triple the integrated luminosity compared to the previous measurement, I also improve the statistical resolution of the measurement by incorporating more information from the decays. Similarly to b-hadrons, the transverse momenta of leptons from W boson decays also depend linearly on the top quark mass, and are mostly independent of the calorimeter jet energy scale. Since the momentum of the leptons is mostly uncorrelated to that of the b-quarks, this is complementary information, and it is an ideal variable to add to the measurement as proposed in Ref. [31]. To illustrate the dependence of these variables on the top quark mass, the expected distributions for $t\bar{t}$ signal events passing our event selection are shown in Figure 6.1.

For my thesis I performed measurements of the top quark mass using both the lepton transverse momentum and the mean decay length variables. Both measurements are performed upon the same events which pass an event selec-

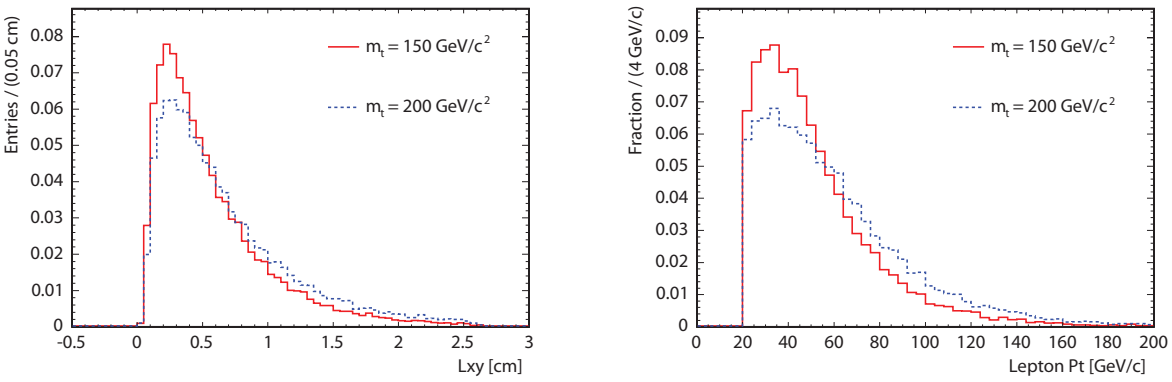

(a) (b)

Figure 6.1: Distributions of the measurement variables for simulated $t\bar{t}$ events after passing our event selection for top quark masses separated by 50 GeV/c^2. These distributions are normalized to unit area.

tion is designed to isolate $t\bar{t}$ events where the W from one top decays hadronically to two quark jets, and the other decays leptonically to an electron or a muon plus a neutrino (the lepton plus jets channel, see Section 4.1). I compute the mean values of Lxy and the lepton transverse momentum for the leptons and the leading two b-tagged jets in the data. I determine the top quark mass through comparisons with the mean Lxy and mean lepton p_T from pseudoexperiments performed for a variety of top quark mass hypotheses. Since the measurements are sensitive to very different event characteristics than typical mass measurements, they will require unique treatments. In particular, it is more important for me to correctly model the boost of the top quarks than for other measurements. Thus, I reweight the simulated signal events to a more accurate parton

distribution function. Further, the Lxy values that are measured in our simulation are sensitive to a variety of possible modeling inaccuracies. I determine correction factors to be applied to the simulated Lxy values by comparing the Lxy measurements between data and simulation for $b\bar{b}$ events, parameterized as a function of the jet energies, in a manner that is similar to what I did with the b-tagging efficiency scale factor, Section 4.3. Unfortunately, this jet energy dependence to the parameterization of the Lxy correction introduces a jet energy uncertainty to the measurement. In order to keep the measurement independent of the jet energy uncertainties of other mass measurements, I developed and applied a tracking based algorithm for measuring jet energies for the Lxy correction parameterization.

I will present my thesis in the following manner. I will start by explaining the event selection used in this analysis and the samples that I studied in Section 6.1. I will then explain the procedures for determining the background shapes and normalizations in Section 6.2, and how the signal simulation is calibrated in Section 6.3. In Section 6.4 I will present the procedures I apply to extract the top quark mass from the final Lxy and lepton momentum distributions and present the results. Finally, in Section 6.5 I will explain in detail how the systematic uncertainties are estimated, and I will conclude with some projections of the future potential for this type of measurement.

6.1 Event Selection

In comparison to the previous publication using only Lxy, I tightened the event selection for my thesis to reduce the systematic uncertainties. The statistical sensitivity of both the lepton momentum and decay length measurements depends linearly on the fraction of $t\bar{t}$ events, and goes as the square root of the number of events selected. This tightened selection improves the former and worsens the latter, and in the end has little impact on the final statistical sensitivity.

The data used in this analysis were collected between March 2002 and May 2007, and correspond to an integrated luminosity of 1.9 fb^{-1}. Events passing the full trigger and event selections described below are studied to determine the expected event counts and uncertainties for each signal and background type. Under the Standard Model both top quarks are virtually certain to decay to Wb. Our selection criteria are designed to accept events where one W decays to an electron or muon plus neutrino and the other decays to two jets. We start from a triggering stream that requires one electron (muon) to have transverse energy (momentum) greater than 18 GeV (GeV/c). Once events are accepted by the trigger, they are saved, reconstructed, and studied in greater detail. Calorimeter towers are clustered together within a cone of radius

$R = \sqrt{(\phi_{tow} - \phi_{jet})^2 + (\eta_{tow} - \eta_{jet})^2} < 0.4$ [22] to form jets. At least three jets must be found under this selection with $|\eta| < 2.0$, and transverse energy greater than 20 GeV after correcting for multiple interactions, and calorimeter response, noise, and detector non-uniformities (C_{abs}, C_{η}, and C_{MI} in equation 5.1). No attempt is made to correct for underlying event or out-of-cone effects, however systematics are still determined for them. To account for the neutrino and suppress the QCD background, a quantity called the missing transverse energy is used, which represents the energy of the transverse component of the four-momentum vector that is needed to conserve momentum in the event [6]. The missing transverse energy in the event must be greater than 20 GeV. Additionally, an electron (muon) must be identified with transverse energy (momentum) greater than 20 GeV (GeV/c). Electrons are formed from a track pointing at a cluster in the calorimeter which matches the expected shape profile. Most of the energy of this cluster is required to be confined in the electromagnetic calorimeter, the track momentum is required to agree with the measured calorimeter energy to within a factor of two, and if the track is consistent with the electron having originated from a photon conversion the electron candidate is vetoed. Muons are formed from partially constructed tracks in the muon chambers which are matched to tracks found in the inner tracking system. The calorimeter energy deposits along the muon trajectory must be consistent with that of a minimum ionizing particle.

Calorimeter isolation is based upon the fraction of energy (f_{iso}) in a cone of radius R=0.4 centered on the lepton, excluding the energy in the calorimeters from the lepton itself. Both electrons and muons must satisfy $f_{iso} < 0.1$. This cut eliminates most fake leptons as well as real leptons which result from b-decays. Further, at least one collision vertex must be reconstructed using tracking for the event, and the track of the lepton must pass within three centimeters of the highest momentum vertex to minimize contamination from multiple interactions and tracking errors. One electron or muon is required to pass these cuts, but to suppress events from the dilepton channel the event is vetoed if any other leptons are found passing a much looser set of cuts. Finally, one or more jets must be identified (tagged) as originating from a b-quark, as explained below. In the case where only three jets pass our selection, two of them must be tagged in order to reduce the large W+jets and non-W QCD backgrounds in this kinematic region.

Jets containing b-quarks are identified using the tight SecVtx algorithm, which relies upon the long lifetime of hadrons originating from b-quarks. As explained in Section 4.2, this algorithm attempts to construct a secondary vertex using tracks that are likely to have originated from a b-hadron decay, which is required to have a transverse distance from the primary vertex projected onto the jet direction (Lxy) with a significance Sig $>$ 7.5. Here, Sig $\equiv$ Lxy/σ, and σ is the uncertainty on the value of Lxy. It should be noted that charmed

daughters of the b-hadron are also likely to travel a significant distance before decaying. The SecVtx algorithm is deliberately designed to be loose enough to attach some tracks from these tertiary decays into one "pseudovertex" at a position that is averaged between two real vertices. Since the boosts of the charm hadrons depend on the boost of the b-hadron, this extra information does not dilute our mass resolution. If the event passes all selection criteria, the lepton transverse momentum and the Lxy of the tagged jet (or of each of its leading two tagged jets if more than one tag is found) are recorded for use in the mass measurements.

6.2 Sample composition

For this analysis I normalize the signal and background distributions using the results of a cross section measurement that was performed on the same data using identical event selection. I briefly describe this measurement in Section 6.2.1, but further details can be found elsewhere. While the cross section measurement procedures have never been published, they are documented in a Ph.D. thesis [7], and bear some resemblance to other published measurements [15] [18]. I devised my own procedures to determine the appropriate

Lxy and lepton p_T shapes for each background sample, which I will deiscuss in Section 6.2.2.

6.2.1 Sample normalization

The first step in the cross section measurement is to determine the number of events for some of the rarest processes from simulation. Backgrounds modeled from simulation include single top production as well as the minor electroweak diboson (WW, WZ, ZZ) and Z plus jets final states. The diboson events are simulated using PYTHIA version 6.216 [41]. The single top and Z plus jets samples are simulated using other programs (Madevent [35] for single top, and ALPGEN version 2.10 prime [36] for Z plus jets) and then showered with PYTHIA. For all of these samples, event counts are determined and scaled according to their theoretical cross sections, branching ratios, and detector and trigger acceptances. The $t\bar{t}$ signal is also simulated using PYTHIA. Its cross section is initially set equal to its Standard Model expectations, and then revised iteratively to determine its true value as explained below. Further corrections are applied to the b-tagging efficiency modeling. I explained the complicated procedures to determine these corrections in Section 4.4.

The simulation also does a poor job of modeling fake leptons, so the non-W QCD contribution is determined from data by performing a fit to the missing

transverse energy spectrum. When studying fake electrons the missing transverse energy of events with isolated electrons passing an inverted selection is used. For muons the standard cuts are kept except that the isolation cut is inverted to require greater than 20% isolation instead of less than 10%. Different binnings, fit ranges, and cuts are applied and the differences in the results are taken as a systematic uncertainty.

The remainder of the observed pretagged events are all taken to originate from W plus jets. The W plus jets sample is simulated using ALPGEN version 2.10 prime [36], and the resulting partons are showered using PYTHIA. This sample is simulated in various bins of jet multiplicity for the $Wb\bar{b}$, $Wc\bar{c}$, Wc, and W plus light flavor final states. These samples are then weighted by their theoretical cross sections and combined. Heavy flavor overlap between the ALPGEN generation and the PYTHIA showering is removed by vetoing ALPGEN events where the two heavy flavor hadrons are generated with a small separation ($\Delta R < 0.4$), and PYTHIA events where the heavy flavor showering is broader than this cut. The fractions of heavy flavor versus light flavor jets in these samples are then directly counted at generator level. However, since the samples are only generated at leading order a correction is needed to reweight the heavy flavor fraction. This correction is taken from comparisons between data and simulation in related samples with higher statistics [18].

By construction, this procedure yields pretagged background and signal normalizations that exactly predict the observed number of data events. After b-tagging, however, fewer events were predicted than were observed, and this discrepancy is attributed to be due to $t\bar{t}$. Thus, the signal cross section was revised upwards and the analysis was repeated in an iterative fashion until the observed number of events agreed with expectations both before and after tagging, at a $t\bar{t}$ cross section of 8.2 pb. The resulting event counts are shown in Table 6.1, along with the number of jets in these events which were b-tagged and included in our Lxy analysis.

Table 6.1: Estimated Signal and Background Contributions

Distribution	Events	Recorded b-Tags
$Wb\bar{b}$	25.4 ± 7.0	37.5 ± 9.7
$Wc\bar{c}$ or Wc	13.9 ± 4.6	15.9 ± 4.7
Wplus light flavor	16.9 ± 3.7	17.9 ± 3.7
non-WQCD	18.8 ± 12.7	20.2 ± 13.2
Electroweak	9.0 ± 0.4	11.3 ± 0.5
Single Top	8.4 ± 0.4	13.4 ± 0.7
ttbar	478.3 ± 40.3	659.3 ± 45.5
Total	570.8 ± 44.3	775.5 ± 50.2

6.2.2 Background shapes

The dominant backgrounds for our analysis are W plus heavy flavor (b and c-jets), W plus light flavor mistags, and non-W QCD events. Along with single top and diboson events, these distributions and the signal account for about

99% of events passing selection. The remaining events come from the Z plus jets background. Since this background is so small, I did not bother to analyze the simulated events and instead used the same Lxy and lepton p_T distributions as for the related W plus jets background.

Single top samples were generated with the masses $m_t = 165,\ 170,\ 175$, and 180 GeV/c^2 in both the s- and t- channels. Results for the s- and t- channels were combined, weighted according to their expected theoretical cross sections, and the final decay length and lepton momentum distributions were fit to Gaussian plus exponential distributions. The trends in the fit parameters were extrapolated to other mass points, and were used to generate new Lxy and lepton p_T distributions for each mass hypothesis.

The W plus jets background is selected in exactly the same way as for the cross section measurements, as is the background for non-W QCD electrons. For non-W QCD background in the muon channel, instead of using non-isolated muons I developed alternative cuts to select fake muons. These cuts were chosen based on studies I performed suggesting that they produced a minimum of bias in the lepton p_T distribution. Resulting signal and background decay length and lepton momentum distributions are shown in Figure 6.2 for events passing full selection. As cross checks the same distributions in the background dominated one and two jet bins are shown in Figures 6.3 and 6.4. These cross check

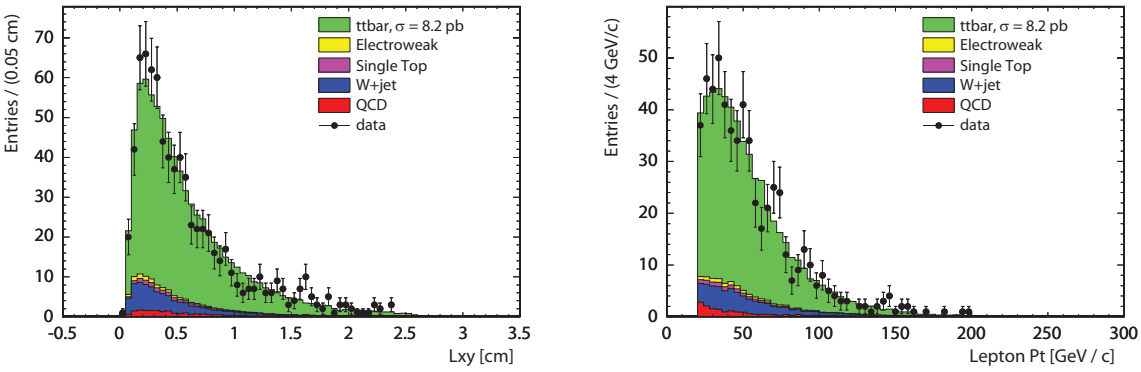

(a) (b)

Figure 6.2: Signal, background, and data for the Lxy and lepton p_T distributions passing full event selection under hypothesis top quark masses that are close to the measured results. The left plot is for the Lxy measurement, using hypothesis top quark mass 168 GeV/c^2, and the right plot is for the lepton p_T measurement, using hypothesis top quark mass 173 GeV/c^2.

samples are used in evaluating the background based systematic uncertainty as described in Section 6.5.

6.3 Corrections to the signal

6.3.1 Parton distribution functions

For this analysis it is vital to have accurate models of the lepton and jet boosts. Since these quantities depend on the energy of the colliding partons, an accurate modeling of parton distribution functions (PDFs) for particles within the proton is very important. The default CDF samples were generated using

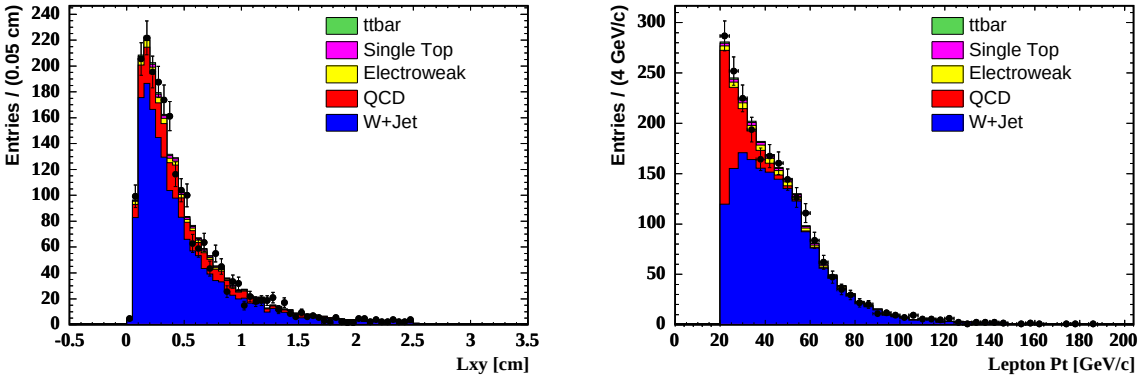

(a) (b)

Figure 6.3: Background prediction compared with data (black points) in the one-jet control region for Lxy (left) and lepton p_T (right).

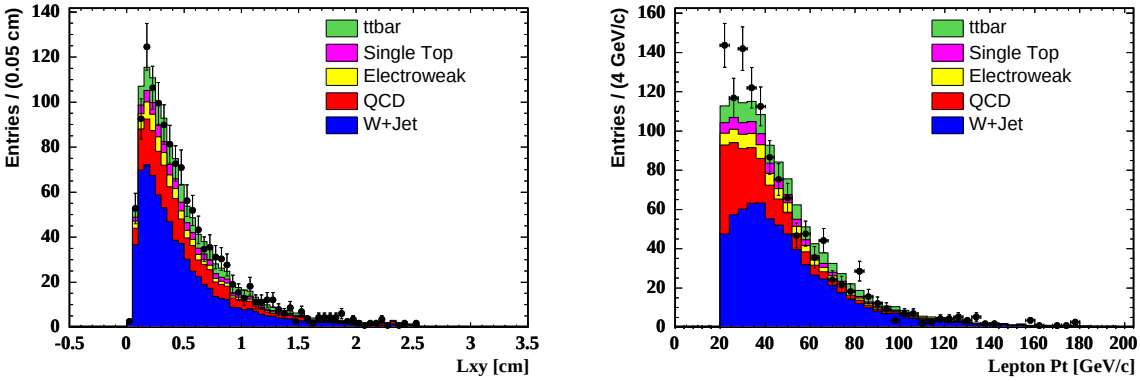

(a) (b)

Figure 6.4: Background prediction compared with data (black points) in the two-jet control region for Lxy (left) and lepton p_T (right).

the leading order CTEQ5L parton distribution function [33]. One drawback of this PDF is that it underestimates the rate of $t\bar{t}$ production by gluon fusion (5% observed versus about 15% expected from theory). Since gluon fusion produces events with slightly smaller boosts on average, this will bias the analysis. In addition, since the CTEQ5L quark PDF is too hard, the $t\bar{t}$ events produced within the quark-antiquark annihilation channel have artificially high boosts.

To compensate for these effects, I had to study different parton distribution functions. Unfortunately, CDF software uses an obsolete package for its PDFs that is no longer maintained, and none of the newer PDFs were available. In order to improve things I installed the Les Houches Accord PDF package [8] and interfaced it with the standard software. I used this software to determine the prior probabilities for each parton involved in the collision to have the generated momentum for a given PDF. One can then reweight the simulated events to construct distributions appropriate to a different PDF as explained in Ref. [10]. I used this prescription to reweight all of the $t\bar{t}$ events to the next to leading order CTEQ6M [38] parton distribution function for gluon fusion and quark annihilation events separately. I applied a further weighting to gluon fusion events to scale the gluon fusion fraction to the value expected for the sample's top quark mass (20% for $m_t = 150$ GeV/c^2, 10% for $m_t = 200$ GeV/c^2). These combined reweightings result in new distributions to be used in our final measurement that

lead to approximately a 1.7 GeV/c^2 shift in the top quark mass for the lepton p_T analysis and a 0.9 GeV/c^2 shift for the Lxy analysis, relative to the values obtained using CTEQ5L. I developed similar prescriptions to reweight to other PDFs and gluon fractions to evaluate systematic uncertainties as I will explain in Section 6.5.

6.3.2 Decay length calibration

There are many effects that may bias the decay length measurement in simulation which I alluded to in Section 4.3. Inaccuracies in the EvtGen [34] values for the hadron lifetime or production fractions would have a direct impact. Similarly, inaccuracies in the PYTHIA fragmentation model would lead to the wrong boost, and thus the wrong average decay length, of the b-hadrons. Entirely different problems may arise from any inaccuracies in the modeling of the tracking system, which could lead to biases in the vertexing results. Just like for the tagging efficiency scale factor, my approach is to calibrate the simulation directly to the data to compensate for all of these biases simultaneously. Systematic uncertainties on this calibration will be discussed in Section 6.5.6. I select the calibration sample just like for the $p_{T,rel}$ method so that it will be enriched in b-jets. I select dijet events where the jets are required to be back to back in ϕ, where both jets are b-tagged, and where one of the jets contains a well resolved

muon with at least 9.5 GeV/c transverse momentum. The fit of Figure 4.4 shows that this sample is composed of approximately 95% $b\bar{b}$, and I performed further fits to determine that the remaining sample is composed of almost entirely charm jets. The charm sample is slightly larger in data than in the simulation, requiring me to apply a small additional correction and uncertainty.

To apply these calibrations, however, one must be convinced that the $b\bar{b}$ jets are good models of the b-jets from $t\bar{t}$. Since jets from $b\bar{b}$ events tend to have a much lower transverse energy than those from $t\bar{t}$ events, the possibility that a different calibration is needed for high energy jets than for low energy jets must be taken into account. Just like for the tagging efficiency scale factor measurement, I bin my $b\bar{b}$ jets according to their energy, and derive the needed correction bin-by-bin. Unlike for the tagging efficiency, however, I will parameterize this correction in the transverse energy of the tagged signal jet. Great care must be taken in doing this, however, as this kinematic based correction directly introduces a jet energy scale uncertainty to the analysis. If, for example, the simulation underestimates the energies of jets, then the average decay length in a given energy bin will be too high in the simulation, and this will artificially bias the calibration, as illustrated in Figure 6.5. This is an unavoidable uncertainty for any kind of calibration using dijets. In fact, even if I could convince myself that it is unnecessary to parameterize this decay length calibration as a

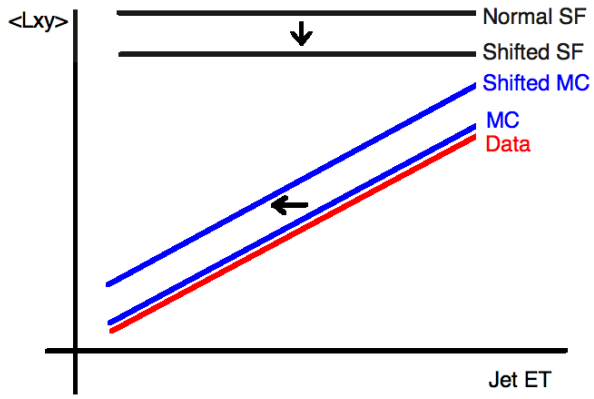

Figure 6.5: Since the mean decay length is proportional to the jet's energy, if the simulation models the wrong reconstructed jet energy this will directly bias the correction factor that is derived.

function of energy, a significant uncertainty based upon jet energy would still be needed based upon determining which jets pass selection in the simulation. At the low energies where jets are selected, the uncertainty on the amount of energy falling outside of our jet cone is especially substantial. As my goal is to keep this analysis independent of the calorimeter jet energy uncertainties, I choose to parameterize the calibration based upon the energies of jets measured using tracking rather than the calorimetry.

6.3.3 Tracking based jet energies

None of the jet experts that I talked with knew of any pure tracking algorithm for measuring jet momentum. I have since discovered that there was such an algorithm, but since I was not aware of it at the time, I invented my

own. My goal was to make the algorithm as simple as possible to minimize systematic complications. I simply measured tracking based jets by adding the four-momenta of all good tracks within a cone size of $R = 0.4$ of the calorimeter jet direction in $\eta - \phi$ space. The calorimeter is used, but only to determine the jet direction for purposes of matching tracks. The tracks themselves are required to be well resolved in both the wire tracking chamber and in the silicon, and they must have a transverse momentum of at least 1 GeV/c. I also require that the tracks pass within 3 cm in the longitudinal dimension of the fitted primary collision vertex, which should eliminate about 85% of the contamination from multiple interactions. I then take the total transverse momenta of the sum of the track four-vectors as the tracking transverse energy. I make no attempt to correct for the missing neutral particles because my goal is not to measure the true jet energy, but rather to measure a quantity that is proportional to the true energy (in this case, the charged transverse momentum) in a manner that agrees well between data and simulation.

I calibrate the energies of the track jets in events where one photon and one jet are found back-to-back, where no extra jets in the event above 3 GeV in E_T are allowed, and strict cuts on photon quality are applied to minimize fakes. Under such a selection, the transverse momentum of the photon should match that of the jet, and the ratio of their transverse momenta can be compared

between data and simulation. This study is done in an identical manner to that used for calorimeter-based jets at CDF [22]. The distributions of the fraction of the photon energy found in the track-based simulated jets show good agreement with data as seen in Figure 6.6(a). The mean energy fractions in data and simulation, along with the ratio of these means, are plotted in Figure 6.6(b) as a function of photon transverse energy. A line is fitted to these ratios and represents the calibration that will be applied to the track jet energies that are measured in our simulation. Specifically, I approximate the true transverse energy of the jet to be the corrected energy as measured in the calorimeter, and use this to look up the calibration that should be applied to the track based energy according to this linear fit. Uncertainties in the simulation of the calorimeter based jet energy measurements are found to have a negligibly small impact on the correction factor that is determined in this manner. There are, however, significant statistical uncertainties on the fitted function that translate into a systematic uncertainty on the results as will be explained in Section 6.5.

When applying these calibrations to the $b\bar{b}$ and $t\bar{t}$ samples, however, it is important to account for the fact that these events are busier than those in the photon calibration sample, and tracks from other jets may fall into our jet cones and bias some of our measured track jet energies upwards. If the decay length calibration parameterization proves to have a trend over jet energy, then such

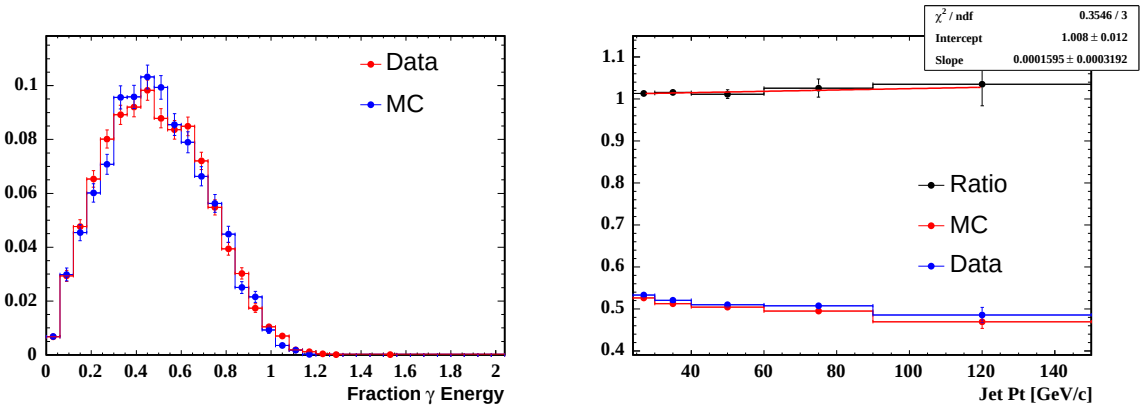

(a) (b)

Figure 6.6: a: The distributions of track jet Pt divided by the photon Pt for events with photon Pt between 30 GeV/c and 40 GeV/c. b: The behavior of the average jet-to-photon Pt ratio for Pythia (red) and the data (blue) as a function of the measured photon transverse momentum (assumed to be the true jet transverse energy). The ratio of these trends between data and simulation is fitted to a line which is then used to correct measured track jet energies in the simulation.

biases in the $t\bar{t}$ sample will have a direct impact on the correction factor that is applied to a given jet. On the other hand, in the $b\bar{b}$ samples there will only be a bias to the extent that PYTHIA does not properly model the amount of overlap that is observed in data. I start by minimizing this problem as much as possible by removing events where a second jet is close to the jet I am studying and might contribute overlapping particles in both data and simulation. For our $b\bar{b}$ events I veto events where the jet is between ΔR of 0.7 and 1.2 of any other jet with greater than 9 GeV of energy. This cut was chosen to remove most events that might have extra jets in a region that could overlap our primary jet without eliminating jets with hard out-of-cone QCD radiation.

For $t\bar{t}$ events, however, such a cut would be inefficient. Instead of vetoing events I develop a correction procedure to remove the effects of tracks from other jets that overlap the jet cones. I do this by plotting the track momentum around $t\bar{t}$ jets as a function of the ΔR between the track and the jet direction as measured in the calorimeter. Since jet shapes depend directly on the jet energy, I make different distributions depending on the generator level energy for b-hadrons in the $t\bar{t}$ sample. I fit these distributions to two components, a "primary" part, and an "overlap" part. The tracks originating from the b-hadron and associated fragmentation products (the primary jet) are modeled by a Gaussian multiplied by a Fermi distribution to force the jet energy to converge

to zero at small ΔR. The contribution from underlying event, minimum bias, and other jets (overlap), turns out to be well modeled by a quadratic function in ΔR, multiplied by a Fermi function to account for damping effects as tracks pass out of the active detector range. The fits are performed separately for each region of b-hadron p_T. Examples of these fits (along with cross check fits for $b\bar{b}$ events) are shown in Figure 6.7. I then use these fit results to extract correction factors, parameterized by the b-hadron's energy, to remove the average amount of momentum from charged particles expected to fall inside the track jet cone from other sources.

There are a number of approximations and assumptions that have gone into these vetoes and corrections. For example, the fitting functions could be inappropriate, or the vetoes applied to the $b\bar{b}$ events might be inadequate. If so, then using an alternate cone size to select track jets will lead to a direct systematic shift. Thus, the systematics for this procedure are evaluated by changing the track jet cone size to 0.7, removing overlap according to the revised jet size, and repeating the analysis, as explained in Section 6.5.6. However, two other cross checks are also run to improve confidence in this procedure. First, I have no physical motivation for using a quadratic function as the base shape of the overlap, so in one cross check I repeat the analysis, modeling the overlap as though it were distributed perfectly uniformly in $\eta - \phi$ space (a line times a

fermi function in ΔR). There is, of course, no physical basis for using this symmetric distribution either, since complicated correlations between jets according to $t\bar{t}$ kinematics and sculpting from the jet clustering algorithm could equally well render the shape asymmetric. This alternate shape is simply a cross check, and in addition to producing a much lower quality fit, it leads to a mass measurement result that is shifted by an amount that is slightly smaller than the systematic I am assigning for this procedure. As a second cross check I look at the tagged jets that are selected to be back-to-back with the muon jet in the $b\bar{b}$ sample. If the overlap fitting function is incorrect, then one would expect the shape of the primary part of the fit to be better modeled in these $b\bar{b}$ events, since the overlap is about a factor of four smaller than in $t\bar{t}$ events. These $b\bar{b}$ comparison distributions are shown as the dashed curves in Figure 6.7. By far the worst agreement is seen in the lowest bin of hadron Pt (shown on the left). In every other bin of b-hadron momentum, the differences between the $b\bar{b}$ and $t\bar{t}$ results can barely be distinguished by eye, as is the case in the right hand plot. But since less than 8% of $t\bar{t}$ jets have hadron momenta that fall into the lowest bin, the systematic mass shifts caused by enforcing the alternate $b\bar{b}$ shapes in the primary distribution will be well less than the systematic that we are claiming.

With the track jets in hand I can evaluate the decay length calibration parameterization in the $b\bar{b}$ sample as described above. The calorimeter is still used to

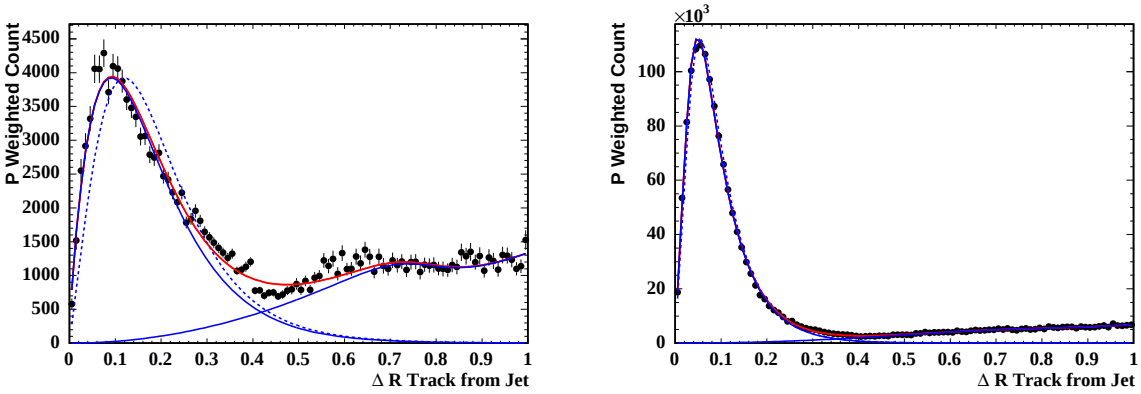

(a) (b)

Figure 6.7: These plots show the charged jet momentum spread in ΔR for $t\bar{t}$ events. Specifically, the ΔR of each track relative to the calorimeter jet direction is plotted as the black points for tagged $t\bar{t}$ jets, weighted by the track momentum. The results are then fit to determine the fractions of primary jet energy and overlap energy from other jets for $t\bar{t}$. The solid blue curves are the fit shapes for the primary and overlap fit components and the solid red curve is the full fit result. The dashed blue curve shows the expected primary distribution based upon fits in the $b\bar{b}$ sample. a: results for b-hadron Pt less than 20 GeV/c. This is the only kinematic range where there is any significant disagreement between the $t\bar{t}$ and $b\bar{b}$ results. Much more typical is b: results for b-hadron Pt between 50 GeV/c and 60 GeV/c.

select jets and determine their direction, but I loosen the calorimeter jet energy cut significantly, and exclude most jets with energies near this cut by applying a track jet energy cut. As a result, calorimeter driven jet energy uncertainties will play a minimal role. Additionally, I apply the overlap vetoes and corrections as described above. Since it is possible for a jet to pass the tracking energy cuts while failing the calorimeter energy selection, fluctuations within standard, calorimeter-based jet energy uncertainties can still cause events to pass in and out of selection. However this is a small effect which is only significant in the lowest energy track jet bins, and thus leads to a small systematic. Figure 6.8 shows the trends in the mean decay length for data and simulation. The ratio of these trends determines the correction that should be applied to the measured decay lengths in data, depending on the measured track jet energy in our signal samples. The distribution of the track jet energies to which the calibration will be applied is overlaid.

6.4 Mass Measurement Method

Given the mean measured Lxy and lepton p_T values from the data I need to determine the associated top quark masses and statistical uncertainties. I

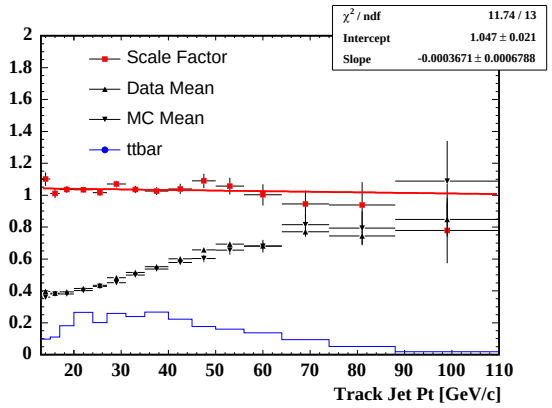

Figure 6.8: The decay length calibration parameterization over track jet transverse momentum is shown here. It is developed from the solid black data and simulated average decay lengths, the ratio of which gives the Scale Factor points, which are then fit to determine the calibration. Also shown is the distribution of $t\bar{t}$ jets to which the calibration is applied. For the Lxy data and simulation the vertical scale represents the mean decay length in cm, while for their Scale Factor ratio and the $t\bar{t}$ distribution this axis is unitless.

simulate experiments under a variety of hypothesis top quark masses and use them to perform each of the measurements as described below.

6.4.1 Single variable measurements

Pseudoexperiment events are drawn separately from signal and background events where the probability of each event is given by its PDF weighting as discussed in Section 6.3. The mean Lxy (p_T) of the tagged jets (leptons) in these events will be used to measure the mass. Samples of $t\bar{t}$ events are generated under 23 hypothetical top quark mass values ranging from 140 GeV/c^2 to a top

quark mass of 220 GeV/c^2, and the decay length results are corrected according to the track jet energies by the fit results of Figure 6.8.

Uncertainties from the background normalization are small and are wrapped into the pseudoexperiments. A total of 92.5 ± 17.1 background events are expected, however due to the iterative nature of the $t\bar{t}$ cross section evaluation, there is a small additional uncertainty on the background normalization of ± 3.7 based upon the theoretical uncertainty on the input $t\bar{t}$ cross section, leading to a total background count uncertainty of ± 17.5. For each pseudoexperiment, the total number of background events is fluctuated according to a Gaussian with the above mean and RMS to determine an expected background normalization. The resulting number is then fluctuated according to Poisson statistics to determine a background normalization for each pseudoexperiment. Given a fixed number of observed data events, the excess is taken to be signal. Further, for each pseudoexperiment the Lxy calibration parameterization is fluctuated within its fitted statistical uncertainties as explained in 6.3.2.

The mean Lxy and lepton p_T pseudoexperiment results are Gaussian in shape and are fit to a Gaussian for each hypothetical top quark mass. To evaluate the top quark mass results for the Lxy and lepton p_T measurements, the central values of the Gaussians are plotted as a function of top quark mass and are fit to a quadratic polynomial. The mean Lxy and lepton p_T values measured in

data are then converted to the measured top quark mass values according to this polynomial. To extract the statistical uncertainties, the central Lxy and lepton p_T pseudoexperiment values are shifted up and down by the standard deviation of the Gaussian fits. Then the difference between the measured top quark mass according to the unshifted polynomial and the top quark mass results resulting from these one sigma shifts are taken to be the asymmetric one sigma statistical uncertainties on the measurements. The fitted polynomials are shown in Figure 6.9. The mean Lxy and lepton p_T values measured in data are also drawn on in the red lines, and the projections used to determine the statistical uncertainties are shown.

6.4.2 Measurement using both variables

The pseudoexperiments from the single variable results are used to plot two dimensional mean Lxy versus mean lepton p_T distributions. The results for the two most extreme mass hypotheses are overlaid in Figure 6.10(a).

The observed data produce a point on this two dimensional mean Lxy versus lepton p_T plane. Given this point, my task is to determine the most likely value of the top quark mass and the associated statistical errors. To accomplish this, I evaluate a likelihood for each mass hypothesis according to the data. The likelihood is simply the probability that if the true mass were the one in our

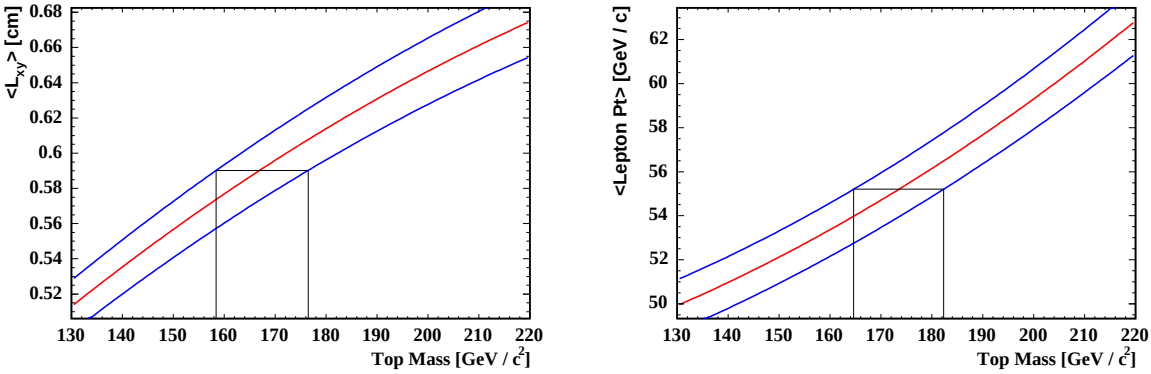

(a) (b)

Figure 6.9: The resulting mean values of the Gaussian fits to the pseudoexperiment results are fit to the quadratic polynomials plotted here in red for the mean Lxy measurement (left) and the mean lepton p_T measurement (right). Polynomials associated with a one sigma statistical shift are also shown in blue. The mean values from data are drawn on in the horizontal black lines. These values are then translated into top quark masses according to their intersections with the red polynomial, and into statistical uncertainties according to their intersections with the blue polynomials.

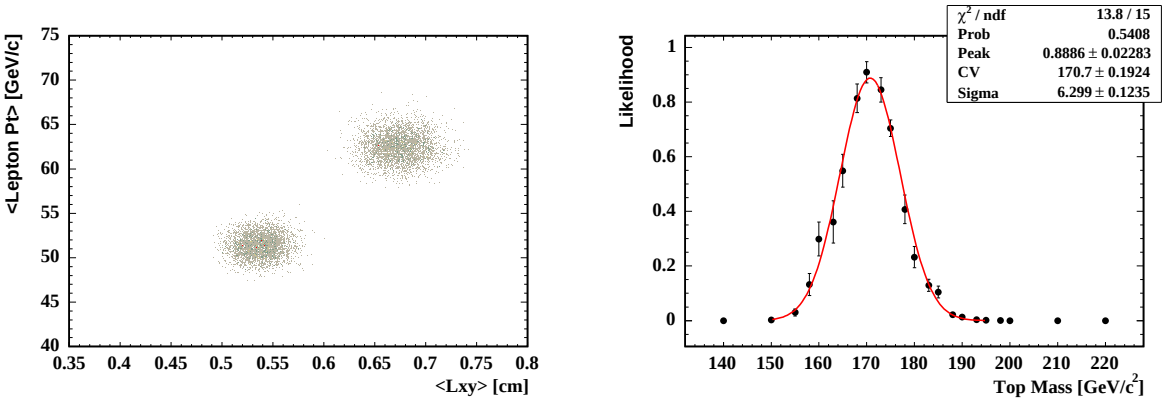

(a) (b)

Figure 6.10: a: distribution of mean Lxy versus mean lepton p_T from pseudo-experiments for extreme mass cases of 140 GeV and 220 GeV. b: the pseudo-experiment results for the 23 mass points considered are used to determine the likelihood of agreement with the data according to the metric of Equation 6.1, and the results are plotted and fitted here. The mean fit result is taken as the measurement result, and the RMS represents our statistical uncertainty.

hypothesis, the mean Lxy and lepton p_T results would fluctuate as far away or farther than the results of the data point. This probability is taken from pseudoexperiment results such as those overlaid in Figure 6.10(a). Specifically, I evaluate a "distance" that the data point is from the expected central value for the mass hypothesis, and take the likelihood that the hypothetical mass is correct to be the fraction of pseudoexperiments which are "farther away" from the expected values than the data point.

For this approach to be meaningful, a reasonable definition of "distance" must be used. I choose the definition so that the distances are equal for points

along the equal probability contours of a two dimensional Gaussian centered at the expected mean Lxy and lepton p_T values, with the expected standard deviations. Here, the expected means and standard deviations are taken from the Gaussian fits to the mean Lxy and lepton p_T pseudoexperiment results that were described for the single variable measurements. Then, the "distance" from the expected value is defined according to the metric in Equation 6.1:

$$D = \sqrt{\left(\frac{\delta P_t}{\sigma_{Pt}}\right)^2 + \left(\frac{\delta L_{xy}}{\sigma_{Lxy}}\right)^2} \tag{6.1}$$

Here, δP_t (δL_{xy}) is the difference between the mean lepton p_T (Lxy) of the data and the fitted central value for the hypothesized top quark mass, and σ_{Pt} (σ_{Lxy}) is the fitted standard deviation of the mean lepton p_T (Lxy) for the hypothesized top mass. In principle this approach could be modified to account for a correlation between the two variables, but this is unnecessary, as the correlations are empirically determined to be zero.

Based on this equation, the fraction of pseudoexperiments for each hypothesis for which the distance metric is evaluated to be larger than that for the data is taken as the likelihood for the hypothesized mass. Finally, the likelihood values for each mass point are plotted with statistical uncertainties determined by adding in quadrature the statistical uncertainties from the pseudoexperiments, and the statistical uncertainty due to the number of simulated events from which

they are drawn. I run enough pseudoexperiments (4000) that the size of our simulated $t\bar{t}$ samples is the primary limitation. These likelihoods are fit to a Gaussian as shown in Figure 6.10(b). The mean of the Gaussian is taken to be the result of our combined measurement with a statistical uncertainty given by the RMS of the fit.

6.4.3 Sanity checks

To verify that the algorithm is unbiased, I choose nine mass points that were not included in the pseudoexperiments used to evaluate the likelihoods. When results are evaluated for the Lxy, lepton p_T, and simultaneous measurements, the resulting residuals (measured results minus expected results) are distributed linearly and without bias as shown in Figure 6.11. Again, the error bars are due to both the statistics from the finite number of pseudodata experiments and to the finite number of events in the mass samples. The pulls (residuals divided by widths) are also examined. In the case of the asymmetric errors for the one dimensional experiments, the error is chosen based upon the direction of fluctuation. If the statistical uncertainty is correct and the distributions of pseudoexperiments are Gaussian in shape then the pull distributions should have a fitted width of 1.0. The mean widths of the pull distributions are shown in Figure 6.12, and show good agreement with this hypothesis. The method was

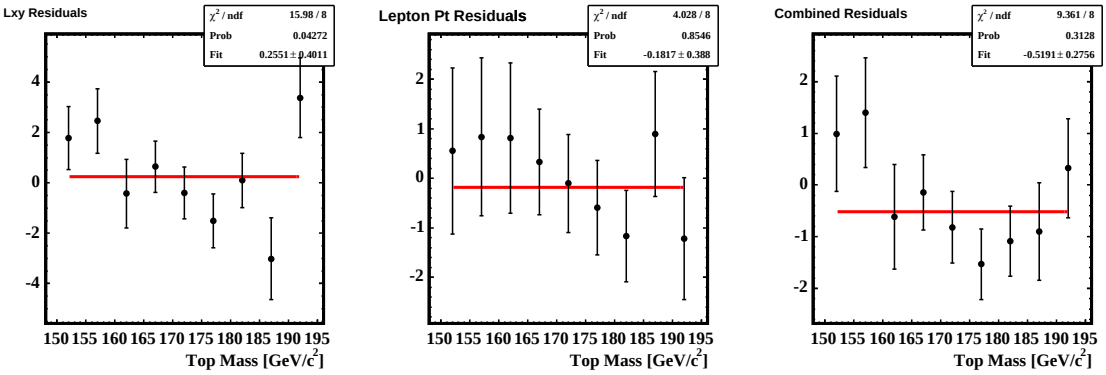

Figure 6.11: Offset between input and output top masses (residuals) expected for Lxy, lepton Pt, and in combination. Statistical errors are based on number of pseudoexperiments thrown and finite Monte Carlo statistics.

further validated by examining the residual and pull width results for ten blind Monte Carlo samples for which the true mass values were kept secret. Again, the results came out consistent with expectations. I concluded that the method has no bias worth considering, and that the statistical errors evaluated are sufficiently accurate [1]

6.4.4 Results

The data events passing event selection have a mean Lxy of 0.596 ± 0.017 cm and a mean lepton p_T of 55.2 ± 1.3 GeV/c. Based upon these values, the mass measurement with the decay length technique yields a result of $166.9^{+9.5}_{-8.5}$ GeV/c^2 ,

[1] I recently realized that the combined algorithm is biased due to the conversion of asymmetric to symmetric uncertainties. This results in a roughly 0.3 GeV/c^2 upwards shift to the top mass, which is far to small to be seen given the limited statistics of the samples I have run here, and is negligible on the scale of my statistical errors. It can safely be ignored.

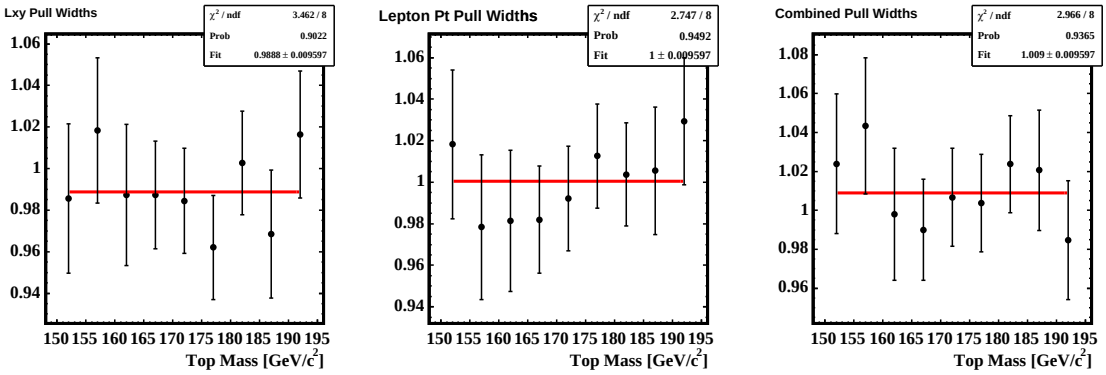

Figure 6.12: Pull widths expected for Lxy, lepton Pt, and in combination. Statistical errors are based on number of pseudoexperiments thrown and finite Monte Carlo statistics.

and with the lepton transverse momentum technique yields a result of $173.5^{+8.8}_{-8.9}$ GeV/c^2, where the errors are statistical only. For the simultaneous measurement with both variables the fit to data is shown in Figure 6.10 and corresponds to a mass result of 170.7 ± 6.3 GeV/c^2.

6.5 Systematic uncertainties

6.5.1 Background uncertainty

As mentioned above, the uncertainty on the background composition and shape is evaluated in the control regions of the one- and two-jet bins, shown in Figures 6.3 and 6.4. The differences between the observed and expected means are shown in Table 6.2. The largest disagreements are observed in the one-jet

bin of Lxy, and in the two-jet bin for lepton p_T. Since both of these worst-case shifts are larger than their uncertainties, they are taken as the uncertainties on the background mean results and are scaled by the background fraction to determine the systematic errors.

Table 6.2: Background shifts and uncertainties in the one and two-jet control regions. The uncertainties account for both statistical effects due to data limitations as well us uncertainties in the relative contributions from the individual backgrounds from the cross section measurement.

Variable	Shift
One-Jet Lxy	-0.0131 ± 0.0082(cm)
One-Jet lepton p_T	0.25 ± 1.25(GeV/c)
Two-Jet Lxy	0.0022 ± 0.0118(cm)
Two-Jet lepton p_T	-2.22 ± 1.05(GeV/c)

6.5.2 QCD radiation uncertainty

Uncertainties in the simulation of QCD radiation could have a significant impact on the results of an analysis like this which is dependent upon an accurate modeling of the boost of the decay products of the $t\bar{t}$ system. Comparisons of the dilepton boost for Drell-Yan events have been made between data and simulation [16] and used to constrain inaccuracies in the modeling of initial state radiation in quark-antiquark interactions. PYTHIA parameters were varied to conservatively bracket the possible disagreement between data and simulated initial state radiation, and analogous parameters were simultaneously varied for

the final state radiation by an equivalent amount. Signal events were generated with these parameters shifted up and down, and we compare these samples with each other and the nominal PYTHIA sample. Half of the largest mass shifts between any pair of these three samples was taken as the QCD radiation systematic uncertainty.

6.5.3 Parton Distribution Function uncertainty

As described in Section 6.3, for this analysis I reweight the events to match both the predictions of the next to leading order CTEQ6M parton distribution function, and the expected gluon fusion top production fractions predicted by theory. The advantage of using the CTEQ6M PDF is that it includes a prescription for estimating PDF uncertainties [38]. The degrees of freedom of the PDF can be parameterized in 20 orthogonal sources of uncertainty that are commonly called eigenvectors. Associated with each of these uncertainty sources are two parton distribution functions, where the parameters associated with these eigenvectors are shifted up or down to cover a 90% confidence interval. I reweight the top quark mass 175 GeV/c^2 sample to each of these forty alternate PDFs, taking half the full mass shift for each pair as a systematic uncertainty, and adding them in quadrature. While these uncertainties are intended to represent a 90% confidence interval, I conservatively use them as one sigma systematics instead.

While I fix the fraction of $t\bar{t}$ events produced by gluon fusion interactions to theoretical expectations, it should be noted that some of these eigenvector variations are expected to change this fraction. Thus, I allow the gluon fractions to float around their expectations for purposes of determining systematic uncertainties on the PDF results.

One uncertainty that is not accounted for by these eigenvectors is the uncertainty on the strong coupling constant, $\alpha_S(m_Z) = 0.1176 \pm 0.0020$ [27]. To study the effects of this uncertainty I reweight to the CTEQ6A and B PDFs [39], which are two different series of PDFs constructed with varying α_S values in intervals of 0.002. I average the mass shifts obtained when varying the PDFs, and arrive at an uncertainty that is roughly half as large as the eigenvector uncertainty. I observe consistency between the A and B series PDFs and add this uncertainty in quadrature to the eigenvector uncertainty to determine our full PDF systematic. As a final cross check to the results of the CTEQ Collaboration, I also reweight to the MRST Collaboration's NLO PDF MRST2004 [37]. I observe agreement well within our stated eigenvector uncertainty when the comparison is made with the corresponding CTEQ6A/B PDFs with similar α_S values.

6.5.4 Generator uncertainty

In CDF top quark mass analyses it is conventional to re-evaluate the top quark mass using samples produced with the HERWIG 6.510 generator [26], and take the shift from the PYTHIA mass result as a systematic uncertainty. Note that many of the differences between these generators will double count our existing systematic uncertainties. The different fragmentation models between PYTHIA and HERWIG will double count our decay length scale, jet energy, and QCD Radiation uncertainties. HERWIG also does not properly handle QED radiation off of leptons from the W decays, which is instead inserted with the PHOTOS [20] program. Differences in these approaches will double count our lepton energy scale uncertainty. Another minor difference is that HERWIG does not account for transverse Fermi motion of the colliding partons within the protons, which is a further artificial inaccuracy, though its effects are quite small. The generators also have minor differences in the applied top width (and HERWIG has a sharp cutoff preventing the presence of top quarks in the high and low mass tails), and only HERWIG properly handles spin correlation between the two top quarks. Though many uncertainties are double counted, I follow the convention of other analyses for consistency by taking the difference between our PYTHIA and HERWIG mass results as a Generator systematic. For the Lxy and Combined

measurements the statistical uncertainty on our mass shift is greater than the shift itself, so we take the uncertainty on the shift as our systematic instead.

6.5.5 Lepton momentum uncertainty

I test the modeling of the lepton momentum in simulation by fitting the invariant mass of Z's in data and simulation, separately for electrons and muons. A Lorentzian is used to model the inherent Z width, which was convoluted with a Gaussian to account for detector resolution. Additionally, to model the kinematic fall off of cross section with higher mass Z production this function is multiplied by a decaying exponential. When a function modeling a QCD background shape is included the fits return zero for its normalization as expected due to the high purity achieved by our lepton selection. The fit distributions are shown in Figures 6.13 and 6.14.

The centers of the Lorentzian fit results, shown in Table 6.3, were compared between simulation and data. To evaluate the systematic, the mean lepton p_T of the signal was scaled by the ratio of the data and simulated means for electrons and muons separately, and the shift in the measured mass results was taken as a systematic (the statistical uncertainties on these fit results are negligible). Clearly, the disagreements in the electron results dominate this uncertainty.

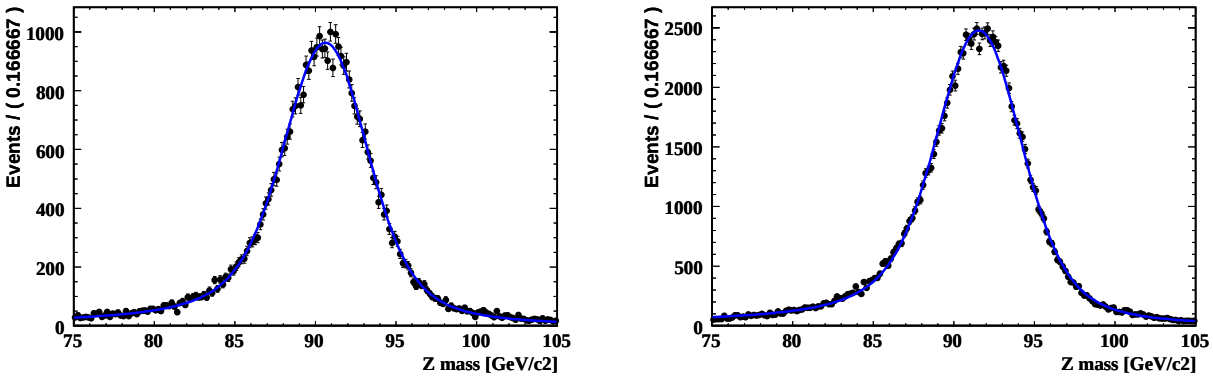

(a) (b)

Figure 6.13: a: fit to the dilepton mass peak for electrons in data. b: fit to the Z mass peak for electrons in simulation.

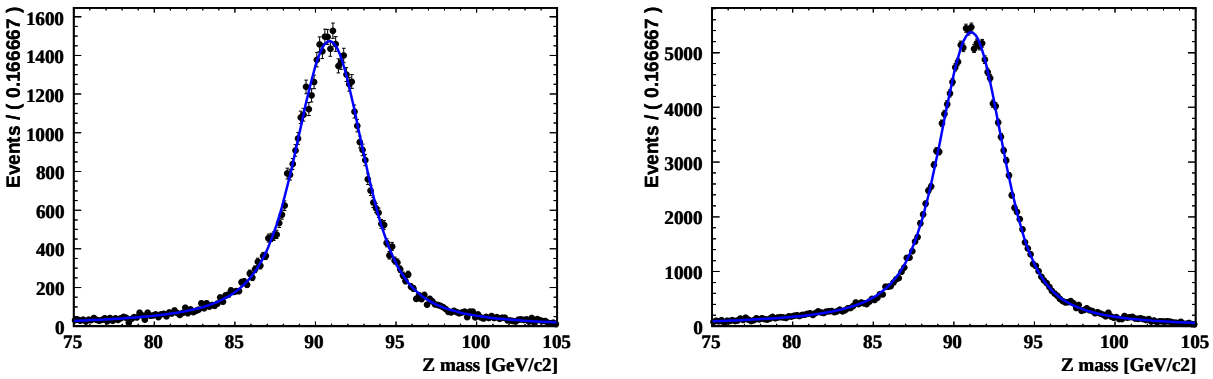

(a) (b)

Figure 6.14: a: fit to the dilepton mass peak for muons in data. b: fit to the Z mass peak for muons in simulation.

Table 6.3: Centers of the Lorentzians from the fits to the Z peaks shown in Figures 6.13 and 6.14

Sample	Mass Peak (GeV/c^2)
Muon Simulation	91.16
Muon Data	90.96
Electron Simulation	91.81
Electron Data	90.84

6.5.6 Decay length related uncertainties

The procedure for calibrating the decay length measurements in our signal sample is described in section 6.3.2. There are many uncertainties which must be considered for this calibration, some due to the modeling of b-jets, and others due to the track jet energy measurements that are used to parameterize the calibration.

The decay length calibration has a statistical limitation due to the data and PYTHIA $b\bar{b}$ sample sizes. This uncertainty is folded directly into the pseudoexperiments as explained in section 6.4, but its contribution is quite small. There are a number of uncertainties on the photon plus jet energy calibration technique. The energy scale calibration curve has an associated statistical uncertainty which propagates through to a mass uncertainty. A systematic uncertainty of 1% is taken on the measured energy of the photon in the simulation, which corresponds to a 1% uncertainty on the measured track jet transverse momentum. Finally, as described in [22], about 30% of the photon plus jets sample in data is com-

posed of QCD dijet production where one of the QCD jets fakes a very clean photon signature through a pion or lambda decay. This contamination has been determined to have a momentum balance discrepancy compared to the photon plus jets signal at the 1% level, and so an additional 0.3% uncertainty is taken on the track jet momentum. Finally, after applying the procedures described in section 6.3 to minimize (in the case of $b\bar{b}$) or correct (in the case of $t\bar{t}$) for jet overlap and underlying event effects, the mass is reevaluated using cone size 0.7 track jets and the shift in the mass results is taken as an additional systematic uncertainty to represent out-of-cone and jet overlap uncertainties.

Another track jet energy uncertainty arises in connection with the simulation of the b-jets. If the EvtGen decay tables do not produce the correct distributions of charged particles then this will artificially bias any measurements of the tracking energy of the b-jets. The DELPHI Collaboration has measured the charged decay multiplicity of b-hadrons [12], excluding the decay products of long lived light-flavor particles and of excited b-hadron to ground state b-hadron transitions, and determined an average of 4.97 ± 0.07. We evaluate this number at generator level in our samples using the same exclusions and arrive at a mean result of 5.047. This discrepancy is very slightly larger than the reported error at DELPHI, and it cannot be explained by uncertainties in the production fractions of different b-hadron types or on the semileptonic decay rate. Under the

assumption that excess tracks will be distributed randomly in the b-hadron rest frame, this discrepancy should directly translate into an equivalent discrepancy on the measured energy of the component of the track jet originating from the b-hadron decay. This leads to an additional 1.1% uncertainty on the measured track jet energies.

In addition to jet energy effects, other uncertainties are considered in relation to the modeling of the physics of the $b\bar{b}$ sample. It is important to minimize and understand any charm contamination in this sample. My $p_{T,rel}$ fits demonstrate that the muon jets in the sample are about 95% likely to be b-jets, and 5% likely to be charm jets. But they also suggest that the simulation slightly underestimates the number of charm jets. I apply a small reweighting to correct for this deficit, parameterized according to jet transverse energy as shown in Figure 6.5.6. I take the fit to the ratio as a charm flavor scale factor that I use to correct the simulated Lxy values before determining the Lxy scale factor. I determine the uncertainty on this charm scale factor by varying the charm fraction within its one-sigma fit uncertainties, arriving at the new correction results shown in Figure 6.5.6. Though the flavor variations are a few percent in scope, the resulting systematic uncertainty is much less than this due to the similarity between the b- and c-jet decay lengths. These effects lead to a small systematic uncertainty.

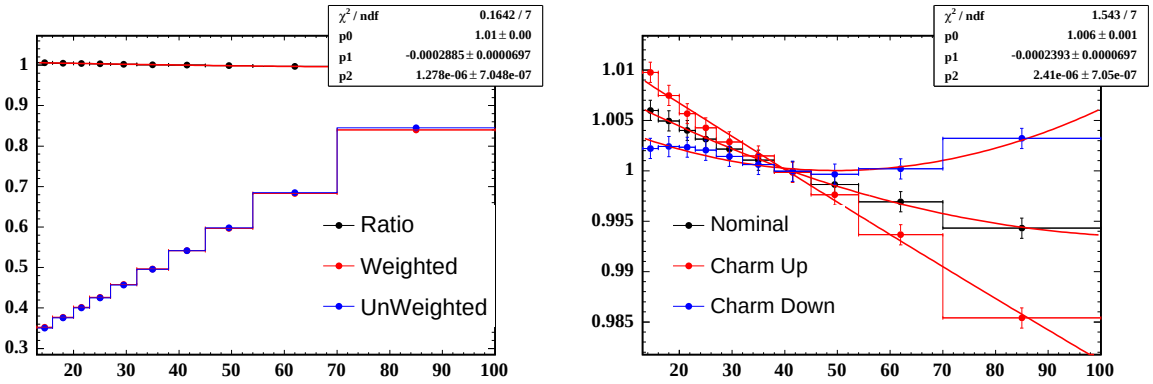

(a) (b)

Figure 6.15: a: The mean Lxy as a function of jet transverse energy for the simulated jets in the scale factor sample using the simulated flavor distributions (red), and after correcting to the charm fraction found in data (blue). Their ratio determines the correction factor that will be applied to the simulation. b: There are minor variations in the ratio when the charm fraction is varied within the uncertainty of the $p_{T,rel}$ fits. The error bars on this plot are artificially set to help the fits converge, and have no physical meaning.

Another small uncertainty arises from the event selection on the muons in the leptonic $b\bar{b}$ sample. If the simulation does not properly model the measurement of the muon momentum then higher or lower energy muons (corresponding to higher or lower decay length vertices) will pass selection. While the Z peak fits above suggest a very accurate modeling of isolated muons, we conservatively take a 1% uncertainty on the muon momentum scale, and evaluate the mass shift that results from the new set of events passing selection. Finally, it is important to understand the uncertainties in PYTHIA's modeling of the b-quark fragmentation. In PYTHIA, the energy carried by the hadron after the fragmentation process is modeled with the Bowler function. As explained in Section 5.3.4, the D0 Collaboration has studied LEP and SLD data [5] and determined the PYTHIA tune required to reproduce their results. Samples of $t\bar{t}$ events were generated according to each of these tunes, and the resulting b-hadron energies were found to be about 2% higher than under the default PYTHIA tune. As expected, this results in a proportionally larger mean decay length of the signal b-hadrons. However, since this effect also occurs in the $b\bar{b}$ samples, I expected the effects to compensate. To test this, I generated $b\bar{b}$ Monte Carlo samples with the same alternate fragmentation tunings, and sure enough, the Lxy fluctuations cancel one another out, illustrating the motivation for the Lxy calibration procedure. The fragmentation fluctuations do, however, produce the following minor fluctu-

ations which are not canceled out. When events are reweighted to the alternate fragmentation distributions it causes small alterations to the measured track jet energies and raises the muon energy distribution slightly. Accounting for all of these effects, I took the larger of the mass shifts between the default PYTHIA sample, and the results after reweighting to the SLD or LEP results as a fragmentation systematic uncertainty.

A summary of calibration systematic uncertainties for the decay length measurement is shown in Table 6.4.

Table 6.4: Calibration based top quark mass uncertainties for the decay length measurement.

Systematic [GeV/c^2]	Lxy	Lepton p_T	Simultaneous
Photon Plus Jet Stats	0.7	0	0.3
Photon Pt	1.4	0	0.6
Photon Background	0.4	0	0.2
Track Jet Cone Size	0.8	0	0.3
N_{trk} from b-hadrons	1.6	0	0.7
Semilep Muon Pt	0.3	0	0.1
Fragmentation	0.6	0	0.3
$c\bar{c}$ Background	0.2	0	0.1
Total SF	2.5	0	1.1

6.5.7 Multiple interactions uncertainty

There are two respects in which other interactions during a beam crossing may result in a systematic bias. These extra interactions are simulated as overlaid minimum bias events, however this modeling may not be accurate, resulting

in biased jet energies. I describe these effects in Section 6.5.8 for calorimeter-based jet measurements, and in Section 6.5.6 for the tracking based jet measurements.

However another effect must be taken into account. The simulation is tuned to an older dataset corresponding to an integrated luminosity of 1.2 fb^{-1}. The newer data included in the measurement were collected at higher instantaneous luminosities with more interactions per bunch crossing than the earlier data. To study this effect I used a specially generated high luminosity $t\bar{t}$ sample with top quark mass 175 GeV/c^2. I then segregated the events according to the number of collision vertices that are reconstructed in the event (which has been shown to be approximately proportional to the luminosity [22]). While there is no statistically significant trend over number of vertices for the Lxy measurement, there is a significant dependence for the lepton momentum measurement. For electrons, a small part of this trend is due to particles from other collisions falling into the electron cluster, however the primary cause of this effect is the isolation requirement for the leptons. Since we require the total calorimeter energy found around the lepton to be less than 10% of the letpon momentum, low momentum leptons are more likely to fail selection in high luminosity events, as illustrated in Figure 6.16.

Like the simulation, we segregate the data based upon the number of reconstructed collision vertices. As expected, the data have about 15% more reconstructed collision vertices per event than the standard $t\bar{t}$ samples. I reweighted the high luminosity simulated events to reproduce the distribution of the number of reconstructed vertices in both the standard $t\bar{t}$ samples and the data in turn. These two reweighted results are equivalent to each other within statistics for the Lxy measurement, however the lepton transverse momentum is significantly higher under the luminosity profile of the data. There are insufficient statistics at very high luminosities to reliably correct for this effect in all of our signal $t\bar{t}$ samples. Instead I take the differences between the associated top quark mass results using the luminosity profiles of the data and the simulation for our high luminosity sample as a systematic uncertainty for each measurement. I should emphasize that this uncertainty is due to the simple logistics of the luminosity profile that was used in our simulation, and is not due to any irreducible physics effect. For the decay length measurement the statistical uncertainty due to the number of generated $t\bar{t}$ events is larger than the observed systematic shift, and so this statistical uncertainty is taken as the systematic instead.

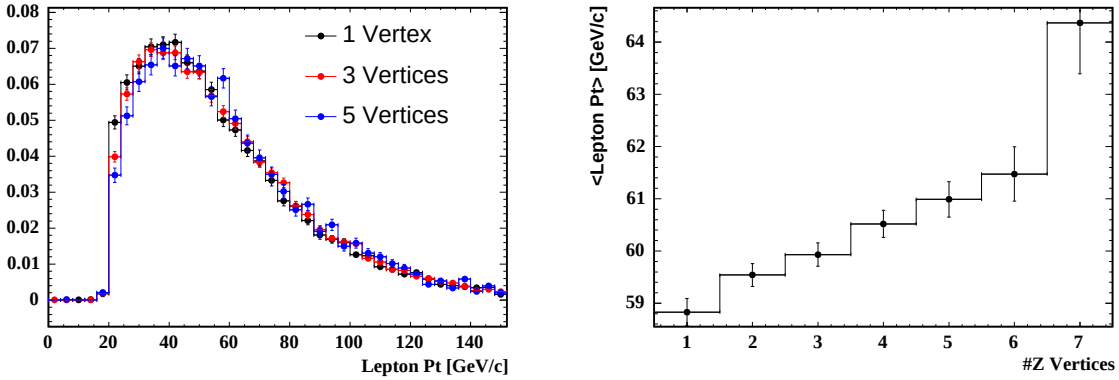

(a) (b)

Figure 6.16: Effects of luminosity on the lepton momentum, evaluated at generator level. Any luminosity dependence is due to the isolation requirements in the lepton selection. a: Generator-level lepton transverse momentum. The higher the luminosity (as measured by the number of vertices), the smaller the number of low energy leptons that pass the isolation requirement. b: The mean lepton transverse momentum trend.

6.5.8 Jet energy uncertainties

The jet energy uncertainties can be broken down into two categories: those arising from the tracking energy measurements which impact the Lxy calibration, and those from the calorimeter measurements that are common to all three analyses. As I explained in Section 5.3, none of the uncertainties represent an uncertainty on the determination of a "true" jet energy. Rather, they represent uncertainties in the modeling of jet energy measurements in simulation. In this section I discuss the evaluation of these uncertainties and explain why they have

minimal correlation with the calorimeter based uncertainties that are claimed by other top quark mass analyses.

The dominant jet energy uncertainties in this analysis arise from the track jets. They are listed in Table 6.4 as Photon Plus Jet Stats, Photon Pt, Photon Background, Track Jet Cone Size, and N_{trk} from b-hadrons. Of these, the Photon Pt (energy bias in our calibration photons) and N_{trk} from b-hadrons (EvtGen decay multiplicity mismodeling) categories are the largest contributions to the jet energy uncertainties. The third largest uncertainty is due to the limited statistics in our photon plus jets data and will not present any difficulty in future high statistics analyses.

Table 6.5: Calorimeter Based Jet Energy Uncertainties. The residual uncertainties result from possible inaccuracies in the cancellation that occurs for our out-of-cone uncertainty.

Systematic [GeV/c^2]	Lxy	Lepton p_T	Simultaneous
Eta Dependent	0.06	-0.08	-0.02
Multiple Interactions	0.17	-0.01	0.07
Calorimeter Response	-0.14	-0.07	-0.09
Underlying Event	0.09	-0.06	0.01
Splash Out	0.15	-0.10	0.02
Base Out of Cone	0.18	-0.28	-0.06
Out of Cone Residual Uncertainties			
bb Semileptonic	0.24	NA	0.24
W plus charm/LF	0.14	0.22	0.30
Final Out of Cone	0.33	0.36	0.24
Tot Calorimeter JES	0.44	0.39	0.32

The fourth largest uncertainty is due to the size of the cone used to construct our track based jets. This uncertainty may have components from a wide variety of physical effects, but minimal correlation to the jet energy uncertainties of other analyses. The most significant component comes from an uncertainty in the overlap of particles from other jets falling into the jet cone, for which no corresponding uncertainty is claimed for calorimeter jets, and thus will remain uncorrelated to the results of other measurements as long as this effect continues to be neglected. The only systematic components for which there are any correlations to calorimeter-based uncertainties are the much smaller underlying event, multiple interaction uncertainties, and out-of-cone uncertainties. Most of the multiple interaction contributions are vetoed by the z-vertex matching requirement. As for the out-of-cone uncertainty, it should be minimal due to the application of our photon calibration procedure. It will only contribute to the extent in which the simulation models out of cone effects in b-jets with a different level of accuracy compared to the light flavor jets on which the calibration is performed. To summarize, of all the track jet energy uncertainties only a small part of the 0.8 GeV/c^2 (0.3 GeV/c^2) systematic uncertainty on the Lxy (combined) measurements that is due to the altered cone size could have any correlation to the calorimeter based jet energy scale uncertainties we are claiming, or to those of other analyses.

As I discussed in Section 5.3, the calorimeter-based uncertainties can be split into six categories (counting splash-out), which are assumed to be independent of one another and are described in [22]. Since these are the same categories into which the jet energy scale corrections are split, these uncertainties are sometimes called jet energy scale uncertainties. As for the tracking based uncertainties, their impact on the decay length measurement arises based on which jets pass our event selection thresholds. These low energy jets which pass in and out of our sample as the jet energies are varied within uncertainties tend to have a small decay length, and thus do contribute a bias to the average decay length of our sample. Unlike track jets, however, for calorimeter jets this effect is present in both our $b\bar{b}$ calibration and our main analysis samples, and it largely cancels in the final mass determination. Such cancellations are what motivated the choice of the Lxy calibration procedure. To evaluate these uncertainties, I fluctuate the calorimeter energies of the jets within these six categories and reevaluate the missing energy of the event, keeping track of which jets and events pass selection. I take the resulting mass shifts as calorimeter-based systematic uncertainties.

Of course, my primary criticism of recent top mass measurements is their application of procedures that mostly cancel out their primary systematic uncertainties without ever accounting for uncertainties on these uncertainties (for example, due to the different jet shapes based upon flavor, which I highlighted

in Section 5.4.1). The same concerns can be applied to the cancellation of my jet energy uncertainties from my calibration technique. This concern would specifically pertain to the out-of-cone jet energy uncertainty, which dominate for the low energy jets that pass in and out of our selection. If the out-of-cone disagreement between simulation and data is different for the jets in our $b\bar{b}$ sample than for the jets in our $t\bar{t}$ sample, then by chance the cancellation may lead to an artificially small systematic result.

If jets near the selection threshold were to have identical properties for the $b\bar{b}$ and $t\bar{t}$ samples, then disagreements between data and simulation would be identical for the samples and the resulting systematic cancellation resulting from assuming identical out-of-cone uncertainties would be appropriate. Fortunately, the differences are small. For our purposes, the only relevant differences between the jets in our samples are that the $b\bar{b}$ jets used in our decay length calibration are required to contain muons, and that some of the tagged jets in our backgrounds are light flavor or charm. In all other respects the simulated jets we use near our selection threshold are similar. There are two systematic cross checks that I ran to address these concerns. The results will be shown at the end of this section.

As I explained in Section 5.3.2, the out-of-cone jet energy uncertainties are parameterized based upon the calorimeter energy measurement. Since lower energy jets are broader, a lower jet energy corresponds to a larger out of cone

Table 6.6: Final Systematic Uncertainties. The Lxy Calibration systematic uncertainty is summarized in Table 6.4. The Calorimeter JES systematic uncertainty is summarized in Table 6.5.

Systematic [GeV/c^2]	Lxy	Lepton p_T	Simultaneous
Background Shape	1.0	2.3	1.7
QCD Radiation	0.5	1.2	0.7
PDF	0.3	0.6	0.5
Generator	0.7	0.9	0.3
Lepton p_T Scale	0	2.3	1.2
Lxy Calibration	2.5	0	1.1
Multiple Interactions	0.2	1.2	0.7
Calorimeter JES	0.4	0.4	0.3
Systematics Total	2.9	3.8	2.6

uncertainty. Since the muon's energy is mostly lost for jets in the $b\bar{b}$ sample, it can be argued that we are overestimating the out-of-cone uncertainty for these jets. To check the impact this would have, we add the muon's energy back in and repeat our systematics evaluation. As a second check, I consider the fact that -jets and light flavor jets have different degrees of mismodelling of the out-of-cone uncertainty, see Figure 5.7. I check the shifts of the uncertainties which occur when I fluctuate the size of our out-of-cone uncertainties for charm and light flavor jets relative to b-jets within conservative constraints. The fluctuations from these two cross checks are taken as residual out-of-cone systematic uncertainties, and I take the full out-of-cone systematic as the quadrature sum of direct and residual out-of-cone uncertainties. The results are summarized in Table 6.5. The

systematic uncertainties from all effects that have been considered are shown in Table 6.6.

Chapter 7

The Wrap

For my thesis analysis I have used 1.9 fb^{-1} of integrated luminosity to measure a top mass of $166.9^{+9.5}_{-8.5}$ (stat) $\pm$ 2.9 (syst) GeV/c^2 using the mean decay length of b-jets, $173.5^{+8.8}_{-8.9}$ (stat) $\pm$ 3.8 (syst) GeV/c^2 using the mean transverse momentum of the leptons from W-decay, and 170.7 $\pm$ 6.3 (stat) $\pm$ 2.6 (syst) GeV/c^2 using both variables simultaneously. Compared to the original Lxy measurement [14], I have run over almost three times the integrated luminosity, and added in a second, uncorrelated variable with similar top mass sensitivity to the analysis that effectively adds another factor of two to the statistics. At the same time, I have improved the systematic uncertainty on the measurement by a factor of 3.3. Thus, the results remain limited by statistics and so will improve significantly by the end of Run II if they are continued at

CDF. Further, if a measurement of this type were performed at the LHC, the systematic uncertainties would be the true limitation as the statistical uncertainties would be negligible.

So those are the official results for this round of the analysis. How much better can they be? I can reliably predict the extent to which the systematics can be easily improved at the Tevatron. If the Monte Carlo were regenerated using the proper luminosity profile the multiple interactions uncertainty would become negligibly small. And if the preliminary lepton momentum calibration techniques described in [25] are applied this systematic would decrease by more than a factor of two. The systematic uncertainties for these measurements at CDF would then be expected to drop to 2.9 GeV/c^2 for the decay length measurement, 3.0 GeV/c^2 for the lepton transverse momentum measurement, and 2.3 GeV/c^2 using both variables simultaneously.

7.1 Caveats

One of the points I have tried to make in this writeup is that potentially significant uncertainties are often overlooked in analyses. In Section 5.4, I explained my reasons for doubting the uncertainties in conventional top mass measurements. Even given all the publications that have used these procedures and the

careers that have been launched with them, I still think that the uncertainties are untrustworthy. Do I really think that I do not have similar mistakes that a careful reader would find in my own analysis?

There are two uncertainties that I can think of that I have not considered that could potentially play a significant role. One is from the use of leading order Monte Carlo simulations. I have seen how a leading order parton distribution function can produce a bias on the order of a couple of GeV/c^2 in my results. I have corrected for this effect, but there are still rare missing categories of interactions involving higher order loops that are not simulated, and that could potentially bias the transverse momentum of the top decay products. Like other analyses, I have attempted to verify that these biases are insignificant by running over a specially generated NLO $t\bar{t}$ sample with a next to leading order PDF. Unfortunately, I also discovered that these samples, though generated with millions of events, used a broken random number generator that was only capable of generating 32 thousand unique events. I doubt that any of these effects are significant, but no one is sure that top mass analyses (not mine, nor the official ones) are not biased from unaccounted for NLO interactions, because there are not enough statistics to be sure of this. In order to evaluate this effect I would have to dig into the code, find and fix the error, and generate a large $t\bar{t}$ sample to run over.

A second possible source of bias comes from "color reconnection" inaccuracies. When bare quarks are produced, how accurately does the simulation model color flow between them to produce hadrons? This was not an issue that I ever thought about until after I had already almost finished my analysis. I had always thought that the Lxy calibration procedure using the $b\bar{b}$ sample would compensate for any overlooked uncertainties such as this one (Section 6.3.2). But color flow is presumably quite different in $t\bar{t}$ and $b\bar{b}$ events, so I cannot be sure that the calibration procedure will properly compensate for these inaccuracies. Today, extensive studies are being done and new, specially generated $t\bar{t}$ Monte Carlo samples have been generated to quantify the size of this uncertainty for standard mass measurements. The next generation of top mass analyses will quote a color reconnection systematic. If, unlike other top mass analyses of my generation, I wanted to quantify the size of this uncertainty, I would have to generate special color-reconnected samples for my $b\bar{b}$ Monte Carlo, and repeat the Lxy calibration with the alternate $b\bar{b}$ samples, and the measurement with the alternate $t\bar{t}$ samples. As for the leading order approximation above, I suspect the size of this effect is small, but I cannot be sure without doing a lot of work to prove it.

In the end, I feel that I have done a solid analysis for my thesis. I believe that its results are more correct than those of other modern top mass publications,

but in order to fully convince that there are no significant oversights I would have to complete the studies I have mentioned above. If I were to perform another round of this analysis, these studies would be on the top of my to-do list.

Bibliography

[1] The ALEPH Collaboration and others, hep-ex/0811.4683 (2008).

[2] Tevatron Electroweak Working Group, hep-ex/0903.2503 (2009).

[3] The CDF Collaboration, CDF Note 9692, http://www-cdf.fnal.gov/physics/new/top/2009/mass/mtm3_p19_public/.

[4] The CDF Collaboration, CDF Note 9694, http://www-cdf.fnal.gov/physics/new/top/2009/mass/TMT_AH_P19_public/TMT_AH_3fb.html.

[5] W. Feborko, PhD. thesis, University of Chicago, FERMILAB-THESIS-2008-42 (2008).

[6] More explicitly, the missing transverse energy is defined as $|-\sum_i E_T^i \vec{n}_i|$, where the sum is over calorimeter towers with transverse energy deposits, E_T^i, and $\vec{n}_i$ is the unit vector from the collision vertex to the tower in the plane transverse to the beam direction. Further corrections are applied to

account for muon tracks and to account for calorimeter non-uniformity and non-linearities.

[7] D. Sherman, Ph.D. Thesis, Harvard University (2007).

[8] M. R. Whalley et al., hep-ph/0508110 (2005).

[9] P. Stephan, hep-ex/0902.0169 (2009).

[10] T. Aaltonen et al. *Phys. Rev. D*, 73:032003, 2006.

[11] F. Abe et al. *Phys. Rev. Lett.*, 80:2767–2772, 1998.

[12] P. Abreu et al. *Phys. Lett. B*, 425:399–412, 1998.

[13] A. Abulencia et al. *Phys. Rev. D*, 71:052003, 2005.

[14] A. Abulencia et al. *Phys. Rev. D*, 75:071102, 2005.

[15] A. Abulencia et al. *Phys. Rev. Lett.*, 97:082004, 2006.

[16] A. Abulencia et al. *Phys. Rev. D*, 73:032003, 2006.

[17] A. Abulencia et al. *J. Phys. G*, 34:2457–2544, 2007.

[18] D. Acosta et al. *Phys. Rev. D*, 71:052003, 2005.

[19] I. J. R. Aitchison and A. J. G. Hey. IOP Publishing Ltd, 1989.

[20] E. Barberio and Z. Was. *Comput. Phys. Commun.*, 79:291, 1994.

[21] R. Barlow and C. Beeston. *Comp. Phys. Comm.*, 77:219–228, 1993.

[22] A. Bhatti et al. *Nucl. Instrum. Meth. A*, 566:375–412, 2006.

[23] I. Blumenfeld et al. *Nature*, 445, 2007.

[24] R. Cahn. *Rev. Mod. Phys.*, 68:951, 1996.

[25] The CDF Collaboration, 2009. CDF Conf. Note 9683.

[26] G. Corcella et al. *J. High Energy Phys.*, 0101:010, 2001.

[27] A. Amsler et al. (Particle Data Group). *Physics Letters B*, 667:1, 2008.

[28] D. Griffiths. John Wiley & Sons, Inc., 1987.

[29] F. Halzen and A. D. Martin. John Wiley & Sons, Inc., 1984.

[30] F. Hartmann. Evolution of silicon sensor technology in particle physics, 2009.

[31] C. S. Hill, J. R. Incandela, and J. M. Lamb. *Phys. Rev. D*, 71:054029, 2005.

[32] J. Incandela et al. Status and prospects of top-quark physics. *Prog. Part. Nucl. Phys. to be published 2009.*

[33] H. L. Lai et al. *Eur. Phys. J. C*, 12:375–392, 2000.

[34] D. J. Lange. *Nucl. Instrum. Meth. A*, 462:152–155, 2001.

[35] F. Maltoni and T. Stelzer. *J. High Energy Phys.*, 0302:027, 2003.

[36] M. L. Mangano et al. *J. High Energy Phys.*, 0307:001, 2003.

[37] A. D. Martin et al. *Phys. Lett. B*, 605:61–68, 2004.

[38] J. Pumplin et al. *J. High Energy Phys.*, 07:012, 2002.

[39] J. Pumplin et al. *J. High Energy Phys.*, 02:032, 2006.

[40] P. Sigmund and A. Achinner. *Nucl. Instrum. Meth. B*, 212:110, 2003.

[41] T. Sjostrand et al. *Computer Phys. Commun.*, 135:238, 2001.

[42] M. Srednicki. Cambridge University Press, 2007.

[43] M. C. Thompson et al. *AIP Conf. Proc.*, 877:903–909, 2006.

Appendix A

A Glimpse of Quantum Field Theory

This section is intended to provide the reader with some idea of how a Lagrangian can be used to make meaningful predictions. Let us consider what is quite possibly the simplest physical interaction in the Standard Model: a process where a spinless charged particle (such as a pion) absorbs a photon, see Figure A.1. This interaction is actually unphysical since it does not conserve momentum, but it is simple enough that it allows some tricks to make the complicated mathematics of Quantum Field Theory a bit more manageable in a small space.

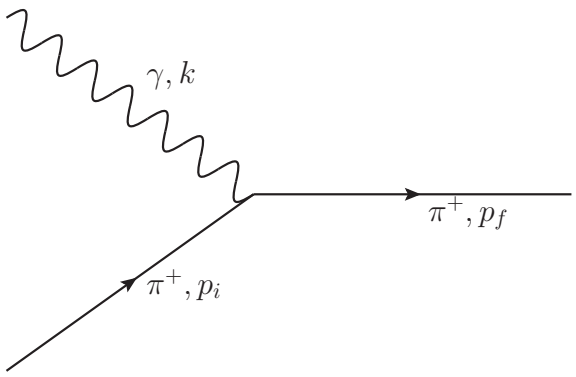

Figure A.1: A simple Feynman diagram for the gamma absorption process I am considering. In reality this process can only be part of a larger diagram that conserves momentum and energy through some additional interaction such as photon emission.

The "derivation" of the probability amplitude that I will show is not at all rigorous, but hopefully it provides a qualitative sense of how Standard Model calculations can be performed. To do this calculation in a formally correct manner one needs to delve deeply into Quantum Field Theory. Also note that I am taking this example from Chapter 5.2 of [19], but I think my explanation will be significantly more compact and conceptually complete, despite skipping most of the steps of the calculations.

If this were a Quantum Mechanics problem we could start with Schrodinger's Equation to determine the time evolution of the state of the system,

$$\hat{H}|\psi(t)>= i\frac{d}{dt}|\psi(t)> \tag{A.1}$$

Here we are working with the relativistic hamiltonian $\hat{H}$, which is composed of both free particle and interaction terms. However, since a precise calculation of the probability amplitude for the above scattering process would be mathematically impossible using this full formalism, we instead apply a perturbative expansion, keeping only the lowest order term:

$$A_0 = -i \int dt < f|e^{i\hat{H}_0 t} \hat{H}_{int} e^{-i\hat{H}_0 t}|i > \tag{A.2}$$

Here, A_0 is our lowest order approximation of the probability amplitude, $|i>$ and $|f>$ are the initial and final quantum states, $\hat{H}_{int}$ represents the interaction Hamiltonian, and $\hat{H}_0$ the free particle Hamiltonian (which in this approximation will govern the time evolution of the states or operators). But this equation has a significant flaw, which is that it is incompatible with Special Relativity. We must modify this expression to put time and space on an equal footing by writing it in terms of the Hamiltonian density, $\hat{\mathcal{H}}(\vec{x}, t)$, where $\hat{H} = \int d^3x \hat{\mathcal{H}}(\vec{x}, t)$:

$$A_0 = -i \int d^4x < f|\hat{\mathcal{H}}_{int}(\vec{x}, t)|i > \tag{A.3}$$

Now our task is to determine an expression for $\hat{\mathcal{H}}$. If we were dealing with a quantum probability amplitude (ϕ) instead of a field, the Hamiltonian density could be found from the Lagrangian density, $\mathcal{L}$, which would be required to obey the relativistic equation:

$$\frac{\partial \mathcal{L}}{\partial \phi} - \partial_\mu \left(\frac{\partial \mathcal{L}}{\partial(\partial_\mu \phi)}\right) = 0 \tag{A.4}$$

Here, $\partial^\mu = (\frac{\partial}{\partial t}, -\vec{\bigtriangledown})$ is the relativistic generalization of the gradient operator from vector calculus. We would also need to require this equation to agree with the expectations for a relativistic spinless massive particle. In particular if the particle were free it would need to reproduce the Klein Gordon Equation: $m^2\phi + \partial_\mu \partial^\mu \phi = 0$. The necessary Lagrangian density can be solved to be

$$\mathcal{L}_0 = \frac{1}{2}\partial_\mu \phi \partial^\mu \phi - \frac{1}{2} m^2 \phi^2 \tag{A.5}$$

However, we are not really dealing with amplitudes, we want to be dealing with quantum fields. The fields will be complex, and will thus have two degrees of freedom (related to particle and antiparticle), $\hat{\phi} = \frac{1}{\sqrt{2}}(\hat{\phi}_1 - i\hat{\phi}_2)$, and $\hat{\phi}^\dagger = \frac{1}{\sqrt{2}}(\hat{\phi}_1 + i\hat{\phi}_2)$, for real fields $\hat{\phi}_1$ and $\hat{\phi}_2$. The generalized free particle Lagrangian is then:

$$\hat{\mathcal{L}}_0 = \partial_\mu \hat{\phi}^\dagger \partial^\mu \hat{\phi} - m^2 \hat{\phi}^\dagger \hat{\phi} \tag{A.6}$$

However, our particles are not free, they interact according to the electromagnetic potential, $\hat{A}^\mu$. If we take this as the only source of potential (the potential is free), it must obey the Maxwell Equations, $\partial_\mu \hat{F}^{\mu\nu} = 0$, where

$\hat{F}^{\mu\nu} = \partial^\mu \hat{A}^\nu - \partial^\nu \hat{A}^\mu$. We are also allowed to enforce the Lorentz gauge, $\partial_\mu \hat{A}^\mu = 0$. Additionally, a key assumption of the Standard Model is that we can enforce "local" gauge invariance at every point in spacetime. That is, we will require that the Lagrangian be invariant under $\hat{\phi}(x) \rightarrow e^{i\alpha(x)}\hat{\phi}(x)$ for any function $\alpha(x)$ (with a similar requirement on $\hat{A}^\mu$). In order to make all of these conditions hold we must modify the derivative as follows $\partial^\mu \rightarrow \hat{D}^\mu = \partial^\mu + ie\hat{A}^\mu$, where e will represent the electric charge of the particle. These gauge invariance requirements are how interactions are created in the theory. As a result, the Lagrangian of Equation A.6 is no longer free, it needs to have interaction terms added to it:

$$\hat{\mathcal{L}}_{int} = -ie(\hat{\phi}^\dagger \partial^\mu \hat{\phi} - (\partial^\mu \hat{\phi}^\dagger)\hat{\phi})\hat{A}_\mu + e^2 \hat{A}^\mu \hat{A}_\mu \hat{\phi}^\dagger \hat{\phi} \tag{A.7}$$

We can then determine the associated interaction Hamiltonian density from this Lagrangian for a single photon interaction:

$$\hat{\mathcal{H}}_{int} = ie(\hat{\phi}^\dagger \partial^\mu \hat{\phi} - (\partial^\mu \hat{\phi}^\dagger)\hat{\phi})\hat{A}_\mu = \hat{j}^\mu_{em}\hat{A}_\mu \tag{A.8}$$

Where we have dropped the term of second order in charge since we are dealing with only the leading order expansion, and $\hat{j}^\mu_{em}$ represents the electromagnetic probability current. We can then Fourier expand our fields:

$$\hat{\phi}(\vec{x},t) = \int_{-\infty}^{\infty} \frac{d^3\vec{k}}{(2\pi)^3 2w} \{\hat{a}(k)e^{-ikx} + \hat{b}^\dagger(k)e^{ikx}\} \tag{A.9}$$

Here, the field, $\hat{\phi}$ is expanded in terms of the raising and lowering operators $\hat{b}^\dagger(k)$ ($\hat{b}(k)$), which can be interpreted as creating (destroying) an anti particle with four-momentum $k = (\vec{k}, w)$, and $\hat{a}(k)$ ($\hat{a^\dagger}(k)$), which creates (destroys) a particle with four-momentum k. Mathematically these operators behave the commutation relations:

$$[\hat{a}(k), \hat{a}^\dagger(k')] = [\hat{b}(k), \hat{b}^\dagger(k')] = (2\pi^3)2w\delta(k-k') \tag{A.10}$$

A similar expansion can be made for the electromagnetic potential operator:

$$\hat{A}^\mu = \sum_\lambda \int_{-\infty}^{\infty} \frac{d^3\vec{k'}}{(2\pi)^3 2w'} \{\epsilon^\mu(k',\lambda)\hat{\alpha}(k',\lambda)e^{-ik'x} + \epsilon^{*\mu}(k',\lambda)(\hat{\alpha}^\dagger(k',\lambda)e^{ik'x}\} \tag{A.11}$$

Where the $\hat{\alpha}^\dagger$ ($\hat{\alpha}$) creates (destroys) photons, and ϵ is the complex polarization vector for the photon. The Lorentz Gauge condition and the Maxwell Equation constraints allow us to fix this polarization vector to be perpendicular to the photon's direction of propagation and to have no time component.

For the diagram of Figure A.1, we have a final state of $|f>= \hat{a}^\dagger(p_f)|0>$, and an initial state of $|i>= \hat{a}^\dagger(p_i)\hat{\alpha}^\dagger(k,\lambda)|0>$. This allows us to write out the full lowest order equation, A.3, as:

$$A_0 = -i \int d^4x < 0|\hat{a}(p_f)\hat{\mathcal{H}}_{int}\hat{a}^\dagger(p_i)\hat{\alpha}^\dagger(k,\lambda)|0 > \tag{A.12}$$

From Equation A.8 we can break this expression into two terms, one for the photon, and one for the pion:

$$< 0|A_\mu \alpha^\dagger(k,\lambda)|0 >, \; < 0|\hat{a}(p_f) j^\mu_{em} \hat{a}^\dagger(p_i)|0 > \tag{A.13}$$

Substituting in the Fourier expansions, A.9 and A.11, it is just a tedious calculation to solve for these expressions. The answers are given by

$$N_k \epsilon_\mu e^{-ikx}, \quad eN_f N_i (p_i + p_f)^\mu e^{-i(p_i - p_f)x} \tag{A.14}$$

For the photon and electron terms respectively. Here the N's are simple normalization terms, and the integrals will reduce the complex exponentials to delta functions that enforce momentum and energy conservation. From this very simple calculation, a number of general rules might be guessed at. For example, we might conclude that a normalization factor, N, is needed for each pion in the diagram, a polarization factor, ϵ, is needed for each photon in the diagram, a factor of $ie(p_i + p_f)^\mu$ is needed for each intersection between a pion and a photon, and a four-momentum conserving delta function must be applied to the final result.

These rules only apply in a very simple case like this with no virtual particles, no spin, and no interactions beyond electromagnetism, but they illustrate the basic principles that are used to build up the mathematical foundation of the Standard Model and many of the speculative physics scenarios that might extend it. First setup up the standard equations of quantum field theory to solve for the states you care about, constraining the fields to obey the laws of physics that you want in your theory (the Maxwell Equations for electromagnetism and gauge invariance in the above case). It will be necessary to expand these equations to some order of approximation. Each of the terms in your approximation can be thought of as a diagram where each of the factors is an interacting particle. After solving enough problems like this it quickly becomes clear that a formal set of rules (called Feynman rules) can be set up whereby someone can skip the first few steps and instead write down all the diagrams that could contribute to the transition between $|i> \rightarrow |f>$ at a given order of approximation, translates the lines, vertices, propagators, and loops into equations according to the appropriate Feynman rules, and solve the equations without going through the tedius initial steps.

For a Leading Order (LO) approximation like the one above, your mathematical expression will have at most two charge factors, meaning that your diagrams will have at most one internal propagator. Most physics simulations at collid-

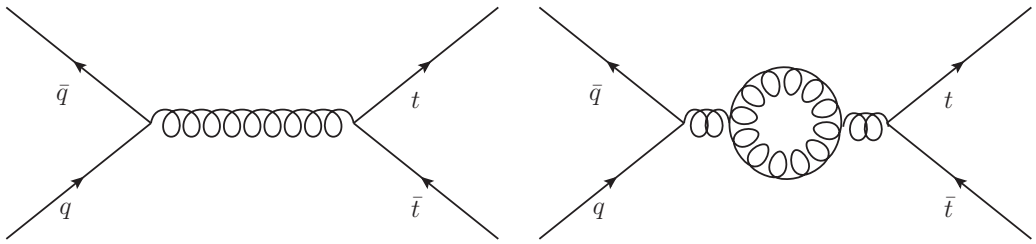

Figure A.2: Feynman diagrams for some processes by which $t\bar{t}$ is produced. The left is an example of a leading order process, and in fact is the most common channel for $t\bar{t}$ production. The right is an example of a next to leading order process, and is not modeled in any commonly used simulation.

ers content themselves with relying on LO calculations. By keeping an extra term in the perturbative expansion (next to leading order, NLO), further layers of interactions will occur, which give rise to many more complicated diagrams with an extra internal line or a loop, and the problem will require much more challenging calculations. If your perturbative approximation is good (meaning that the coupling is small), the contribution of these more complicated diagrams to your final probability distributions will be quite small. There are contexts, however, where this higher level of accuracy is needed to give results with the required level of accuracy. My thesis concentrates on the $t\bar{t}$ production mechanism, which proceeds through the much more complicated interactions that are governed by the strong force rather than those of Electromagnetism (explained shortly). Diagrams for examples of LO and NLO $t\bar{t}$ production processes are shown in Figure A.2.

Appendix B

The Philosophy of Uncertainties

Scientists are expected to publish measurement results along with one-sigma uncertainties (or 90/95% confidence intervals). When dealing with systematic uncertainties, however, it is usually impossible to determine what a one-sigma variation is. Usually you can only do some cross checks under different assumptions and see how much your answer changes. There often will be no clear answer for how you should extract uncertainties from the cross checks. Other times, you just can't think of any good cross checks to run that will not give you a grossly overestimated upper bound on your uncertainty. What can you do in those cases? Do you ignore the potential uncertainty, or do you quote a greatly exaggerated systematic to be sure you are not underestimating anything? How sure do you really need to be that an uncertainty is not underestimated?

Is it better to publish results that claim one-sigma uncertainties that are more likely to be two sigma bands or to risk an underestimation? These are questions of philosophy and honesty, and there are wide variations in how authors and collaborations decide to answer them. But there is a constant pressure to publish with smaller uncertainties than the competing analyses. As long as your results are not obviously discrepant with previous results and your procedures are not obviously wrong, then you can often get your result into journals with little critical review.

Fortunately, most scientists are very honest people, and some (perhaps most) experiments err on the side of being too conservative. Personally I don't see a problem with this. Uncertainty bands that are too large will lead researchers to work harder to be sure of things before coming to conclusions. Uncertainty bands that are too small can lead them to jump to false conclusions. And there are a minority of cases where flagrantly wrong results make their way into official publications. For several decades, the Particle Data Group (PDG) [27] has been taking published measurements of physical quantities and combining them into world average results and uncertainties. In each edition they take a look back at their historical results to determine some of their published quantities that have changed in surprising ways from one edition to another. I am copying their latest plots of the most discrepant quantities into Figure B.1. Some of

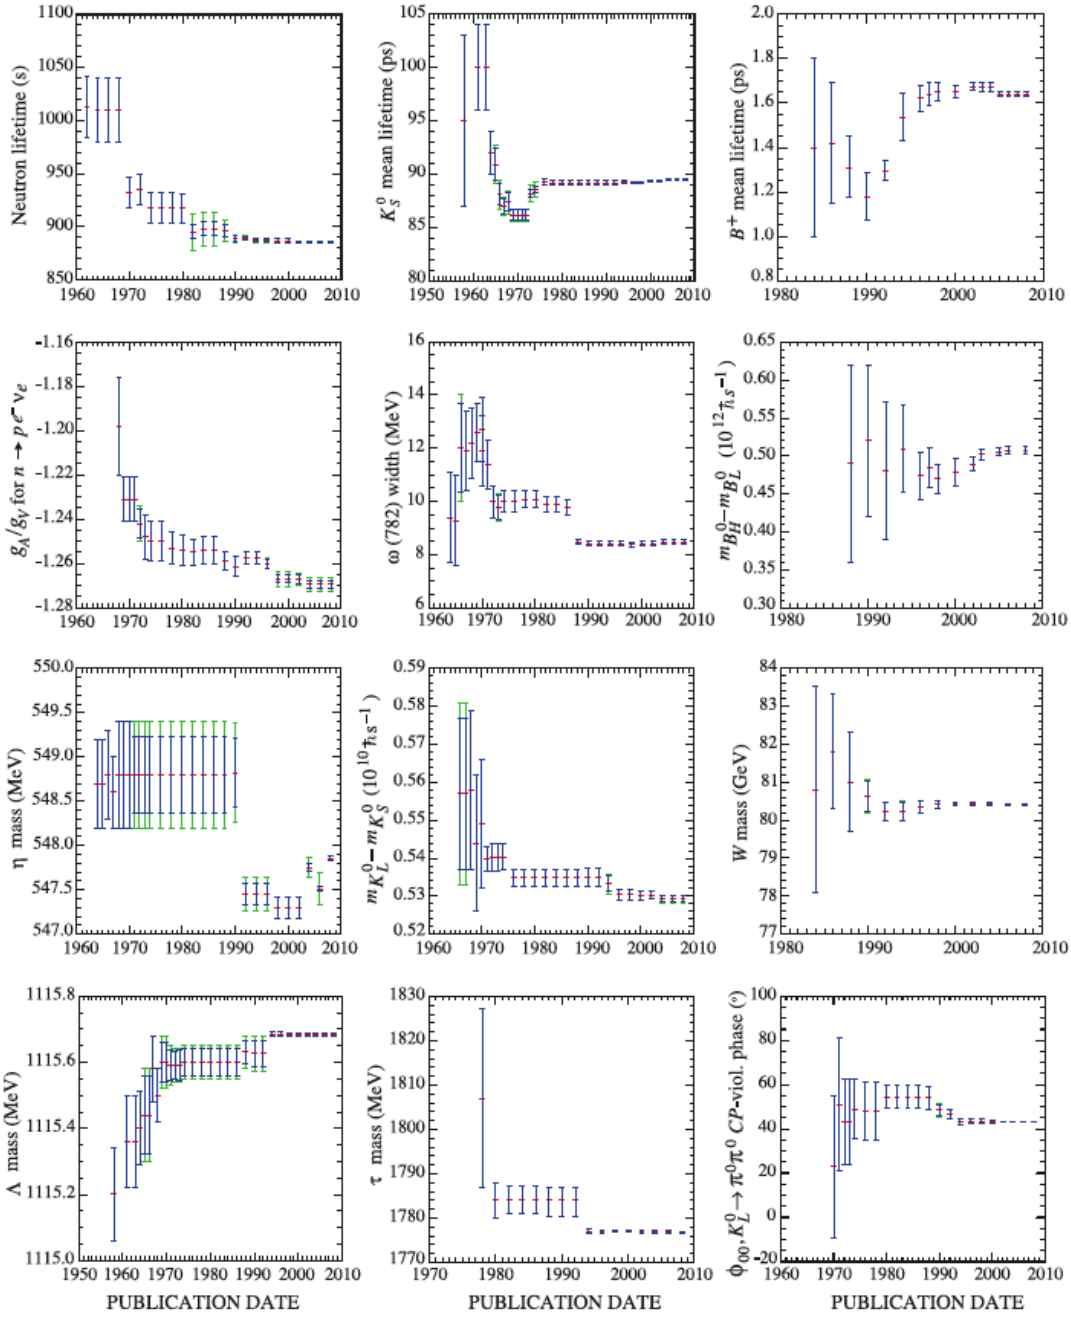

Figure B.1: The world averages of some physical quantities have changed dramatically between different editions of the PDG.

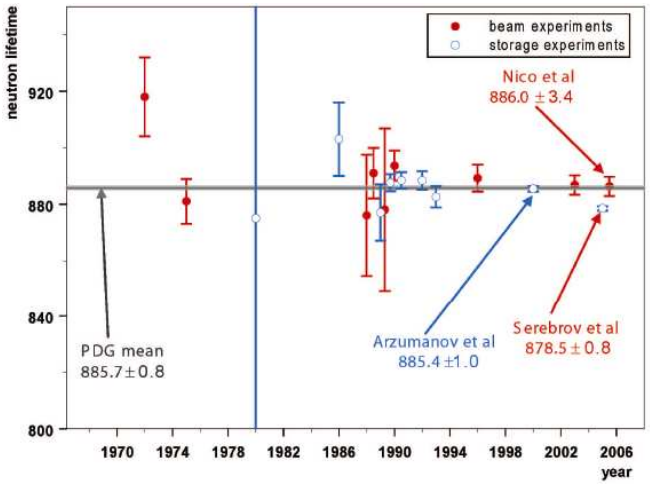

Figure B.2: Even after all of the past problems have been exposed, the most precise measurement of the neutron mass is not even close to being compatible with the world average.

these fluctuations are indicative of vastly greater oversights than anything I have been worried about for the top mass measurements, where the world average measurement has needed to be revised by several standard deviations.

Past embarrassments may encourage authors to be more careful, but not always enough to prevent further embarrassments in the future. Some of these quantities have turned out to be wrong, and have clearly been corrected to a very different value that has also proven to be badly wrong. Consider the neutron lifetime measurements. There is a major modern embarrassment that does not even show up in this figure that is detailed in reference [9]. Even after all of the past mistakes, if the recent results of Serebrov et. al are correct then the PDG

world average neutron lifetime is still wrong by more than six sigma as shown in Figure B.2. I am sure the authors of the neutron lifetime measurements are all honest people, but these trends should serve as a warning against the power of wishful thinking and illustrate the need to critically examine results that you want to be correct.