GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN

II. Physikalisches Institut

Selected Methods of High Energy Physics Experiments at Hadron Colliders

von

Adam Roe

An assortment of methods used to verify, monitor and calibrate components of detectors used at collider accelerators will be described here. Monitoring of the *DØ* calorimeter and its level-1 trigger will be discussed. Verifying and installing the readout electronics for the CMS cathode strip chambers in the endcap muon system will be described. First data taken with these chambers, using cosmic muons, are analyzed and their comparison with Monte Carlo expectations are shown. Measuring the precise geometry of the ALFA detector at ATLAS, necessary for a precise luminosity measurement, is explained.

Post address:
Friedrich-Hund-Platz 1
37077 Göttingen
Germany

II.Physik-UniGö-Dipl-2008/04
II. Physikalisches Institut
Georg-August-Universität Göttingen
September 2008

GEORG-AUGUST-UNIVERSITÄT

GÖTTINGEN

II. Physikalisches Institut

Selected Methods of High Energy Physics Experiments at Hadron Colliders

von

Adam Roe

Dieser Forschungsbericht wurde als Diplomarbeit von der Fakultät für Physik der Georg-August-Universität zu Göttingen angenommen.

Angenommen am:	1 September 2008
Referent:	Prof. Quadt
Korreferent:	Prof. Klute

II.Physik-UniGö-Dipl-2008/04

Contents

1 Introduction

The work done by physicists at detectors built at accelerators has the same ultimate goal as that done by any other scientist: to understand the patterns of nature more fundamentally than had been previously. This requires verifying the work of previous scientists, testing and falsifying theories and looking into unexplored realms. Experimental high energy physics is no exception to this. Within the field, the main methods used are experiments associated with a particle accelerator and those which rely on astrophysical processes for their data. As the physics is the same here on earth as in outer space, many of the methods used are similar. In this paper, selected methods of high energy physics, specifically those by detectors used at accelerators, will be explored.

There are several types of accelerators used and accordingly a variety of apparatuses employed to record the data. There are fixed target experiments where a particle beam is collided into a stationary object and there are colliding beam experiments. Beams have been made of helium nuclei (α), electrons and positrons ($e^{\pm}/\beta^{\pm}$), photons (γ), protons and antiprotons (p and $\bar{p}$), muons (μ), neutrinos (ν) as well as of other, composite particles.

The examples used in this paper are taken from experiments at hadron colliders, colliding protons with antiprotons or protons with protons. The experiments at other types of accelerators, such as lepton colliders or proton-electron colliders, use similar methods. The composite nature of the proton means that the exact collision conditions are not known at a hadron collider while the higher mass of the proton (with respect to the electron) means that higher energies can be achieved without great loss to synchrotron radiation. This trade off means that lepton and hadron colliders necessarily complement one another. While the examples in this paper are taken from experiments at hadron colliders, the methods employed are similar in any particle detector whether at a lepton or hadron collider or on a satellite.

An overview of the current status of particle physics will be given. This will include a brief description of the Standard Model of Particle Physics followed by a few of the open questions still posed, questions which motivate the field today. The physical processes underlying the detection of particles at high energy physics experiments will be described, namely how particles interact with matter at this scale. A chapter is then dedicated to the accelerators and experiments pertinent to the rest of the methods discussed here: Fermilab's Tevatron accelerator and CERN's Large Hadron Collider (LHC), and the corresponding experiments of $D\emptyset$ at the Tevatron as well as CMS and ATLAS at the LHC. Three examples of methods used will be explored in depth. The first is the Level-1 calorimeter trigger at the $D\emptyset$ experiment, with emphasis on the monitoring of the upgraded system used in the Tevatron's Run IIb. Next is the endcap muon system at CMS, comprised of cathode strip chambers, with emphasis on the processes of commissioning and calibration of the chambers. This is accompanied by an analysis of first data taken with the chambers using muons of cosmic origin. The final method discussed will be used to determine the absolute calibration scale for measuring luminosity delivered to the ATLAS experiment using a dedicated detector known as ALFA. Following these sections, a brief summary and outlook will be given.

2 The Standard Model of Particle Physics

The Standard Model of Particle Physics uses quantum field theory to describe the phenomena observed on a sub-microscopic scale in terms of three forces: Electromagnetic, Weak and Strong. Each of the three are described by a quantum field theory, and the fields of the Electromagnetic and Weak forces are unified in a single Electroweak field at high energies.

The Standard Model contains quarks, leptons - both being fermions - and force mediators, the vector bosons. This is shown graphically in Table 2.1. Quarks interact strongly, weakly and electromagnetically, while the charged leptons only interact weakly and electromagnetically. The uncharged leptons, the neutrinos, only interact weakly. The mediator of Electromagnetism is the massless photon, γ. There are three massive mediators of the Weak force, the two charged $\mathrm{W}^{\pm}$ and the uncharged Z^0. There are eight massless gluons mediating the strong force, which each carry two color charges. The Electroweak theory also necessitates the existence of a symmetry breaking mechanism, done so via the Higgs boson in the Standard Model. Direct experimental evidence is lacking for this elusive particle, a topic of much interest in the physics community today. For a more detailed explanation of the Standard Model, see for instance [1, 2, 3, 4].

Quarks			**Mediators**
$\begin{pmatrix} u \\ d \end{pmatrix}$	$\begin{pmatrix} c \\ s \end{pmatrix}$	$\begin{pmatrix} t \\ b \end{pmatrix}$	γ
			g
Leptons			Z^0
$\begin{pmatrix} \nu_e \\ e \end{pmatrix}$	$\begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix}$	$\begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix}$	W^+
			W^-

Table 2.1: The particles of the Standard Model, with fermions shown in term of their isospin doublets. As direct experimental evidence for the Higgs boson is lacking, it is not included here.

The theory of Quantum Electrodynamics (QED) was the first successful quantum field theory. It was proposed in order to understand the behavior of Electromagnetism in a quantum framework which also preserved causality by forbidding action at a distance. Developed by many people over many years, the 1965 Nobel Prize in Physics was shared amongst Richard Feynman, Julian Schwinger and Sin-Itiro Tomonanga for their profound contributions. QED is considered to be the most precise scientific theory which exists. It is characterized by α, the dimensionless electromagnetic coupling constant, which is defined in SI as $\alpha = e^2/4\pi\epsilon_o\hbar c$, where $\alpha \approx 1/137$ in any units at low energies. The field described by QED is now understood as a lower energy manifestation of the Electroweak field.

In the mid 1960's, work was done to unify the Weak force with the Electromagnetic force [5, 6, 7]. Since the 1930's, Enrico Fermi's 4-point theory of the Weak interaction had been used, but had also been understood to only be an effective theory: it diverged at sufficiently high energies. The proposal was to make the Weak force mediated by a massive vector boson. Letting the "charge" of the Weak force be approximately the same as that of the Electromagnetic force, they achieved Electroweak unification, where the combined force is described by four boson fields: one isospin triplet and one singlet. The central point of the theory is that the Weak

force is not weak (with respect to Electromagnetism) because the coupling is small but rather because the mass of the mediators, the $W^{\pm}$ and Z^0, are large. The early successes of the theory led to the 1979 Nobel Prize in physics being given to Glashow, Weinberg and Salam for their contributions to the field. The theory's main outstanding feature is a lack of discovery of the Higgs boson, the quanta of the Higgs Fields, necessary to break Electroweak symmetry [8, 9]. It should also be noted that the recent discovery of neutrino mass does not fit into the theory well. The recently discovered oscillations between neutrino flavor and mass eigenstates [10] can be easily incorporated to the Electroweak theory, however the implication of this is that neutrinos are massive; it is extremely difficult to include neutrino mass in the Standard Model theory and most theoretical solutions involve complicated extensions of the theory.

The theory of Quantum Chromodynamics (QCD) has been developed to explain the Strong force [11, 12]. In QCD there are eight quanta, the gluons, mediating the Strong force between quarks. Gluons have no electric charge but they do carry a color charge; they are also massless. The most notable physical feature of the theory is asymptotic freedom: at higher energies (or smaller distances) the interaction weakens and thus the quarks become "free". This leads to what is known as "infrared divergence". At lower energies, such as the rest energy of the nucleons, quarks are confined; a free quark is never observed. In order to account for infrared divergence and quark confinement, in QCD the coupling constant, α_s, "runs": it is a function of energy. This is also true of the electromagnetic coupling constant, α, but running is a more prominent feature of QCD. The strong coupling constant is best measured at the Z pole, $\alpha_s(M_Z) \approx 0.18$. QCD is also co-linear divergent, meaning that the theory does not coherently predict interactions where outgoing particles have a small angle between their momenta.

The mathematics of group theory is used in order to exploit the symmetries of the theories used to describe nature. The Standard Model is described by $\mathrm{SU}(3)_C \otimes \mathrm{SU}(2)_L \otimes \mathrm{U}(1)_Y$, corresponding to the symmetry of color in the Strong force, the symmetry of left-handed isospin partners in the Weak force and the symmetry of hypercharge, which the Electromagnetic force takes on when it is unified with the weak force into the Electroweak theory. The symmetry of the Electromagnetic force becomes charge after the Electroweak force's symmetry is spontaneously broken by the Higgs mechanism, creating the two apparently separate forces: Electromagnetic and Weak. In terms of group theory, this symmetry breaking can be envisioned as $\mathrm{SU}(2)_L \otimes \mathrm{U}(1)_Y \rightarrow \mathrm{U}(1)_q$ [13].

The Standard Model describes an extremely wide range of phenomena, but there are still open questions. Some of these questions are:

- Is the Higgs the Electroweak symmetry breaking mechanism or is there something else?
- How can neutrino mass be accounted for?
- What is the nature of Dark Matter, whose existence is strongly implied from astrophysical models?
- Can the Electroweak theory be unified with QCD at higher energies?
- What accounts for nucleon spin?

These are some of the questions which drive particle physics today. These issues and others are the subject of many analyses at hadron collider experiments and other high energy physics experiments.

3 The Framework of Experimental High Energy Physics

The use of scattering to probe matter at atomic and sub-atomic distances has been used since the famed Gold Foil experiment of 1909 conducted by Hans Geiger and Ernest Marsden under Ernest Rutherford. In the Gold Foil experiment a beam of alpha particles was used as an incoming beam against a stationary gold target. Such fixed-target experiments dominated the field for more than fifty years. The development of finer accelerator techniques allowed for colliding beam experiments to be used in place of a single beam incident upon a stationary target. Since their inception in the 1970's, colliding beam experiments have included the e^+e^- colliders LEP at CERN and SLAC at Stanford University, the $p\bar{p}$ colliders $Sp\bar{p}S$ at CERN and Fermilab's Tevatron, as well as the eP collider Hera at DESY. The new Large Hadron Collider (LHC) at CERN, opening the summer of 2008, will be a pp collider with a center of mass energy of $\sqrt{s} = 14$ TeV. Such colliding beam accelerators and the detectors built there will be the topic of this section. The examples examined in detail in this thesis come from detectors at such accelerators as well.

3.1 Physical Parameters at Colliders

At a colliding beam experiment, the accelerator determines the basic parameters of the entire experimental apparatus. The most important parameters that depend on the accelerator are the center of mass energy of the collision, $\sqrt{s}$, the luminosity, L, and the interaction rate, R. The interaction cross section, σ, is a measure of the likelihood of interaction and Γ is the branching ratio of a specific process. These factors, σ and Γ, are important measurable quantities which are predicted by theory and measured at detectors.

Center of Mass Energy

The relativistically invariant four-momenta[1] squared of a system of two colliding particles, $p^2 \equiv p^\mu p_\mu$, is:

$$p^2 = (\vec{P}_1 + \vec{P}_2)^2 - (E_1 + E_2)^2 = -m_1^2 - m_2^2 + 2\vec{P}_1\vec{P}_2 - 2E_1E_2. \tag{3.1}$$

The system in which the total three momentum is 0 is the center of mass system (cms), such that its total energy, $E^{*2} = -p^2$. This is the same as one of the Mandelstam Variables, denoted by s. For the case of a stationary target experiment with one energetic beam, let E_2 represent the target's energy such that $E_2 = m_2$ and E_1 is the incoming beam's energy. In the cms,

$$E^{*2} = -p^2 = m_1^2 + m_2^2 + 2m_2E_1 \tag{3.2}$$

In a symmetric colliding beam experiment, which will be subject of the rest of this thesis, there are two beams of the same type of particles with the same beam energy and equal opposite

[1] A metric with $g_{00} = -1$ and $g_{11} = g_{22} = g_{33} = +1$ is used here, following [1].

momenta ($\vec{P}_1 = -\vec{P}_2$, $E_1 = E_2 = E$),

$$E^{*2} = -p^2 = m_1^2 + m_2^2 + 2(E_1E_2 + P_1P_2) \simeq 4E^2. \tag{3.3}$$

Thus, in a colliding beam experiment the available energy goes as 2E whereas in a fixed-target experiment $E^* \simeq \sqrt{2m_2E_1}$. There are clear energetic advantages - which come with technical difficulties - for a beam-beam collider. This paper will consider the DØ experiment at the Tevatron's $p\bar{p}$ collider with $\sqrt{s} = 1.96$ TeV and the ATLAS and CMS experiments at the LHC, a pp collider with $\sqrt{s} = 14$ TeV.

The Parton Nature of the Proton

When protons collide into one another only a fraction of the total energy of the proton participates in the collision. This is due to the parton nature of the proton: it is not a point-like particle but is composed of constituents. The fraction of momentum carried by each parton is called *Bjorken x*. The proton appears as a point-like particle at low energies, such as in the Gold Foil experiment. In this case there is one parton and x=1. Higher energies resolve that there are three quarks, two up type and one down type, which make up the proton. These are called "valence" quarks, which each carry approximately 1/3 of the proton's momentum. At higher energies still, the interactions between these quarks are resolved in the form of gluons. Gluons of high energy can split into $q\bar{q}$ pairs, creating "sea quarks". The result is that as energies get higher the number of partons in the proton rises and the fraction of momentum carried by each falls. This is a result of the asymptotic freedom of quarks. At high enough energies the valence quarks are no longer bound as they were at lower energies and are on near-equal footing as sea quarks and gluons. This is parameterized by a Parton Distribution Function, F. In general, $F(x, Q^2)$ where Q^2 is the proton's total momentum and x is the fraction carried by the parton.

Cross Section

In classical mechanics, the interaction cross section, σ, corresponds to the area of the particles involved in an interaction. In a quantum situation it implies a probability of interaction; it still corresponds to "area" however in a quantum sense. Thus the cross-sectional area of a particle is different with respect to the different forces. The total cross section for various processes is plotted below. In the case of a proton, the cross section has a flat minimum at a value of about $\sigma_{pp} \simeq \pi(2R_p)^2 \simeq 12 \times 10^{-30} m^2 = 1.2$ mbarn. As shown in Fig. 3.1, $\sigma = \sigma(\sqrt{s})$, the total cross section is a function of the center of mass energy for any process. At higher energies there are more interacting partons, thus increasing the total cross section.

Luminosity

The instantaneous luminosity, L, is defined as [1]

$$L = fn\frac{N_1N_2}{A}\ cm^{-2}s^{-1}. \tag{3.4}$$

Here, N_1 and N_2 are the number of particles in each bunch, n is the number of bunches in each beam, f is the revolution frequency and A is the cross-sectional area of the beams. Luminosity on the order of $\sim 10^{32}$ cm$^{-2}s^{-1}$ are achieved at the Tevatron, and $\sim 10^{34}$ cm$^{-2}s^{-1}$ are expected at the LHC.

We can also define an Interaction Rate, $R = \sigma L$, in the units of Hz. For an individual process, a cross section is extracted by counting the rate at which a process occurs and measuring the

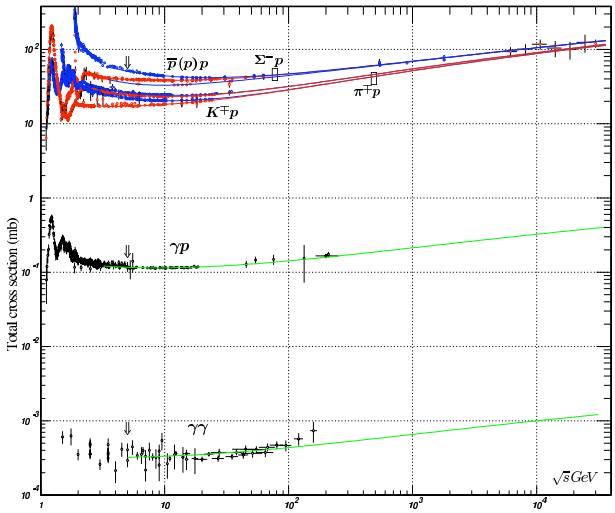

Figure 3.1: Total cross sections of various processes as a function of $\sqrt{s}$.

luminosity. At a hadron collider, luminosity measurement is not trivial and will be discussed in Chapter 7. A total interaction rate is also an important quantity, as it defines the speed at which a detector, its trigger and readout must function. This will be on the order of 1 GHz for the LHC, presenting an impressive technical challenge. Integrated Luminosity is a measure of the complete data set, $\mathcal{L} = \int Ldt$. The Tevatron hopes to deliver 8 fb^{-1} by 2009, and the LHC expects $\sim$ 8 fb^{-1} per year.

Branching Ratio and Process Cross Section

The Branching Ratio of a process defines the likelihood that a specific state will result. Thus, $\Gamma_i \times \sigma_{tot} = \sigma_i$ for some specific process i. Naturally, $\sum \sigma_i = \sigma_{tot}$, so $\sum \Gamma_i = 1$ by definition. The determination of Γ_i for various processes is one of the foremost goals of collider experiments.

Detector Coordinates

The standard coordinate system for a detector at an accelerator is defined in polar coordinates. It is generally right-handed system with z on the beamline and either x or y pointing towards the center of the accelerator. Coordinates are defined as $r = \sqrt{x^2 + y^2}$, $\phi = \tan(\frac{x}{y})$ and $\eta = -\ln\tan(\frac{\theta}{2})$.

3.2 Accelerators

The principal method of studying high-energy collisions is to accelerate particles such as protons, electrons or ions, collide the beam into another beam or a stationary target, and record the outcome. Following the earlier argument concerning $\sqrt{s}$, two colliding beams produces the highest possible center of mass energy. Such colliding beam accelerators will be considered here, as this is the method used at Fermilab's Tevatron and CERN's LHC. Accelerator complexes are staged using various types of accelerators. While the largest accelerator is ultimately a synchrotron, various other methods such as radio frequency (RF) cavities and electrostatic potentials are used as early stages in the accelerator complex. A basic outline of a contemporary high-energy colliding beam accelerator is as follows: particles are extracted from their source via

an electrostatic potential - such as separating protons and electrons from hydrogen gas - and are then accelerated through a combination of electrostatic potentials (such as a Cockroft-Walton) and RF Cavities (such as an Alvarez Linac) and then injected into a synchrotron for acceleration up to collision energies. The methods and structure of a proton synchrotron are considered here, while the exact structure at the Tevatron and LHC are considered later.

Synchrotron Accelerators

The main principle of a synchrotron is to accelerate a beam of particles using RF cavities while increasing the magnetic field to keep the particles moving on the same circular path. This synchronized increase of the magnetic field with the energy of the particles gives the type of accelerator its name. In general, the magnetic field must rise linearly with the proton momentum, by $p = 0.3B\rho$, where p is the proton momentum and ρ is the ring's radius. This is achieved using dipole bending magnets. The use of superconducting magnets in a synchrotron was first demonstrated at Fermilab in the early 1980's, allowing for much higher magnetic fields - and therefore beam energies - to be reached than with conventional dipoles. In a $p\bar{p}$ collider, such as $Sp\bar{p}S$ or Tevatron, a single vacuum in a magnetic field with two beams moving in opposite directions is used. Each beam exhibits helical motion, so the pair of beams weave a double-helix. At the LHC a far more complicated magnetic field with two separate vacuua is necessary because the collisions are pp. The difficulties of using $\bar{p}$ will be discussed below.

A second set of focusing magnets are used in order to collimate the beams for collision at interaction points. This is generally done with sets of quadrupole magnets around the collision point. Each quadrupole magnet creates a magnetic field with a minimum in one direction normal to the plane of beam motion that has a strong gradient, focusing the beam in one plane. Magnets alternate between orthogonal planes, thus focusing the beam [1]. Finer focusing can be achieved with higher n-pole magnets; octupole magnets, for instance, will be used at the LHC.

For a proton collider (pp or $p\bar{p}$) groups or "bunches" of protons, rather than individual protons, are collided. Colliding two such bunches significantly increases interaction probability over colliding protons one-on-one. The exact dynamics of each interaction are not known as a result of the parton structure of hadrons. The momentum fraction x carried by an individual parton cannot be known a priori at high energies. Accordingly, analyses at proton colliders exploit symmetry in the $(r,\ \phi)$ (i.e. $x - y$) plane: while the beam is moving in the z-direction it is at rest in the $(r,\ \phi)$ plane. A practical result of the parton nature of the proton is that at LHC energies it is expected that the difference between pp and $p\bar{p}$ will be minimal, as the momentum fraction carried by the valence quarks becomes small as the beam energy rises. The difficulties of working with antiproton beams are considered worse than the trouble of having two vaccua with opposite dipole bending, and accordingly the LHC is designed to be a pp, not a $p\bar{p}$, machine.

Antiprotons at Accelerators

Antiprotons are unstable when in contact with normal (i.e. non-anti) matter. Accordingly they must be created in a lab in $p\bar{p}$ pairs and then extracted in order to create a beam, which circulates in an extremely tight vacuum. The first time this was done was for the $Sp\bar{p}S$ accelerator at CERN in the 1980's [15]. A proton beam struck a target, then the antiprotons were selected using a magnetic and optical filtration system. The process works but is very inefficient: At the $Sp\bar{p}S$ a proton beam with bunches consisting of $\sim 10^{13}$ protons of E $\approx$ 26 GeV incident on a Cu target yielded bunches of $\sim 10^7$ antiprotons with E $\sim$ 3.5 GeV [1]. The resulting $\bar{p}$ beam was widely spread in phase space. The necessity of reducing the $\bar{p}$ phase space led to the invention

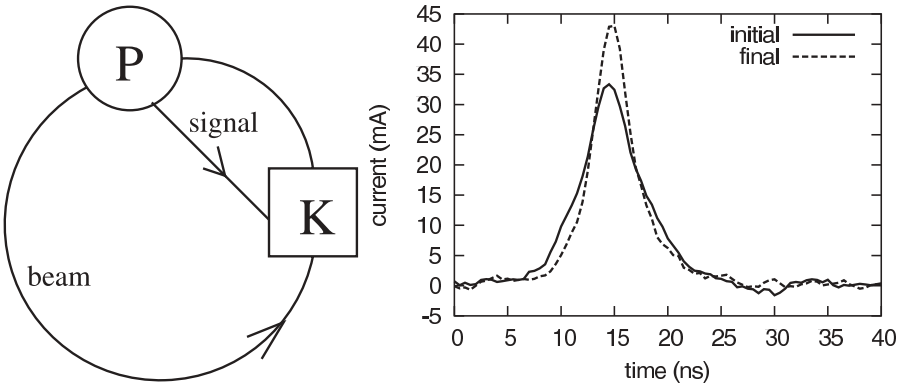

(a) A schematic of a stochastic cooling ring. The current is read at the pickup and a correction is applied at the kicker.

(b) The effects of Stochastic Cooling at RHIC. The beam pulse is shown to be quicker and stronger after cooling.

Figure 3.2: Stochastic Cooling. Images from [14].

of stochastic cooling, leading to half of the 1984 Nobel Prize in physics being given to Simon van der Meer.

In stochastic cooling, the current produced by a $\bar{p}$ bunch is measured by a "pickup". The signal is sent on a chord across the accelerator ring, shown schematically in Fig. 3.2 (a). As the signal propagates it is processed and an adjustment is delivered to the $\bar{p}$ bunch by a "kicker". In the antiproton storage ring used at the $Sp\bar{p}S$, two seconds of cooling resulted in a decrease in transverse and longitudinal spreads of the beam by an order of magnitude. The effectiveness of stochastic cooling can be clearly seen in Fig. 3.2 (b), showing the current before and after cooling at RHIC in a ring presently used for heavy ions.

3.3 Detection Methods

Detectors are built around the interaction points of the colliding beams of an accelerator. The ways in which particles interact with matter and the physics goals of an experiment drive the detector implementation. As described earlier, there are four known forces affecting particles: Strong, Weak, Electromagnetic and Gravitational. Particles which interact Electromagnetically are generally detected by ionization and excitation, as well as by γ radiation. Particles interacting via the Strong nuclear force are detected via gluon-mediated scattering. These interactions are used in tracking, calorimetry and muon systems. Particles which only interact weakly or gravitationally, such as neutrinos, are not detected directly.

3.3.1 Particle Interactions in Matter

The Electromagnetic interaction of charged particles moving through matter can be split into three energy regimes where different effects dominate. At low energies, e.g. $\beta\gamma \lesssim 1$, poorly understood nuclear effects dominate and the average energy loss in a material (the "stopping power") goes as $1/\beta^2$. This is marked as the Anderson-Ziegler range in Fig. 3.3. The well studied Bethe-Bloch Formula describes stopping power over a wide momentum range:

$$-<\frac{dE}{dx}>= Kz^2\frac{Z}{A}\frac{1}{\beta^2}[\frac{1}{2}ln\frac{2m_ec^2\beta^2\gamma^2T_{max}}{I^2}-\beta^2-\frac{\delta(\beta\gamma)}{2}], \tag{3.5}$$

where $K \simeq 0.307$ MeV g^{-1} cm^2, T$_{max}$ is the maximum amount of energy transferable to an electron, I is the mean excitation energy, and $\delta(\beta\gamma)$ parameterizes the density effect. The range in which the Bethe-Bloch Formula is accurate is within the gray bands in Fig. 3.3.

Minimum Ionization, Density and Radiative Effects

The density effect is a relativistic effect whereby the electric field of a charged particle seems flattened and therefore extended in the rest frame of the material, causing an increased probability of interaction. The material becomes polarized which then truncates this effect. The result is that $\delta(\beta\gamma) \simeq \ln(\beta\gamma)$, which approximately cancels the dominant log rise in Bethe-Bloch, causing the "Fermi Plateau". Any particle in this wide, flat range is called a minimum ionizing particle (MIP). Strictly speaking, the point of minimum ionization occurs at $\beta\gamma \simeq 3.5$, but practically speaking any particle on the plateau is considered a MIP. The plateau is very large for muons, ranging over 3 orders of magnitude in $\beta\gamma$, which means that muons leave little signal in most detection systems and require special treatment.

In the plot shown in Fig. 3.3, radiative effects take over for $\beta\gamma \gtrsim 1000$. These effects are not included in the Bethe-Bloch formula. Radiative effects take over for different particles in various media at different energies, however the general form of the graph is the same. Energy loss, with and without the density effect, are plotted in dashed lines which do not take radiative effects into account. It should be noted that while radiative energy losses dominate over ionization losses at high energies, this generally occurs for muons at such high energies that the radiative losses are small compared to the particle's energy. For the situation shown in Fig. 3.3, energy loss from a 1 TeV muon will be on the order of tens of MeV.

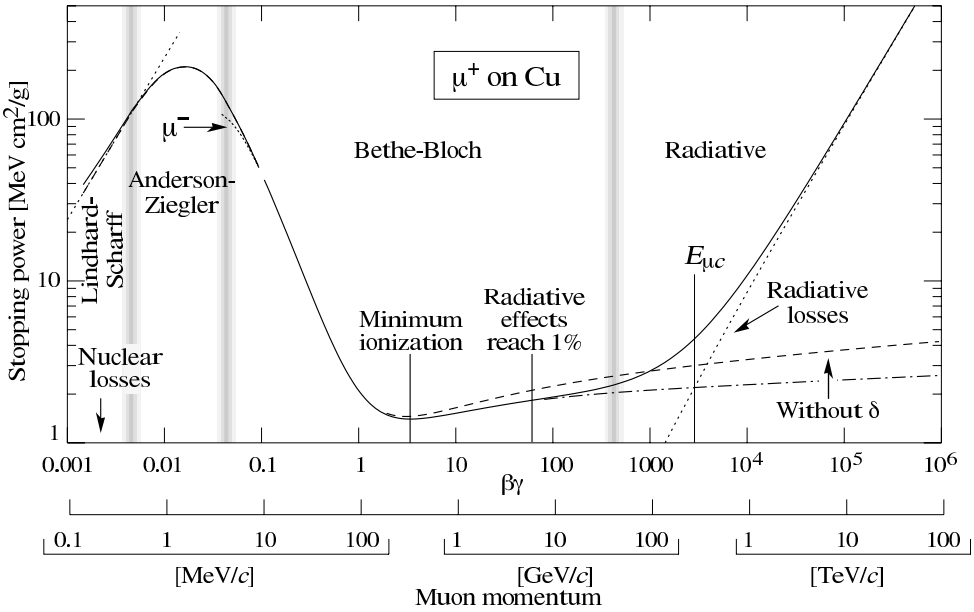

Figure 3.3: $1/\rho\langle -dE/dx\rangle$ for muons in copper, where ρ is the density of copper. This plot includes high-energy radiative effects which are not included in the Bethe-Bloch formula.

"Knock-On" δ Radiation and Mean vs. Mode Energy Loss

When the energy transfered by an incident particle to an electron in the material is large, e.g. $I \ll T \leq T_{max}$, the electron becomes a secondary particle which can ionize, excite and radiate. Such δ-electrons are produced rarely enough that they do not have a Gaussian distribution but are common enough that their effects must be taken into account. Energy loss caused by δ-electrons are well described by a Landau distribution, so the mean loss is not the same as the mode. This is plotted for a specific example in Fig. 3.4 (b).

Electrons and Photons in Matter

Electrons and photons behave similarly to one another in matter. Electrons will radiate photons via bremstrahlung, losing almost all of their energy in this manner for $E \gtrsim 10$ MeV (see Figure 3.4 (a)). Photons will annihilate in an e^+e^- pair for $E_\gamma \geq 2m_e \approx 1$ MeV. The similarity in behavior prompts characterization in terms of the radiation length, X_0. In general, $X_0 \sim \alpha r_e^2$, where α is the electromagnetic coupling constant and r_e is the classical radius of the electron. In practice, the radiation length is approximated by $X_0 \approx 180\ A/Z^2$ [g cm^{-2}], where A is the number of nucleons and Z is the number of protons in the atoms of the material. In these units, X_0/ρ has units of length. For an electron, this characteristic length is the average distance over which it looses all except for $1/e$ of its energy, and for a photon is 7/9 of its mean free path. In effect, electrons radiate photons while photons split into electron-positron pairs. This sequence creates a cascade when an energetic electron or photon enters a material. On average 2^n particles can be expected from a single energetic electron or photon, where n is the number of radiation lengths the particle has passed through.

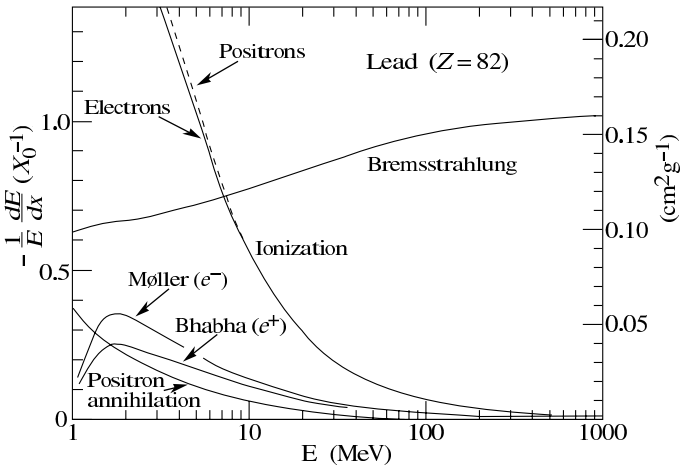

(a) Fractional energy loss of e in lead due to various processes. At $E \approx 10$ MeV, bremsstrahlung takes over at the critical point.

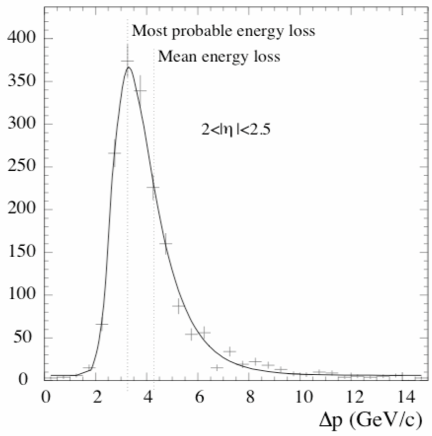

(b) The difference between mean and mode energy loss due to δ electrons, simulated for an incoming μ of $P = 50$ GeV for part of the ATLAS muon system.

Figure 3.4: Energy losses from electromagnetic interactions in matter.

Strong Interaction in Matter

Particles which dominantly interact strongly have a longer characteristic interaction length than those which interact predominantly electromagnetically. Strong force interactions have a smaller cross section than electromagnetic interactions for a hadronic particle passing through matter. This is resultant from the relatively large size of the electric field compared to the gluon field. Strongly interacting particles are forced to cascade in a similar fashion as electromagnetically interacting particles, creating similar cascades and energy deposits. The two types of showers have different shapes in general, which can help to tell them apart in a calorimeter.

3.3.2 General Detector Design

Detector experiments at particle accelerators all have relatively similar designs to one another. Built around the interaction point, an experiment generally consists of a tracker followed by a calorimeter, first electromagnetic then hadronic. This is generally surrounded by a muon system. The overall design is driven by the ability to measure quantities necessary for physics analyses. Generally, both the tracker and muon system are in magnetic fields, which allows for sign determination of charged particles. Trackers use technology such as scintillating fibers and semiconductor detectors for fine spatial resolution, allowing charge and momentum measurements within a magnetic field. Electromagnetic calorimeters detect the cascade described above, generally reconstructing electrons, photons and neutral pions (which decay into two photons very rapidly). Most hadrons will interact primarily via the Strong force in matter and are therefore detected in the hadronic calorimeter, placed beyond the electromagnetic calorimeter as a result of the longer interaction length of the Strong force with respect to the Electromagnetic force. Most muons are in the range of minimum ionization and so require special detector components. Outside of the hadronic calorimeter, the muon system is found. It consists of alternating layers of stopping material (generally iron) and detectors such as drift tubes and strip chambers. Muon detectors are generally in a magnetic field, allowing for sign assignment and finer momentum resolution. The combination of these layers allows for measuring many of the properties of particles resultant from a collision and in turn discerning the underlying physics. With the notable exception of the neutrinos, all known particles or their decay constituents are measured at hadron collider detectors. The identification of neutrinos at such a detector relies on finding missing energy in the plane transverse to the beam line, the plane in which the energy sums to zero. This general design was pioneered by the UA1 and UA2 experiments at CERN's $Sp\bar{p}S$, the first 4π detectors used at accelerators.

Tracking

Tracking is used to measure a charged particle's position as it moves through the detector. From this trajectory its momentum and charge can be found if the tracker is in a magnetic field. A charged particle moving in a magnetic field is bent according to the Lorentz force, $\vec{F} = q(\vec{E} + \vec{v} \times \vec{B})$. In general, the radius of curvature due to a magnetic field is proportional to the momentum component transverse to the direction of motion of the particle and inversely proportional to the magnetic field. In the case of a solenoidal magnet found in the inner detectors of each experiment discussed here, the field is oriented is the $\pm z$ direction. In this case a charged particle's momentum transverse to the beamline is the component transverse to the field. With $p_T = \sqrt{p_x^2 + p_y^2}$ in GeV/c, the radius of curvature in meters is

$$r = \frac{p_T}{0.3B}, \tag{3.6}$$

where B is in Tesla. Momentum resolution from tracking therefore depends on spatial position resolution. A larger magnetic field and a smaller transverse momentum yield a smaller radius of curvature, allowing for more position measurements and a finer reconstruction of the momentum. In the limiting case where a track looks straight within the detector the sign of its charge cannot be found.

Calorimetry

The principal design of a sampling calorimeter is to have a signal board sandwiched by dense (e.g. uranium) plates, with gaps between the board and plates filled with highly ionizable matter (e.g. liquid argon). Showers develop as a result of the plates and the ionization is caused by the shower constituents. The ionization is then measured on the signal board. Because ionization is being measured, calorimeter energy resolution is dependent on the number of particles counted. The number of particles recorded has a Poisson distribution, so $\Delta N \; \alpha \; \sqrt{N}$. Accordingly,

$$\frac{\Delta E}{E} \alpha \frac{\Delta N}{N} = \frac{1}{\sqrt{N}} \alpha \frac{1}{\sqrt{E}} \tag{3.7}$$

Clearly energy resolution improves with shower cascades containing more particles and with a more ionizable gas, as well as with incident particles of higher energy.

4 Experimental Setup

The leading hadron accelerators, Fermilab's $p\bar{p}$ Tevatron collider and CERN's soon to be operational pp collider, the LHC, are described below. The $D\emptyset$ experiment at the Tevatron as well as the CMS and ATLAS experiments at the LHC are described.

4.1 The Tevatron Accelerator

The Tevatron Collider at Fermilab near Chicago, IL, has been operational since 1983, first providing beam for fixed target experiments and later as a beam-beam collider. Since 1992 it has been providing $p\bar{p}$ beam collisions, first at $\sqrt{s} = 1.8$ TeV during Run I (1992 to 1995) and later at $\sqrt{s} = 1.96$ TeV during Run II (ongoing since 2001). There are two interaction points at the Tevatron. The two experiments, $D\emptyset$ and CDF, are built at these two points.

The Tevatron was the first synchrotron to make use of superconducting technology. It has been upgraded several times since its inception. A schematic of the Tevatron in Run II is shown in Fig. 4.1, and is described in detail in [16]. A 750 keV Cockroft-Walton provides protons to the linear accelerator, which in turn sends 400 MeV protons to the booster. The booster is an 8 GeV synchrotron which outputs $\sim 5 \times 10^{12}$ protons. These protons are sent to the Main Injector. This is a synchrotron with 150 GeV maximum energy which serves several purposes. Proton beams of 120 GeV from this synchrotron are used in antiproton production. Such a beam strikes a nickel target, and creates antiprotons with an efficiency of $\sim 15 \times 10^{-6}$ $\bar{p}/p$. The incoming beam of $\sim 5 \times 10^{12}$ p can be expected to produce $\sim 10^8$ $\bar{p}$. Antiprotons are stacked in this synchrotron, and both protons and antiprotons are accelerated here up to 150 GeV for injection into the main Tevatron synchrotron.

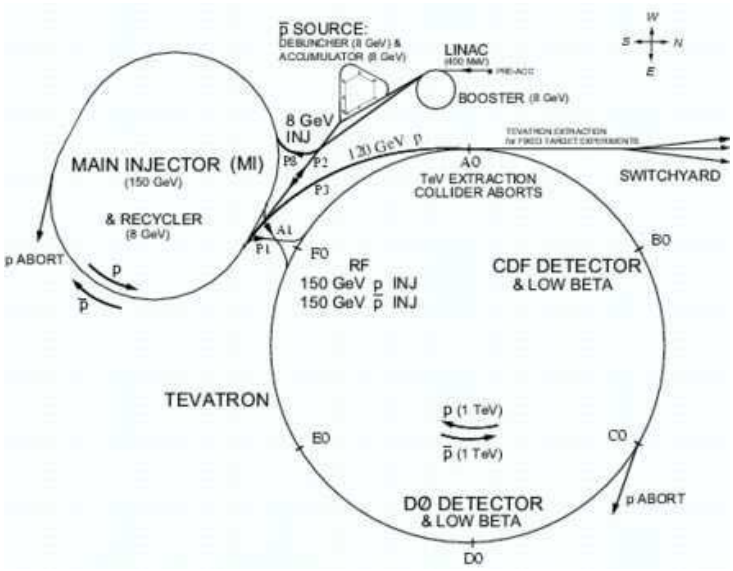

Figure 4.1: The Fermilab Accelerator chain in its Run II configuration.

The Tevatron has beams of p and $\bar{p}$ which travel in the same magnetic field in opposite directions to one another weaving a double-helix. Each beam has 36 bunches at 980 GeV spaced at 396 ns. The Tevatron's luminosity during Run II is shown in Fig. 5.4. It has recently reached sustained instantaneous luminosities of 2.7×10^{32} cm^{-2} s^{-1} [17]. A more detailed explanation of the Tevatron's current setup can be found in [16].

4.2 The LHC Accelerator

The Large Hadron Collider (LHC) is a pp collider at CERN in Geneva, CH, and is set to begin operations in the autumn of 2008. Built in the 27 km circumference tunnel that was dug for the LEP accelerator, the LHC will collide proton beams at $\sqrt{s} = 14$ TeV [18]. The LHC has a design peak instantaneous luminosity of $\sim 10^{34}$ cm^{-2} s^{-1}, about two orders of magnitude larger than the Tevatron. This luminosity will be achieved by having 2,808 proton bunches with $\sim 10^{11}$ protons per bunch. At this energy and bunch density, approximately 20 interactions per bunch crossing are expected, with crossings every 25 ns. Accordingly, bunch crossings will occur at 40 MHz and interactions at ∼1 GHz. There are four interaction points at the LHC. At these points are the CMS and ATLAS general purpose detectors, the ALICE detector dedicated to heavy ion collision studies and LHCb, an experiment dedicated to b-physics.

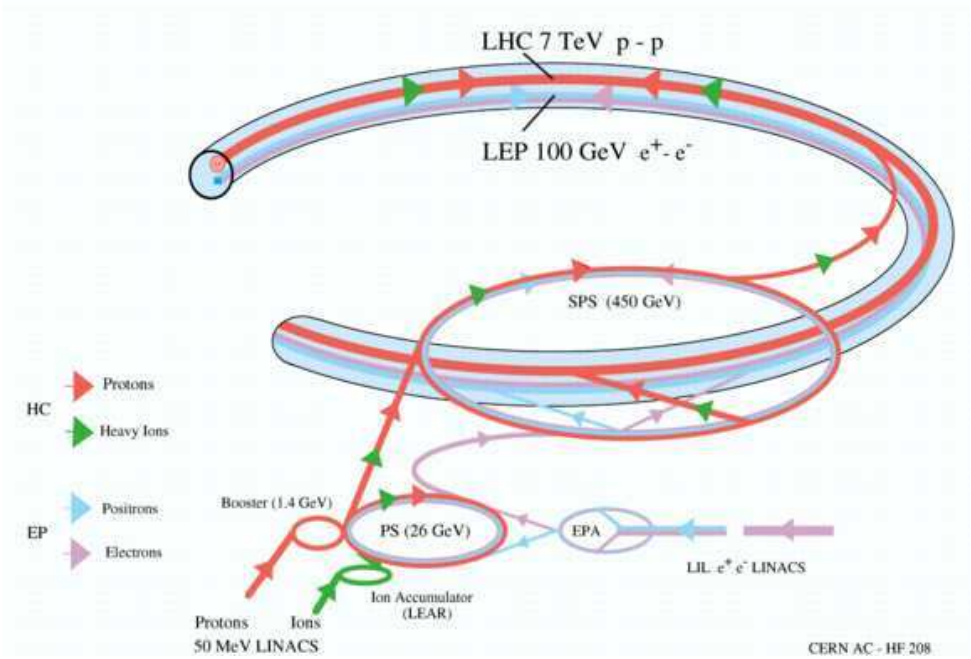

Figure 4.2: A schematic of the LHC accelerator complex for protons and lead ions.

In the LHC accelerator chain, protons are accelerated through a linac to 50 MeV for injection into the Proton Synchrotron Booster (PSB), which accelerates protons to 1.4 GeV. They are then sent to the Proton Synchrotron (PS) for acceleration to 25 GeV, then to the Super Proton Synchrotron (SPS) to be accelerated to 450 GeV. Protons are then sent to the final accelerator, the Large Hadron Collider superconducting synchrotron, to be accelerated to collision energies of 7 TeV. The LHC accelerator has 1,232 bending dipoles. A system of quadrupole, sextupole

and octupole magnets is used to focus beams at interaction points. The LHC's accelerator chain is shown in Fig. 4.2.

At times the LHC will provide lead-lead collisions for heavy ion studies. Dedicated experiments such as ALICE will take data then, as will the more general purpose CMS and ATLAS detectors. Lead ions follow a chain that is similar in principal to protons, however the first stage after a linac is the Low Energy Ion Ring (LEIR), formerly used as LEAR. Ions then go to the PS, then SPS and LHC. Lead beam collisions will occur at $\sqrt{s} = 1.15$ PeV. A more detailed explanation of LHC parameters can be found in [18].

4.3 The DØ Experiment

The *DØ* detector is a layered, multi-purpose detector residing at one of two interaction points at Fermilab's Tevatron. It examines the $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV provided by the Tevatron. The *DØ* detector consists of a tracker closest to the interaction point, followed by both electromagnetic and hadronic calorimeters, surrounded by a muon system. A schematic of the detector is shown in Fig. 4.3.

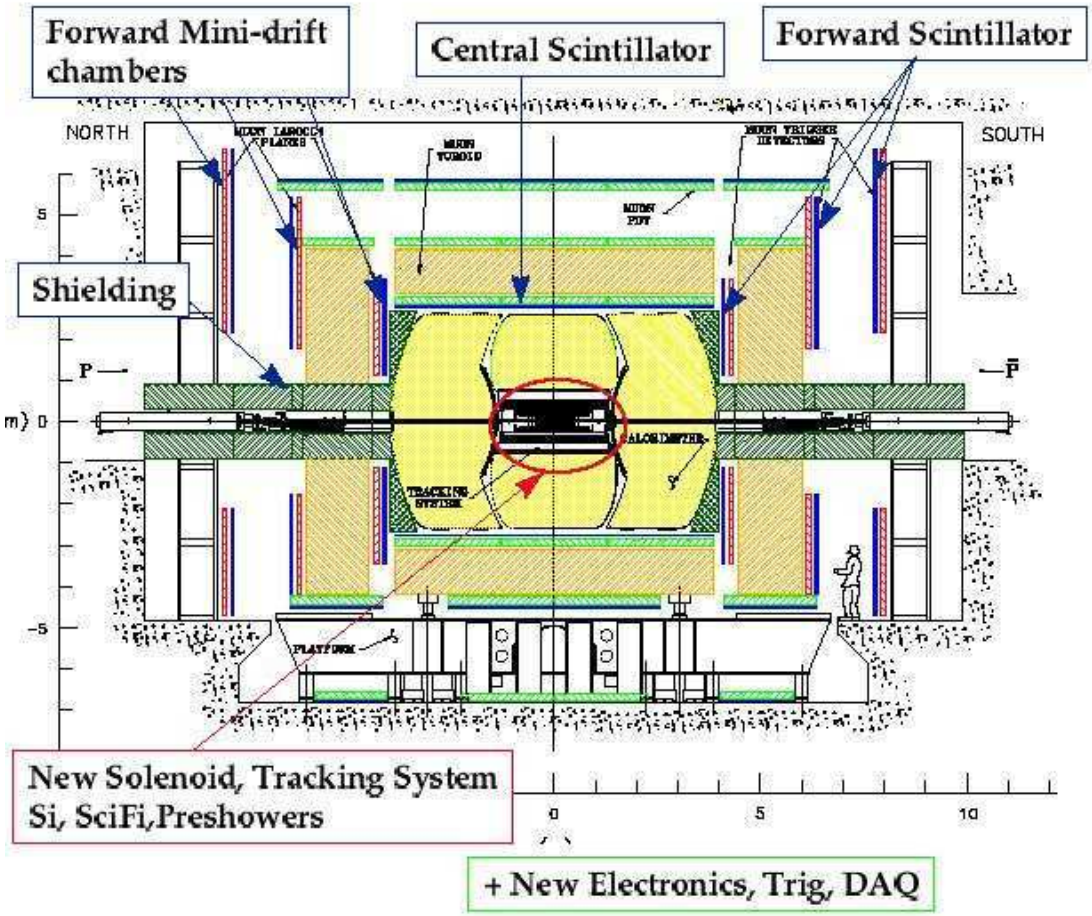

Figure 4.3: A schematic of the Run II *DØ* detector along the beamline.

Inner Detector

The *DØ* inner detector is comprised of the Silicon Vertex Detector and the Central Fiber Tracker, both within a 2.7 m long solenoidal magnet providing a 2 Tesla field. Tracking within the central region allows for vertex reconstruction, a tool essential to the bottom and top physics programs at *DØ*.

Calorimeter System

The calorimeter is a uranium-liquid argon sampling calorimeter which is comprised of a central calorimeter covering $|\eta| < 1.3$ and two endcaps extending out to $|\eta| < 4.2$. The calorimeter measures energy deposits of charged particles in the electromagnetic calorimeter and of strongly interacting particles in the hadronic calorimeter. The calorimeter is segmented into regions approximately 0.1×0.1 in $\eta \times \phi$ space. The calorimeter readout path and trigger system is described in detail in this thesis. Schematics of the calorimeter are shown in Figure 5.1.

In the manner described, most particles are detected. Muons from collisions are generally minimum ionizing particle and thus deposit low amounts of energy in the calorimeter, however they can be detected here. Neutrinos are not detected as they only interact weakly.

Muon System

The *DØ* muon system is comprised of two regions: central and forward. The central muon system covers $|\eta| < 1.0$ and contains three layers of proportional drift tubes as well as a layer of scintillators, used for triggering. The forward muon system contains three layers of mini drift tubes, extending coverage out to $|\eta| < 2.0$. The closest of the three layers of both the forward and central muon systems are found within a 1.8 Tesla toroidal field, which allows for momentum measurement and sign identification of the muons. The outer two layers are found outside of the toroid.

Trigger

The *DØ* trigger uses a three-tier system: Level 1, Level 2 and Level 3. A trigger is necessary in order to contain the amount of information recorded. At the bunch crossing rate of 396 ns provided by the Tevatron not all events can be recorded. In addition, detector dead-time due to a full system readout is far greater than 396ns, but the limiting factor is still offline reconstruction. The third level of trigger must reduce the event rate from ∼2.5 MHz interaction rate to ∼100 Hz for offline reconstruction and permanent storage. The *DØ* trigger system is described in detail later in this thesis. A more detailed description of the current *DØ* experimental setup can be found in [19].

4.4 The CMS Experiment

The Compact Muon Solenoid (CMS) detector is also a layered, multi-purpose detector. It resides at one of the interaction points, Point 5, of CERN's LHC. CMS is designed to examine the pp collisions provided by the LHC at $\sqrt{s} = 14$ TeV. As of writing, the detector is near completion in the underground cavern and is expected to be ready for the first collisions provided by the LHC in October 2008. CMS consists of an inner tracking detector, followed by both an electromagnetic and hadronic calorimeter, and then surrounded by a muon system. Most of the detector is

within the detector's 4 T solenoidal field. A schematic is shown in Fig. 4.4. The CMS detector is more fully described in other documents such as [20].

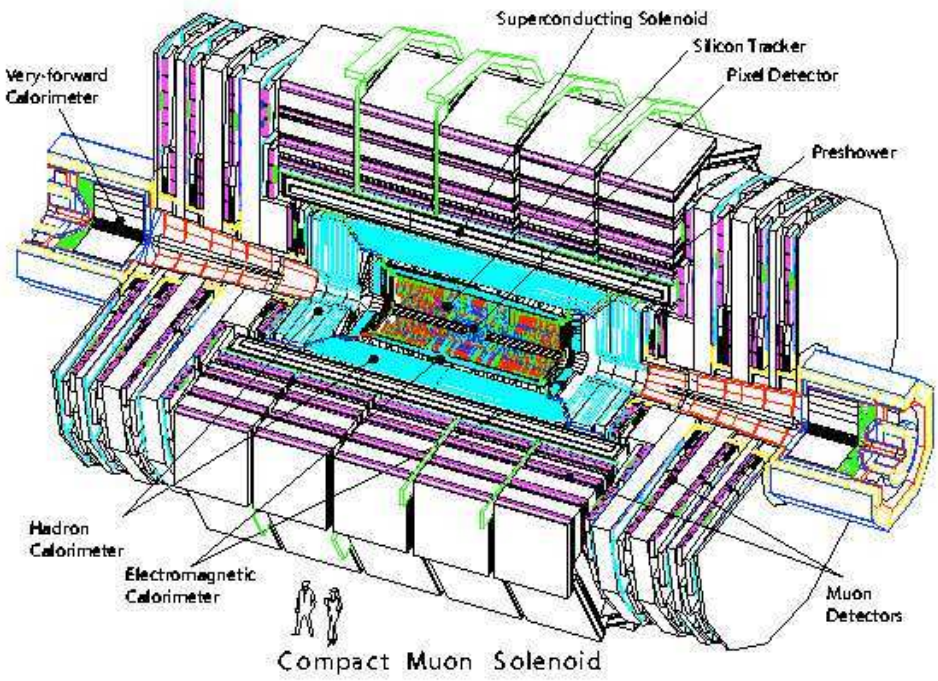

Figure 4.4: A schematic view of the CMS detector.

Inner Detector

The CMS detector subsystem closest to the beamline is the inner tracking system which is entirely within the 4 T solenoidal magnetic field. The component closest to the interaction point is a semiconducting silicon pixel detector comprised of pixels with size $\approx 100 \times 150\ \mu\text{m}^2$. The speed and precision of a pixel detector makes it ideal for the region with the highest particle flux so it is used in the central region. The rest of the inner detector uses silicon microstrips, segmented into a closer region where strips of $\approx 10\ \text{cm}^2 \times 80\ \mu\text{m}^2$ are used and a farther region where strips of $\approx 25\ \text{cm}^2 \times 180\ \mu\text{m}^2$. The full inner detector system covers up to $|\eta| = 2.4$.

Calorimeters

The calorimeter system in CMS is a non-compensating calorimeter comprised of a crystal electromagnetic calorimeter (ECAL) and a scintillator hadronic calorimeter (HCAL). The ECAL uses 61,200 lead tungstate ($PbWO_4$) crystals in the central barrel and an additional 7,324 in each of the 2 endcaps. While these crystals are fast and radiation hard they have a low light yield, requiring photodetectors with gain that can operate in a magnetic field; silicon avalanche photodiodes are used for this. The barrel region covers $|\eta| < 1.479$ while the endcaps extend to $|\eta| < 3$. Each crystal in the barrel ECAL corresponds to 0.0174×0.0174 in $(\eta \times \phi)$ space.

The hadronic calorimeter uses brass absorbers and silicon scintillators which are read out with wavelength shifting fibers. The barrel region covers $|\eta| < 1.4$, segmented as 0.087×0.087 in $(\eta \times \phi)$ space. The endcap region extends to $|\eta| < 3$ while the forward calorimeter covers $3 < |\eta| < 5$.

Muon System

The CMS muon system consists of three types of components: drift tubes (DTs) and cathode strip chambers (CSCs) are used for precision measurements while resistive plate chambers (RPCs) are used for fast readout. The DTs are used in the barrel region, covering $|\eta| < 1.2$, arranged in four radial layers. The endcap region, extending to $|\eta| < 2.4$, is covered by CSCs. There are four layers of CSCs going outward in z. The muon system, especially the CSCs, will be covered in more detail below.

Trigger

The CMS trigger has been designed for the LHC's peak design instantaneous luminosity of $\sim 10^{34}$ cm^{-2} s^{-1} [21]. The ~ 1 GHz interaction rate has to be reduced to ~ 100 Hz of events to be stored. This is done using the trigger. CMS uses a two-level trigger system, the Level-1 Trigger and the High Level Trigger. All events are examined by the Level-1 Trigger, which reduces the data flow by a factor 10^4 and passes events to the High Level Trigger at a rate of 100 KHz. The High Level Trigger will retain events at a rate of 100 Hz which will be stored permanently and used for analysis. It should be noted that the high event rate combined with the physical size of the CMS detector posed significant challenges to the designers of the trigger system, which have largely been met.

4.5 The ATLAS Experiment

The ATLAS (A Toroidal LHC Apparatus) detector is the other layered, multi-purpose detector at the LHC. It resides at interaction Point 1 of the collider. Much like CMS, ATLAS is designed to study the pp collisions from the LHC at $\sqrt{s} = 14$ TeV. The ATLAS detector is expected to be operational for the first test collisions in October 2008. ATLAS has an inner tracking detector, comprised of three separate systems designed to work together, followed by an electromagnetic and hadronic calorimeter, and lastly a muon system. A schematic view of the entire detector is shown in Fig. 4.5. The ATLAS inner detectors are within a 2 T solenoidal field and the outer muon system is inside a toroidal field that has eight-fold symmetry around the beam pipe, shown in Fig. 4.6. A smaller toroidal magnetic system, also with eight-fold symmetry and staggered with respect to the barrel magnets, is used for the outermost regions of ATLAS. A more complete description of the ATLAS detector can be found in [22, 23].

Inner Detector

The ATLAS inner detector is contained within a cylinder which is 7 m long and has a radius of 1.15 m residing within a 2 T solenoidal field. The entire inner detector is designed to give coverage out to $|\eta| < 2.5$. The inner detector has a pixel vertex detector with three staggered layers in the barrel and five endcaps on each side, a total of $\sim 2 \times 10^8$ channels. Beyond the pixel vertex detector is a series of silicon strip layers in the barrel and wheels on the ends. The outermost component is the Transition Radiation Tracker (TRT), designed to distinguish pions from electrons passing through the tracker and entering the electromagnetic calorimeter. The

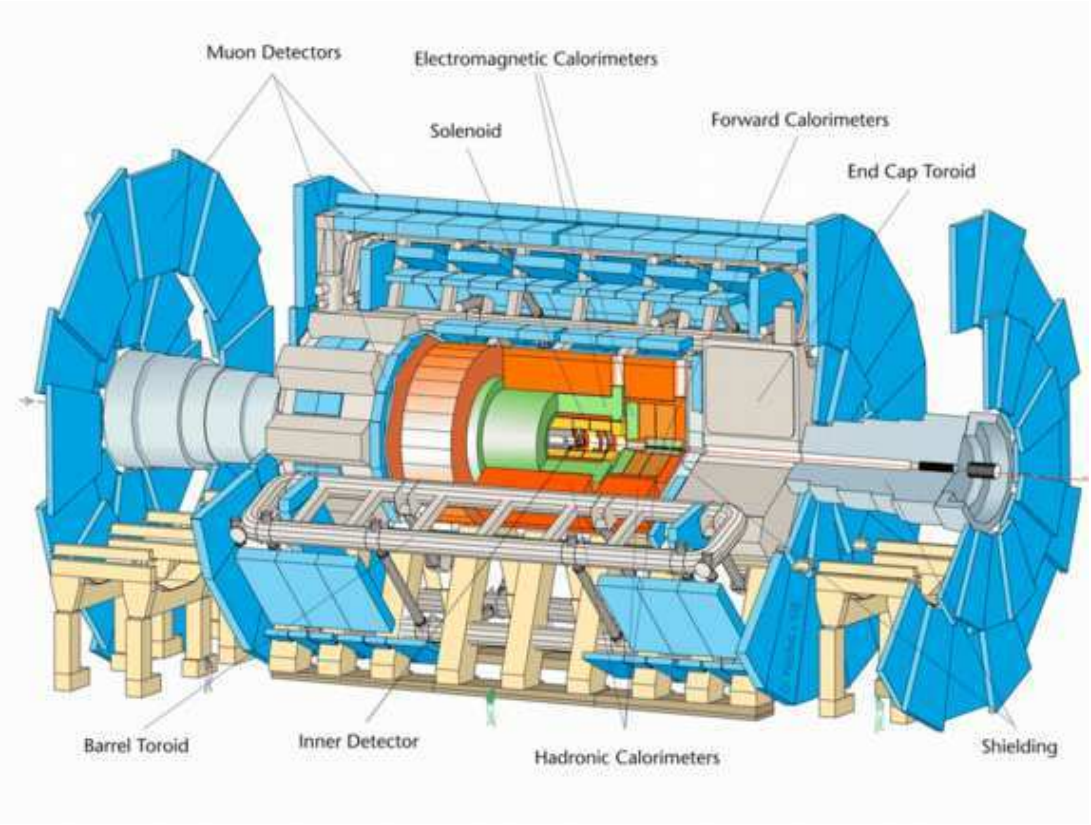

Figure 4.5: A schematic view of the ATLAS detector.

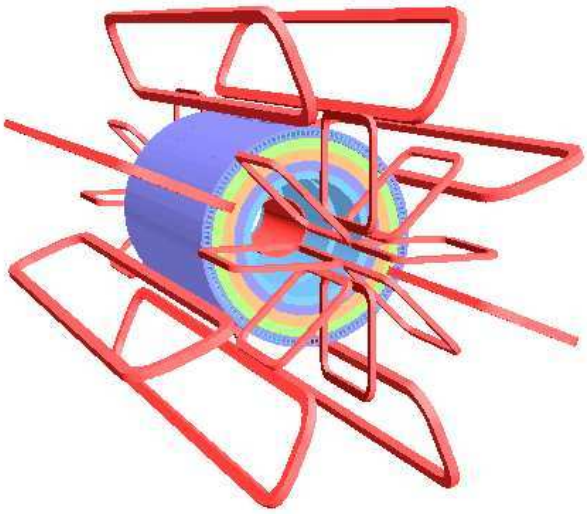

Figure 4.6: A visualization of the ATLAS magnets.

TRT also adds additional tracking points for all charged particles giving a more precise overall tracker. An image of the inner detector design is shown in Fig. 4.7.

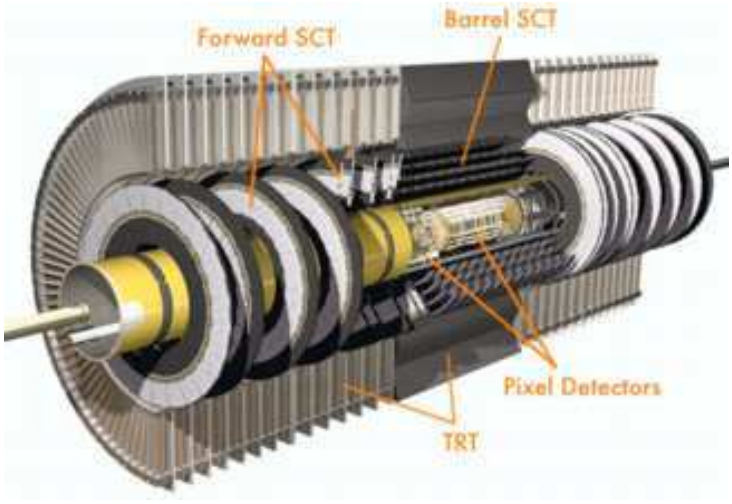

Figure 4.7: The ATLAS inner detector layout.

Calorimeters

At ATLAS there is a barrel electromagnetic sampling calorimeter covering up to $|\eta| < 3.2$ and then three sections of the hadronic calorimeter: a barrel system covering $|\eta| < 1.7$, endcap covering $1.5 < |\eta| < 3.2$ and a forward system covering $3.1 < |\eta| < 4.9$. The electromagnetic system uses lead as an absorber and liquid argon as a sampling material. The hadronic system uses two different methods for the various components. In the barrel region iron is used as an absorber and scintillating fibers are used as an active material. Both the endcap and forward systems use a lead and liquid argon system. In the barrel region, a presampler is used to correct for energy losses.

Muon System

The Muon System at ATLAS, like that at CMS, is comprised of a variety of types of chambers. The barrel system, covering $|\eta| < 1.0$ uses 3 layers of monitored drift tube chambers for precision readout and uses resistive plate chambers for trigger information. The endcap system uses cathode strip chambers for precision readout and thin plate chambers for triggering. The endcap system resides on four disks which are not back-to-back but are at a distance of 7, 10, 14 and $\sim$ 22m from the interaction point. This system covers $1.0 < |\eta| < 2.7$ with trigger information up to $|\eta| = 2.4$.

Trigger

The ATLAS trigger has been designed for the LHC's peak design instantaneous luminosity of $\sim 10^{34}$ cm^{-2} s^{-1} [22]. The trigger is used to reduce the full data flow from the $\sim$ 1 GHz interaction rate to $\sim$ 100 Hz of events to be stored. The ATLAS trigger system does this using a three-level system. The first system, the Level 1 trigger, makes cuts based on crude information and sends a maximum of $\sim$ 75 KHz to the Level 2 system which makes use of object position as well. The Level 2 system brings the data flow to $\sim$ 1 KHz. The High Level Trigger will then use finer information from offline reconstruction to reduce the data flow to the acceptable rate of 100 Hz.

5 Monitoring in the Run IIb Level 1 Calorimeter Trigger Upgrade at $D\emptyset$

Some of the work presented in this chapter has been published as a part of a larger paper in [24], attached in full in Appendix A.

The Tevatron accelerator has delivered data in three different phases since it began beam-beam collisions in 1992: Run I, Run IIa and Run IIb. For Run II beginning in 2001, the Tevatron has been delivering significantly higher luminosity as well as a slightly higher center of mass energy [16] than in Run I. The experiments made upgrades to meet physics goals as well to keep up with the higher luminosity. In May 2006, Fermilab's Tevatron began the Run IIb. This marked a dramatic increase in the instantaneous luminosity delivered by the Tevatron to the experiments. In turn, this required upgrades of the detectors.

The Run IIb Tevatron upgrade stresses $D\emptyset$ in two main ways [19]. First the additional integrated luminosity delivered to $D\emptyset$ during the course of Run IIb increases the total radiation dose accumulated by the silicon detector. Best estimates indicated that this increased dose would compromise the performance of the inner layer of the detector, affecting the ability of $D\emptyset$ to tag b-quarks, which is necessary in much of the experiment's physics program. Secondly, the increased instantaneous luminosity stresses the trigger system, decreasing the ability to reject background while maintaining high efficiency for signal events.

The increased radiation to the silicon detector led $D\emptyset$ to propose the addition of a radiation-hard inner silicon layer to the tracking system. The stress on triggering required changes to various aspects of the trigger system. These additions and modifications, collectively referred to as the "$D\emptyset$ Run IIb Upgrade" were designed and implemented between 2002 and 2006 and were installed in the experiment during the 2006 Tevatron shutdown.

The changes to the Level 1 calorimeter trigger will be briefly reviewed here, while more substantial information can be found in [19]. These changes included installing new off-detector electronics (in the counting house) to run the new trigger algorithm, the Sliding Windows Algorithm. Monitoring of these new electronics and of the calorimeter's performance in Run IIb will be described in this thesis.

5.1 Preexisting System, Motivations and Upgrade

The $D\emptyset$ calorimeter consists of three sampling calorimeters (one barrel and two endcaps), using liquid argon as the active medium and depleted uranium, uranium-niobium alloy, copper or stainless steel as the absorber. The calorimeter has three longitudinal sections (going outwards in r): electromagnetic (EM), fine hadronic (FH) and coarse hadronic (CH). Each of these is subdivided into several layers. It is segmented laterally into cells of size $\Delta\eta \times \Delta\phi \approx 0.1 \times 0.1$, arranged in pseudo-projective towers. The calorimeter system provides coverage out to $|\eta| = 4.2$. It is shown schematically in Fig. 5.1.

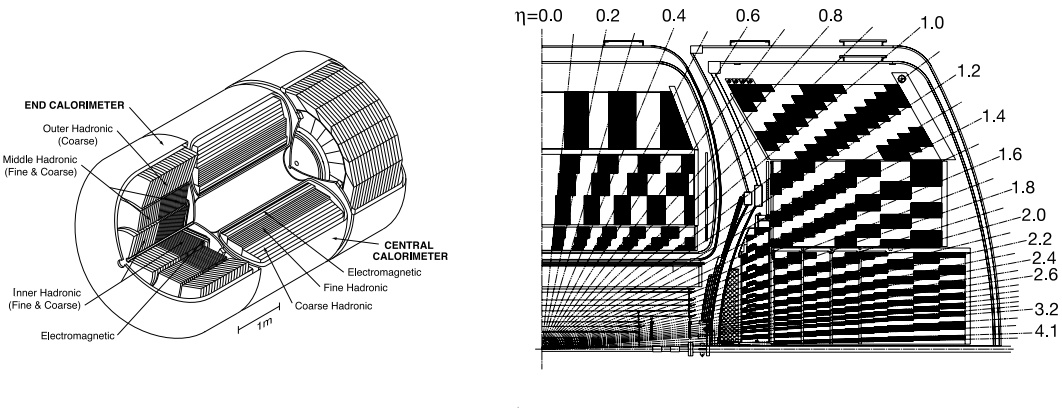

(a) Schematic cross section.

(b) One quadrant, showing segmentation.

Figure 5.1: The *DØ* Calorimeter.

Calorimeter Readout Path

Charge is collected in the calorimeter and is then transmitted to preamplifiers located on the detector[1]. The charge integrated output of these preamplifiers has a rise time of ∼450 ns, corresponding to the electron drift time across a liquid-argon gap, with a fall time of ∼15 μs. These signals are sent over ∼ 25 m long cables to Baseline Subtractor (BLS) cards, of which there are 1152.

On the BLS cards, the preamplifier signals are split into two paths: the precision readout and the trigger sum pickoff. Precision readout path signals for each calorimeter cell are stored in a set of switched capacitor arrays awaiting Level-1 and Level-2 trigger decisions. Signals on the trigger sum pickoff path are shaped to a triangular pulse with a fast rise and a linear fall over 400 ns. They are then passed to analog summers that add signals in different cells, weighted appropriately for the sampling fraction and capacitance of each cell to form EM and HAD trigger towers (TTs). The granularity leads to 1,280 EM TTs and 1,280 HAD TTs forming a 40×32 grid in $\eta \times \phi$ space, which covers the entire azimuthal region for $|\eta| < 4.2$. Due mainly to overlapping collisions which complicate the forward environment only the region $|\eta| < 3.2$ is used for triggering. The EM and HAD TT signals are transmitted differentially to the Level 1 (L1) Calorimeter Trigger. This output from the calorimeter did not change in any fundamental way during the Run IIb upgrade. The output from the calorimeter to the L1 trigger was the same before the upgrade as after. The trigger system itself, however, was changed.

Trigger System Overview

The *DØ* experiment uses a three level trigger system to select ∼ 100 Hz of interesting events from the 2.5 MHz seen in the detector. The L1 trigger system examines data from the detector for every bunch crossing, deciding whether or not it is "interesting". This is implemented in custom hardware, which takes data from various separate elements of *DØ*: from the calorimeter (L1Cal), scintillating fiber tracking (L1CTT), muon (L1Muon), and forward proton (L1FPD)

[1] The following sections lean heavily on [24].

systems. An addition for Run IIb is an element that matches tracks and calorimeter clusters at L1 (L1CalTrk). The trigger system is shown schematically in Fig. 5.2.

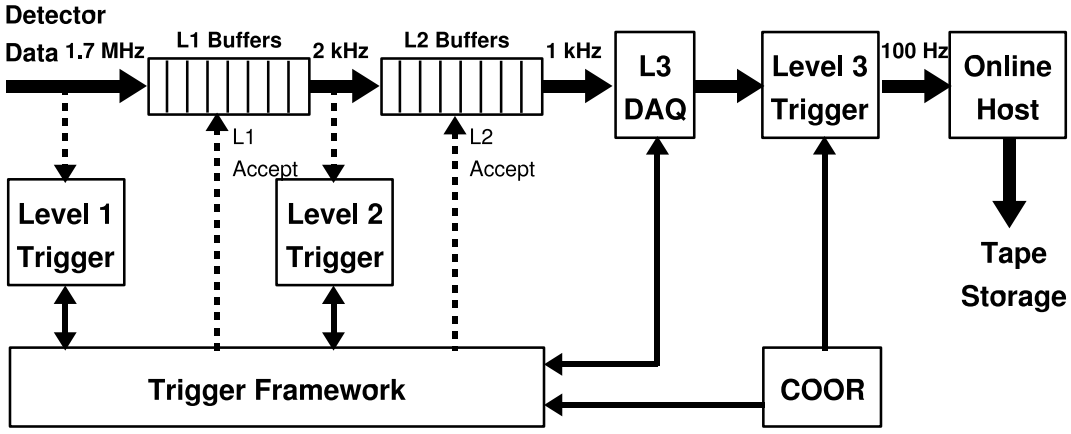

Figure 5.2: An overview of the *DØ* trigger and data acquisition system.

Each L1 trigger element sends its decisions on a set of criteria (for example, the presence of two jets with transverse energy above a threshold) to the trigger framework (TFW). The TFW uses these decisions, referred to as the *and/or* terms, to decide whether the event should be accepted for further processing or rejected. Because of the depth of data pipelines in the detector's front end electronics, L1 decisions from each of the trigger elements must arrive at the TFW within 3.7 μs of the bunch crossing which the data is associated with. This pipeline depth was increased from its Run IIa value of 3.3 μs in order to accommodate the extra latency needed for the L1CalTrk system.

If the TFW decides that the event passes L1 criteria, an *L1 Accept* is issued. After an L1 Accept, data are transferred off of the pipelines. This transfer process includes a certain deadtime. The maximum allowable L1 accept rate, generally around 2 KHz, is set by the desire to limit this deadtime to the 5% level.

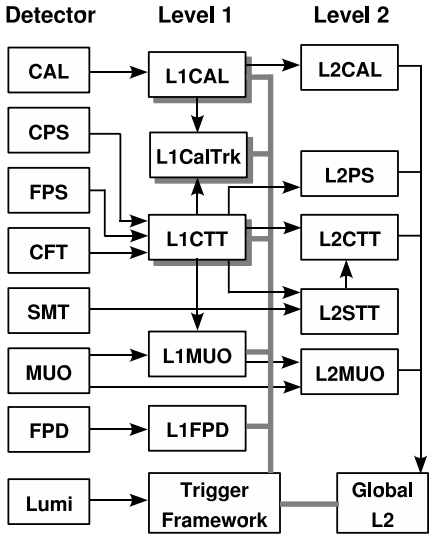

Figure 5.3: A block diagram of the *DØ* L1 and L2 trigger systems.

The L2 system receives data from the detector and from the L1 trigger elements when an L1 Accept is issued. The L2 system consists of detector-specific pre-processor engines for calorimeter (L2Cal), preshower (L2PS), scintillating fiber (L2CTT) and silicon (L2STT) tracking, as well as muon (L2Muon) data. Processed data from each of these elements is transmitted to a global processor (L2Global) that selects events based on detector-wide correlations between its input elements. This is shown as a block diagram in Fig. 5.3. The L2 trigger operates at a maximum input rate of 2 kHz and provides L2 accepts at a rate of up to 1 KhZ.

The final stage in the *DØ* trigger system, Level-3 (L3), consists of a farm of PCs that have access to the full detector readout on L2 accepts. These processors run a simplified version of the offline event reconstruction software and make decisions based on physics objects and the relationships between them. An event with an L3 accept is kept for permanent storage at a rate of up to 150 Hz (typically, 100 Hz). The configuration of the entire *DØ* trigger system is accomplished under the direction of the central coordination program (COOR), which is also used for detector configuration and run control.

5.1.1 Run IIb Trigger System Upgrades

The motivation for the L1Cal IIb project was to make the most of the planned Tevatron upgrades. While the original goal for Run II at the Tevatron was an integrated luminosity of 2 fb^{-1}, the upgrades for Run IIb were designed to deliver an additional 2-6 fb^{-1} of data. That exceptional goal has been largely met since May 2006.

The trigger for the *DØ* experiment in Run IIa had been designed for a maximum luminosity of 1×10^{32} cm^{-2} s^{-1}, while the peak luminosities in Run IIb were anticipated to go as high as 3×10^{32} cm^{-2} s^{-1}. Indeed, a peak luminosity of 3.2×10^{32} cm^{-2} s^{-1} was achieved while a sustained peak luminosity of 2.7×10^{32} cm^{-2} s^{-1} has been delivered recently, as shown in Fig. 5.4.

In the three-level trigger system employed by *DØ*, only the L3 trigger can be modified to increase its throughput; the maximum output rates at L1 and L2 are imposed by fundamental features of the subdetector electronics. Thus, fitting L1 and L2 triggers into the bandwidth limitations of the system can only be accomplished by increasing their rejection power. While an increase in the transverse energy thresholds at L1 would have been a simple way to achieve higher rejection, such a threshold increase would be too costly in efficiency for the physics processes of interest. The *DØ* Run IIb trigger was designed to achieve the necessary rate reduction through greater selectivity, particularly at the level of individual L1 trigger elements.

The L1Cal trigger used in Run I and in Run IIa was based on counting individual trigger towers above thresholds in transverse energy (E_T). Because the energy from physics objects (electrons, photons and especially jets) tends to spread over multiple TTs, the thresholds on tower E_T had to be set low relative to the desired object's E_T. For example, an EM TT threshold of 5 GeV is fully efficient only for electrons with E_T greater than about 10 GeV, and a 5 GeV threshold for EM + HAD tower E_T only becomes 90% efficient for jet transverse energies above 50 GeV.

The primary strategy of the Run IIb upgrade of L1Cal trigger was to improve the sharpness of the thresholds for electrons, photons and jets by forming clusters of TTs and comparing the transverse energies of these clusters, rather than individual tower E_T, to thresholds.

The design of clustering using sliding windows in Field Programmable Gate Arrays (FPGAs) meets the requirements of this strategy, and also opens new possibilities for L1Cal, including: sophisticated use of shower shape and isolation, algorithms to find hadronic decays of tau leptons through their characteristic transverse profile and requirements on the topology of the electrons,

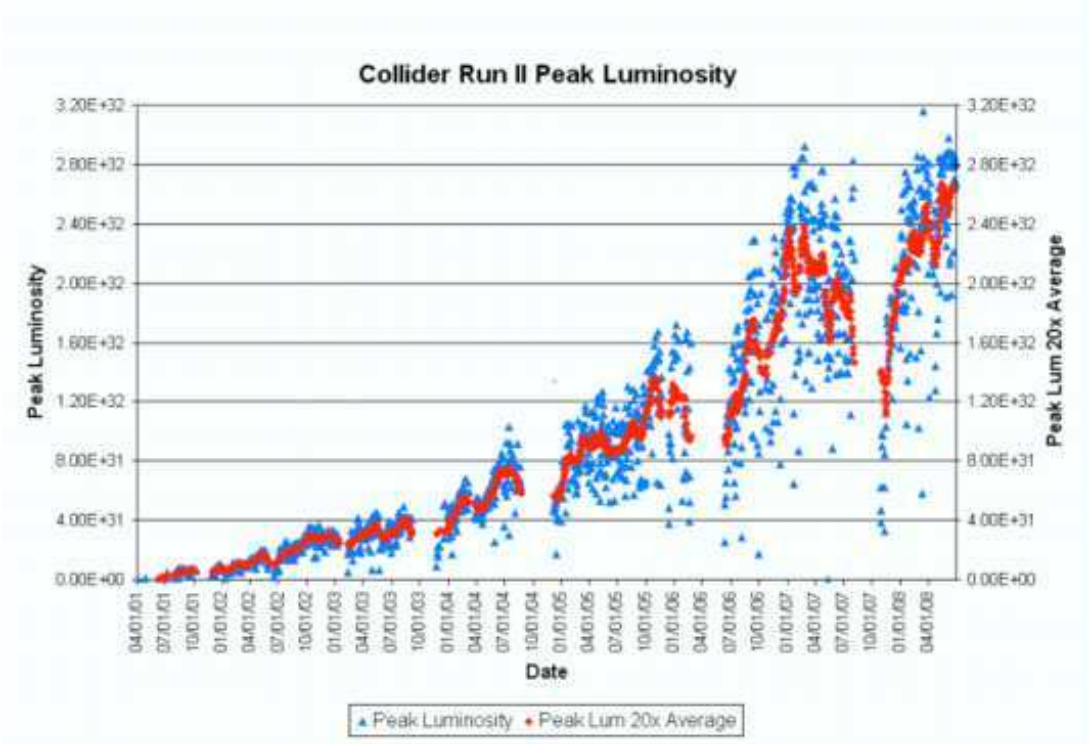

Figure 5.4: Tevatron Run II peak luminosity. The shutdown for and start of Run IIb can be seen near the summer of 2006 [17].

jets, taus, and missing transverse energy in an event.

Object Finding Algorithms

The clustering of individual TTs into EM and jet objects is implemented in the Run IIb L1Cal by the use of a sliding windows algorithm (SWA). This algorithm performs a parallel cluster search in which groups of contiguous TTs are compared to nearby groups to determine the location of local maxima in E_T deposition. This has been shown to be extremely efficient and will be used in the trigger systems of CMS at ATLAS at the LHC.

The sliding windows algorithm in the $D\emptyset$ calorimeter trigger has three phases. In the first phase, the digitized transverse energies of several TTs are summed into Trigger Tower Clusters (TTCL). These TTCL sums, based on the size of the EM or jet sliding window, are constructed for every point in trigger tower space and are indexed by the (η, ϕ) coordinate of one of the contributing TT. This process yields a grid of TTCLs that share energy with their close neighbors, as shown in the first and second panels of Fig. 5.5.

In the second phase, the TTCLs are analyzed to determine locations of large energy deposits called local maxima (LM). These LM are chosen based on a comparison of the magnitude of the E_T of a TTCL with that of its adjacent TTCLs. Multiple counting of jet or EM objects is avoided by requiring a spatial separation between adjacent local maxima as illustrated in the third panel of Fig. 5.5. In the third phase, additional information is added to define an output object. In the case of jet objects, as shown in the fourth panel of Fig. 5.5, energy of surrounding TTs is added to the TTCL energy to give the total jet object energy. EM and Tau objects are also refined in this phase using isolation information, as described below.

Results for the entire calorimeter can be obtained very quickly using this type of algorithm by performing the LM finding and object refinement phases of the algorithm in parallel for each TTCL.

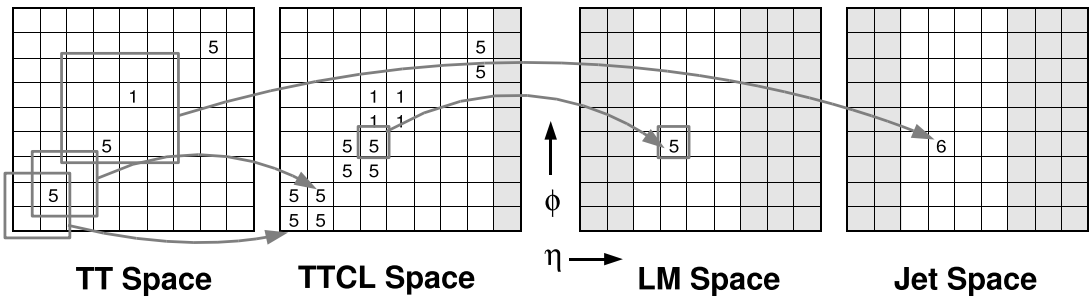

Figure 5.5: The stages of algorithm flow for the sliding windows algorithm. This example corresponds to the Run IIb jet algorithm in which a 2×2 TTCL is used, indexed by the position of its smallest (η, ϕ) TT. Baseline subtracted TT energies are indicated by numbers, and local maxima are required to be separated by at least 1 TT. Jet objects are defined as the E_T sum of the 4×4 TTs centered on the TTCL. Light gray regions in the diagrams indicate areas for which the object in question cannot be constructed because of boundary effects.

Jets

Jets at Tevatron energies have lateral sizes on the order of one unit in (η, ϕ) space and deposit energy in both the electromagnetic and hadronic portions of the calorimeter. Accordingly, jet objects in the *DØ* L1Cal are defined using the sum of the EM TT and HAD TT energies as the input to the TTCL sums. The TTCLs are 2×2 in trigger tower units, corresponding to a region of 0.4×0.4 in $\eta \times \phi$ space. Local maxima are required to be separated by one trigger tower and the final energy sums are 4×4 in TT space, corresponding to a region of 0.8×0.8 in $\eta \times \phi$ space. The values of these clustering parameters were determined by optimizing jet object energy and position resolution.

Jets from Taus

Tau leptons that decay hadronically look similar to jets from quarks but have narrower cones. This allows extra efficiency for processes containing taus to be obtained by relaxing E_T threshold requirements on potential tau objects with respect to those on jets from quarks thresholds. The SWA exploits the difference in cone radius by requiring that small amounts of energy surround the tau candidate. Thus taus can be triggered on in an efficient manner. In Run IIb, L1Cal uses the results of the jet algorithm as a basis for tau objects but also calculates the ratio of the 2×2 TTCL to the 4×4 total jet object E_T. Large values of this isolation ratio as well as large E_T are required in the definition of a tau object. Because of data transfer constraints in the system, however, the E_T associated with the tau object is taken from the jet object closest in ϕ to the LM passing the tau isolation cut.

Electrons and Photons

Electrons and photons together make up EM objects. These have lateral shower profiles that are much smaller than the TT size and tend not to deposit energy in the hadronic calorimeter. For

this reason EM TTs are input directly to the local maximum finding algorithm with a TTCL size is 1×1 in TT space. Because electrons or photons may deposit energy close to the boundary between TTs, the final EM object is comprised of two adjacent trigger towers, as shown in Figure 5.6. These can be oriented horizontally (containing two TTs in η) or vertically (containing two TTs in ϕ), where the first tower is the local maxima and the second is the neighboring tower with the highest E_T.

Cuts can also be applied on the electromagnetic fraction (EM/HAD deposits) and isolation of the candidate EM object. The electromagnetic fraction is the ratio of the sum of the energies in EM TT comprising the EM object and the corresponding two HAD TTs directly behind it. The isolation region is composed of the four EM TTs adjacent to the EM object; cuts are placed on the ratio of the total E_T in the EM-isolation region and the EM object E_T. In both cases, the ratio cut value is constrained to be a power of two in order to reduce latency in the divide operation as implemented in digital logic. This algorithm was chosen based on an optimization of the efficiency for triggering on electrons from $W \rightarrow e\nu$ and $J/\psi \rightarrow e^+e^-$ decays.

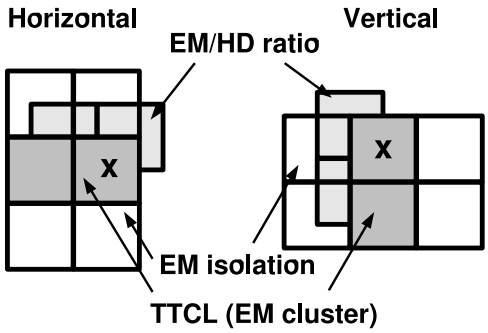

Figure 5.6: Definition of EM trigger objects.

Sum E_T and Missing E_T

Scalar and vector E_T sums are computed for the EM + HAD TTs. In constructing these sums, the η range of the contributing TTs can be restricted and an E_T threshold can be applied to the TTs entering the sums to avoid noise contamination.

Topological Triggers

Because of its increased processing capabilities, the Run IIb L1Cal can require spatial correlations between some of its objects to create topological trigger terms. These triggers can be used to distinguish signals that have numbers of objects identical to those observed in large backgrounds but whose event topologies are much rarer. An example of such a topology occurs in associated Higgs production in which the decay $ZH \rightarrow \nu\bar{\nu}b\bar{b}$ yields two jets acoplanar with respect to the beam axis and large missing transverse energy. Since the only visible energy in such an event is reflected in the jets, it is difficult to distinguish this process from the overwhelming dijet QCD background. The Run IIb L1Cal contains a trigger that specifically selects dijet events in which the two jets are required to be acolinear in the transverse plane. Other topological triggers that have been studied are back-to-back (in the transverse plane) EM object triggers to select events containing J/ψ mesons, and triggers that select events with jet-free regions of the calorimeter containing small energy deposits, for triggering on mono-jet events.

Simulation and Predictions

Two independent methods of simulating the performance of the L1Cal algorithms have been developed: a module included in the overall *DØ* trigger simulation for use with Monte Carlo or real data events, called TrigSim, and a tool developed to estimate and extrapolate trigger rates based on real data accumulated during special low-bias runs, called the Trigger Rate Tool. Both of these methods were used to develop a new Run IIb trigger list that will collect data efficiently up to the highest luminosities foreseen in Run IIb, instantaneous luminosities of up to 3×10^{32} cm^{-2} s^{-1}.

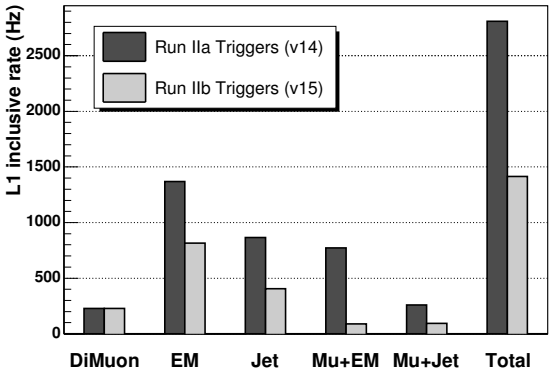

Figure 5.7: Predicted rates for Run IIa (v14) and Run IIb (v15) trigger lists, extrapolated to a luminosity of 2×10^{32} cm^{-2} s^{-1} from trigger-unbiased data collected at lower luminosity.

Predictions of the impact of the new L1Cal sliding windows algorithms on the L1 trigger rates and efficiencies were obtained using simulations of dijet events and various physics processes of interest in Run IIb. After trying different configurations that gave the same rate as those experienced during Run IIa, the most efficient configurations were chosen and put in an overall trigger list to check the total rate.

Figure 5.7 shows the predicted rates at a luminosity of 2×10^{32} cm^{-2} s^{-1} estimated using the Trigger Rate Tool, for trigger lists based on Run IIa algorithms (v14) and their Run IIb equivalents (v15). Both trigger lists were designed to give similar efficiencies for physics objects of interest in Run IIb. However, the Run IIb trigger list yields a rate approximately a factor of two smaller than that achievable using the Run IIa algorithms.

5.2 Monitoring the Calorimeter and the Level 1 Calorimeter Trigger for Run IIb

The algorithms described previously are implemented in a number of custom electronics boards designed for the new trigger. An overview of the main hardware elements of the Run IIb L1Cal system is given in Fig. 5.8. These elements can be divided into three broad groups: the Analog to Digital Filter (ADF) system, which contains the elements that receive and digitize analog TT signals from the BLS cards and perform TT-based signal processing, the Trigger and Global Algorithm Board (TAB/GAB) system, where algorithms are run on the digitized TT signals to produce trigger terms, and the Readout system, which inserts L1Cal information into the *DØ* data path for permanent storage.

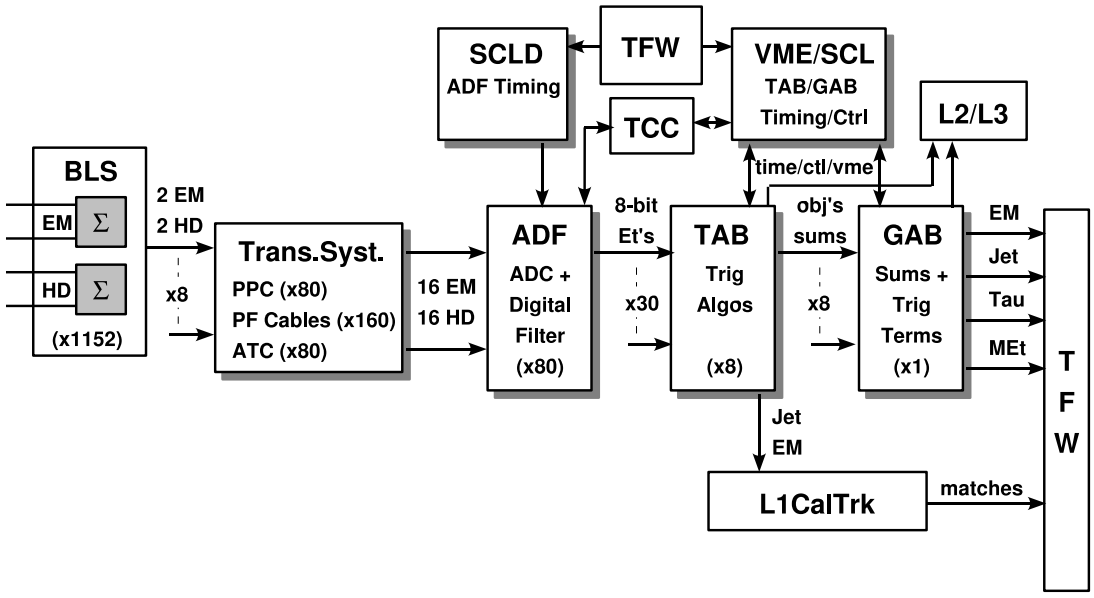

Figure 5.8: A block diagram of the main hardware elements of the Run IIb L1Cal system and their interconnections.

The L1Cal also communicates with other elements of the *DØ* trigger and data acquisition (DAQ) system, including the Trigger Framework (TFW), which delivers trigger decisions and synchronizes the entire *DØ* DAQ. From the L1Cal point of view, the TFW sends global timing and control signals to the system over Serial Command Links (SCL) and receives the L1Cal and/or terms. L1Cal also interfaces with the L1Cal Trigger Control Computer (L1Cal TCC), which configures and monitors the system, and the Level-1 Cal-Track Match (L1CalTrk) trigger system, a somewhat separate L1 trigger system that performs azimuthal matching between L1CTT tracks and L1Cal EM and jet objects.

The ADF cards are responsible for sending the best estimate of the transverse energy E_T in the EM and HAD sections of each of the 1,280 TTs to the eight TAB cards for each Tevatron beam crossing. The calculation of these E_T values by the 80 ADF cards is based upon the 2,560 analog trigger signals that the ADF cards receive from the BLS cards, and upon the timing and control signals that are distributed throughout the *DØ* data acquisition system by the Serial Command Links (SCL). The ADF system is configured and monitored, over VME, by a Trigger Control Computer, as described below. Data from each ADF are sent to TABs, described below.

Trigger and Global Algorithm Boards

Trigger algorithms are implemented in the L1Cal in two sets of cards: the Trigger Algorithm Boards (TAB) and the Global Algorithm Board (GAB). The TABs identify EM, jet and tau objects in specific regions of the calorimeter using the algorithms described above. They also calculate partial global energy sums. The GAB uses these objects and energy sums to calculate and/or terms, which the TFW uses to make trigger decisions. Finally, the VME/SCL card, located in the L1Cal Control Crate, distributes timing and control signals to the TABs and GAB. It also provides a communication path for their readout.

There are 8 TABs and one GAB. In each TAB's main trigger data path, LVDS cables from 30 ADFs are received at the back of the card. The data from these cables are extracted and sent as individual bit-streams for each TT to ten TAB sliding windows algorithm (SWA) FPGAs for

processing. These chips also pass some of their data to their nearest neighbors to accommodate the data sharing requirements of the sliding windows algorithms. The algorithm output from each SWA is sent to a single TAB global FPGA. The global FPGA calculates regional sums and sends the results out the front of the board to the GAB, over the same type of cables used for ADF to TAB data transmission. This data transmission occurs at a clock rate of 636 MHz.

The path of trigger data through the TAB/GAB system is dictated by the necessary inputs for the sliding windows algorithm and by timing contraints. Each of the eight TABs receives data covering a 40×12 region in $\eta \times \phi$ space. Each of the ten SWA chips on each TAB finds objects that are 4×4 in an $(\eta \times \phi)$ grid. To do this, it requires a 9×9 region of input TTs.

The GAB receives data containing regional counts of Jet, EM, and Tau objects calculated by the TABs and produces a menu of and/or terms, which is sent to the *DØ* trigger framework. LVDS receivers each receive the output of two TABs, synchronizing the data to the GAB 90 MHz clock. The synchronized TAB data from all four LVDS FPGAs is sent to a single GAB S30 FPGA. This calculates and/or terms for the whole system and sends them to the trigger framework. There are five 16-bit outputs on the GAB, although only four are used by the framework. The fifth is used as a status register, the use of which is described in detail in this thesis. Upon receiving an L1 accept the GAB S30 sends data to L2 and L3. A total of 64 and/or terms are sent from the GAB to the Trigger Framework.

Monitoring via the Online System

Most components of the *DØ* trigger and data acquisition system are programmable. The Online System allows this large set of resources and parameters to be configured to support diverse operational modes, including those used during proton-antiproton collisions in the Tevatron and those used in the absence of colliding beams for calibration and testing modes. This forms a large set of resources as well as parameters which need to be configured before collecting data. The L1Cal fits seamlessly into this online system, with its online control software hiding the complexity of the underlying hardware, while making the run time programming of the L1Cal Trigger accessible to all *DØ* users in simple and logical terms. Several monitoring client programs, consisting of software that may run on a number of local or remote computers, display information useful for tracking L1Cal operational status.

The L1Cal TCC is the essential tool for monitoring and controls. The TCC interfaces to the 80 ADF cards, the 8 TABs and 1 GAB in each crate. There are three systems which use the TCC in order to access the system: the COOR Interface, the L1Cal Expert Interface, and the Monitoring Interface. The first two of these are used to configure and control L1Cal operations globally (COOR) or locally when performing tests (Expert). The Monitoring Interface collects monitoring information from the hardware for use by Monitoring Clients, described in this thesis. The TCC Control Software collects the monitoring information from the L1 Cal boards and makes it available to client programs. Three monitoring programs, now fully integrated to the *DØ* Online System and used by shifters, will be described here. All three are graphical user interfaces (GUIs) written with Python's TKinter class. One is a program which monitors the calorimeter trigger towers themselves and two which monitor the L1Cal IIB electronics, one for the TAB and GAB boards and one for the power system of the crates containing the TAB and GAB boards as well as other crates for other electronics.

5.2.1 Monitoring the Calorimeter Trigger Towers (TTs)

The Trigger Tower Monitoring GUI receives and displays the ADC counts of every trigger tower for an entire turn (i.e. 36 bunch crossings) of interactions about once every five seconds. It

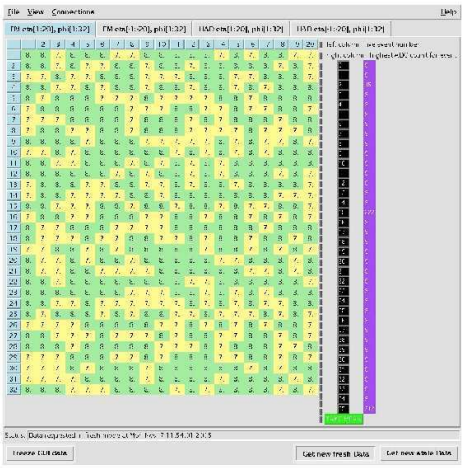

Figure 5.9: Monitoring information of the L1Cal trigger system made available via the TCC. On four panels, the ADC count read by each TT for a given bunch crossing can be seen. Data shown are from a mock TCC in 2005.

is ensured that this turn has an L1 accept. This is meant to be a random, representative sample of data that contains an interesting event which a shifter in the control room can use to quickly identify problems. The actual value that is displayed corresponds to the transverse energy measured by the TT in question. Each ADC count corresponds to 0.25 GeV of E_T and the baseline is set to 2 GeV or 8 counts.

Each of four panels displays a 20 × 32 ($\eta \times \phi$) grid, representing the TTs in a section of the detector. The four sections are positive and negative η in Electromagnetic (EM) and Hadronic (HAD) towers. The data for a single bunch crossing are shown at any given moment, while the user can scroll through all 36 bunch crossings in the turn. This GUI then has 20 × 32 × 4 × 36 = 92,160 ADC counts at any given moment, far too much for a user to take in at once. A variety of tools are provided to make this process more efficient for the shifter. The bunch crossing which corresponds to an L1 accept is flagged, so that the shifter can easily identify the "interesting" event and contrast its behavior to the others. A screenshot of the GUI is shown in Fig. 5.9.

Another useful tool is that the single highest TT ADC count for each bunch crossing is shown next to the button which brings up that data. An abnormally high count is then easily spotted. The value of ADC count for each TT determines the color that will be displayed behind the number itself: green corresponds to the nominal pedestal value of 8 counts, yellow to a value slightly below or above the pedestal and red to a value far outside. The precise settings of these bounds can be changed depending on the calorimeter's performance. Beneath the list of bunch crossings, a button "take mean" can be seen. This does a simple (unweighted) average of the data taken by each TT for all 36 bunch crossings, a tool which can be especially useful in identifying hot towers.

Data is automatically requested by the GUI once every five seconds in normal running. If the shifter's eye is caught, the program can be "frozen" so that the data shown can be examined. The shifter can then restart in either the normal mode of an automatic request every five seconds ("fresh" data) or can only request a single data set ("stale"). This flexibility allows the shifter to analyze abnormalities at their own pace. A status bar across the bottom displays the time

of the most recent data request so that the shifter can identify what they are looking at. This program was written and commissioned before the Run IIb Tevatron shutdown using mock data from the TCC and later edited to include several other functionalities.

5.2.2 Monitoring the L1Cal IIb Trigger Electronics

The electronics used by the L1Cal IIb Trigger include many custom boards which need to be monitored for errors. A stand-alone program has been written for this. The boards are also supplied power by a low voltage supply which needs to be monitored. This was written separately to be later included in the overall low voltage power supply monitoring GUI of the calorimeter system.

Monitoring the TABs and GABs

A program displays the the contents of all error and status registers in the TABs and GAB, associated with each (SWA) and Global chip on the TABs and with the LVDS and S30 (global) chips on the GAB. These registers indicate, among other information, synchronization errors on data transfer links, parity errors on each transfer, and bunch crossing number mismatches at various points in the TAB/GAB signal processing chain. The TCC reads the status registers from the TABs and GABs which is then sent to the GUI. Most bits in the status registers are latched, meaning that once set to an error state they remain so until readout, which returns them to the errorless state. The TCC keeps a count of these bits, which is an important feature: this allows for the GUI and the TCC to be asynchronous. The GUI is designed for both quick and accurate determination of errors in the system.

There is a 16-bit register on each chip which is the status register. There are 10 SWA chips on each TAB and a total of 8 TABS, as well as 4 LVDS chips on the GAB. This means that there are $(16 \times 10 \times 8) + (16 \times 4 \times 1) = 1{,}344$ bits to read out and display, not including the global chips on each board. To do this in an efficient fashion, the main page presents an "OR" of each bit on every chip for a given TAB so that a 16×8 grid is used. If an error appears, the user may click on the appropriate TAB to bring up a subpage which has the status of each bit on each chip on the TAB. In addition, an "or" of all errors on each TAB's global chip is displayed. The main page also displays the full readout of each bit on each of the 5 chips on the GAB. This includes the 4 LVDS chips and the S30 chip.

Each of the values displayed in this GUI actually has three different pieces of information about the error bits, two in number and one in color. The number on the left of any member represents the number of errors that the TCC has counted since the last COOR init. The number on the right is a local counter of the number of errors since the last user reset. The color of the member represents whether or not there was an error the last time data was read, where green corresponds to no error and red corresponds to an error.

Monitoring the Low Voltage Power Supply

The crates of electronics used by the L1 Cal IIB Trigger system are powered by supplies made by Wiener. There are 4 ADF power supply crates, one crate for the TABs and GAB, one for a control crate and one for the L1 Cal Track Match crate. Each of these power supplies needs to be monitored for the output voltage and current on each channel as well as the temperature and fan speeds. Four different GUIs were written, one for each crate type, which have since been integrated to the main monitoring page for the system.

The monitoring system includes alarm levels, set via the EPICS system control. Boxes containing the readout values have a green background during normal running and a yellow background if it is getting close to a dangerous level, a minor alarm. If a major alarm is reached, such as a voltage being off by more than 3 hundredths of the nominal value, the background will turn red and the crate power will trip off. As the powered devices are essential to good data taking, a power supply trip will cause the run to be paused.

5.3 Operations and Outlook

The trigger list for Run IIb was designed to select all physics processes of interest for the high luminosity running period and to run without prescaling at all instantaneous luminosities below 3×10^{32} cm^{-2} s^{-1}. The entire Run IIb L1 trigger menu normally produces an accept rate of up to 1800 Hz. It includes a total of 256 and/or terms, of which 64 come from L1Cal. These include one- two- and three-jet terms with higher jet multiplicity triggers requiring looser E_T cuts, single- and di-EM terms without isolation requirements for triggering on high energy electrons, single- and di-EM terms with isolation constraints designed for low energy electrons, tau terms which select jets with three different isolation criteria, topological terms - such as a jet with no other jet directly opposite to it in ϕ - which targets specific signals that are difficult to trigger and lastly missing E_T terms. These terms can be used individually or combined, and together form the $D\emptyset$ L1 triggers in the TFW.

Calibration

There are three main methods used to ensure that the E_T measured by individual trigger towers is accurate: online adjustment, a calorimeter pulser and offline gain calibration. Calibration here means that one output count corresponds to 0.25 GeV of E_T and that the zero-E_T baseline is set to eight counts. This calibration process has ensured smooth running. Deviations from the zero-E_T baseline can be seen visually via the Trigger Tower Monitoring GUI, discussed above.

Measured Rates

Since Run IIb has begun, measured trigger rates using the new algorithms and trigger list have been consistent with those based on extrapolations of Run IIa data to Run IIb instantaneous luminosities, shown in earlier Fig. 5.7. The total trigger rate observed using the new Run IIb list, to which L1Cal contributes more than 50% of the events, fits into the bandwidth limitations of the experiment. A Run IIa trigger list, designed to give the same selection efficiency as the Run IIb list above, would have exceeded these limits by a factor of two or more. The project is thus considered a success. So far it has performed exceptionally well, achieving background rejection factors sufficient to fit within the bandwidth limitations of the experiment while retaining the same or better efficiencies as observed in Run IIa for interesting physics processes.

Outlook

The use of a sliding windows algorithm for the $D\emptyset$ Level-1 Calorimeter Trigger in the Run IIb at the Tevatron was the pioneering effort in application at a hadron collider experiment. Its success has been crucial in the identification of interesting physics events at $D\emptyset$ in the increased luminosity running of the Tevatron.

The large increase in statistical power lent by the larger data set has opened new possibilities for physics at the Tevatron, including greater precision in critical measurements like the top

quark mass and W boson mass, the possibility of detecting or excluding very rare Standard Model processes (including production of the Higgs boson), and greater sensitivity for processes beyond the Standard Model. *DØ* has recently achieved sensitivity to the some of the rarest of Standard Model processes, including the production of a single top quark together with a b quark[25] and ZZ production [26, 27]. In a combined result with CDF, the Tevatron experiments have recently published the accelerator's first direct exclusion of certain masses of the Higgs Boson [28]. These advances hinge on larger data sets, made available by Run IIb, and exploited using the selection methods of the trigger upgrade which has been described.

The results attained recently bode well for the first level calorimeter trigger systems at ATLAS and CMS at CERN's LHC, both of which use similar methods. At the LHC experiments, as in *DØ*, monitoring and verification of the trigger system will be essential to the efficient identification of interesting physics processes. The monitoring programs described here have been shown to be useful and effective to accomplish this.

6 Calibrating, Commissioning and First Data with the CMS Cathode Strip Chambers

The CMS muon system is designed to be able to reconstruct the sign of a muon with momentum up to 1 TeV when trajectories are matched to tracks from the inner detector [20]. With this goal in mind, the muon system has been designed with three types of chambers: drift tubes (DTs) and cathode strip chambers (CSCs) for precision readout and resistive plate chambers (RPCs) for fast response. A schematic showing one quadrant of the muon system is shown in Fig. 6.3. The DTs are used in the barrel region, covering $|\eta| < 1.2$. The barrel system is mostly outside of the 4 T solenoidal field. The CSCs are used on the endcaps, covering up to $|\eta| < 2.4$, mostly within the magnetic field. For both the endcap and barrel, RPCs are used for triggering. A detailed description of the full muon system can be found in [29]. Work involving the preparation of the CSCs is described here. This includes commissioning the on- and off-chamber electronics and calibrating the chambers. A first analysis of cosmic muon data taken with a portion of the CSCs is then presented.

6.1 The Cathode Strip Chambers (CSCs)

The cathode strip chambers (CSCs) are an essential part of the CMS detector, providing precision information on muons passing through the endcap region of $1.2 < |\eta| < 2.4$. There are 468 chambers in total, arranged in four layers going with increasing $|z|$ on each side of the detector. Each chamber has 6 layers (described below), allowing for a directional vector to be reconstructed in each chamber and an overall flight path to be reconstructed with several chambers. As the system is inside the solenoidal magnetic field, this allows for momentum reconstruction.

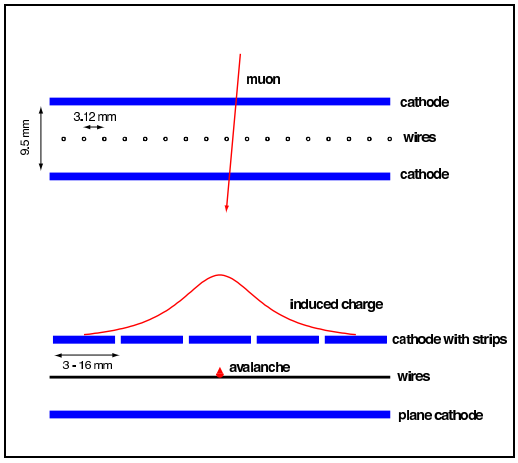

Figure 6.1: The charge avalanche caused by a muon passing through a CSC.

Detection Principle and Design

A CSC detects a charged particle passing through by the measurement of its ionization of a gas [29]. Planes of anode *wires* and cathode *strips*, oriented perpendicular to one another, are separated by a gas mixture of Ar-CO_2-CF_4 (30%-50%-20%). In each chamber this repeats six times, each one called a *layer*. A muon passing through ionizes the gas, and the high potential of $\approx$ 3kV creates a drift of the ionized particles. A charge then builds up on the wires and strips, allowing for a determination of the muon's position on each of six planes within a chamber. This is shown graphically in Fig. 6.1. The design of a chamber is shown graphically in Fig. 6.2 (a) and the layers' segmentation is shown in Fig. 6.2 (b). When the chambers are installed, the cathode strips are oriented in the radial direction and anode wires are oriented in the azimuthal direction. The strips provide a precision measurement in ϕ while the wires provide precision timing and a coarse radial measurement.

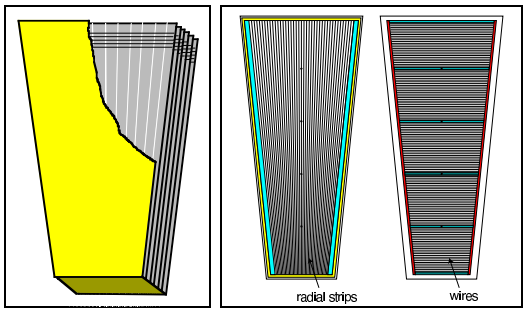

(a) The layers of a CSC.

(b) The anode and cathode planes.

Figure 6.2: The Cathode Strip Chamber design.

The endcap muon system is built on three *disks* on each side of the CMS barrel system. The CSC endcap system has four layers of chambers, each one called a *station*, going with increasing $|z|$. The first and third disk each have one station while the second disk has two, one on the front and one on the back. The disks themselves are made of iron, acting as stopping material for the muons. These disks double as the return yoke of the solenoidal field. The stations are comprised of *rings* of CSCs going outwards in r. Each station is symmetric in ϕ. This system creates a nearly seamless detection system for $1.2 < |\eta| < 2.4$. In this range, nearly 100% of the muons are expected to pass through 3 chambers, allowing full position and momentum reconstruction. A schematic with chamber nomenclature is shown in Fig. 6.3.

It should be noted that the chamber size varies slightly from station to station. In addition to these small variations, there are so called "half chambers" used for the first ring of the first station, ME1/1. These chambers are still within the stronger region of the 4 T solenoidal field. The smaller design is necessitated by the iron return yoke of the magnet. A summary of the basic properties of the various types of chambers is presented in Table 6.1.

6.2 Commissioning the CSCs

The entire process of commissioning the CSCs required a large group of people from many institutes around the world. The chambers themselves were built in Dubna and in the US,

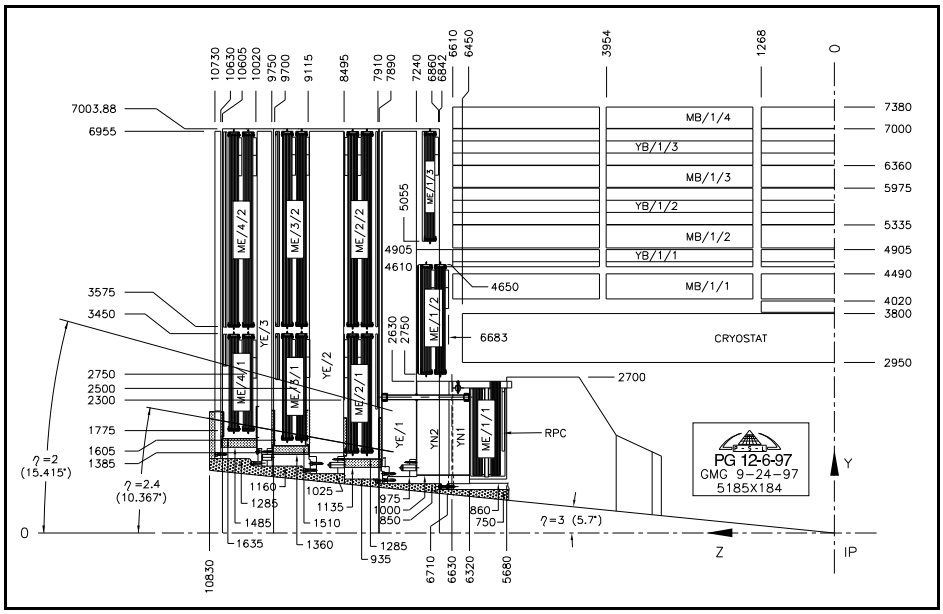

Figure 6.3: A schematic view of one quadrant of the CMS muon system. Here, ME stands for "Muon Endcap", and is followed by a sign representing the $+z$ or -z side of the detector. The first number following this corresponds to the station, the next to the ring. Therefore ME+3/2 corresponds to ring 2, station 3, on the $+z$ side. Chamber number can follow, as can layer and strip number, yielding a unique identifier for position within the system.

Table 6.1: Selected relevant chamber parameters.

Parameter	ME1/1	ME1/2	ME1/3	ME2/1	ME3/1	ME4/1	ME23/2
Active Region							
width (top), mm	487	819	933	1254	1254	1254	1270
width (base), mm	201	511	630	534	617	685	666
length, mm	1505	1635	1735	1900	1680	1500	3215

while custom readout electronics where built by three universities in the United States: The University of California at Los Angles (UCLA), Ohio State University (OSU) and Rice University in Houston, Texas. Commissioning at CERN was the first time that complete sets of the various components were tested together.

Front End Electronics

The CSC front end boards (FEBs) are designed to acquire precision measurement for full reconstruction and analysis as well as to provide information for triggering. The FEB readout path is shown in Fig. 6.4 (b). The anode FEBs (AFEB) and cathode FEBs (CFEB) send and receive information to a motherboard which is also on the chamber. A more complete description can be found in [29].

Peripheral Crates

The off-chamber electronics are in several places, but the first step is the *peripheral crate*. Each crate contains inputs from the motherboards of up to nine CSCs. The crate contains a slow control board (CC), generally connected via a coaxial optical cable directly to a computer during the commissioning process, a Clock Control Board (CCB), a Trigger Timing Control (TTC) board and a L1 muon trigger sorting board (MPC). Each peripheral crate contains a pair of boards for each chamber: a Trigger Mother Board (TMB) and Data Mother Board (DMB). In the back of each crate there is one card for each CSC which receives trigger signals from the RPC portion of the muon system, called the RPC to AFEB Transition (RAT) boards. The components interface with the crate's backplane to communicate with one another. The TMBs receive fast information from the AFEBs on each chamber, while the DMBs receive more detailed information from the CFEBs. The RAT boards receive input from the RPCs, which is sent to the appropriate TMB. The TMBs are synchronized by the TTC, while the all crate electronics - and the on-chamber electronics - are synchronized by the CCB. The TMBs send local trigger information to the MPC, which interfaces to the L1 trigger system. This system is illustrated in Fig. 6.4 for both the trigger readout path (a) and data readout path (b).

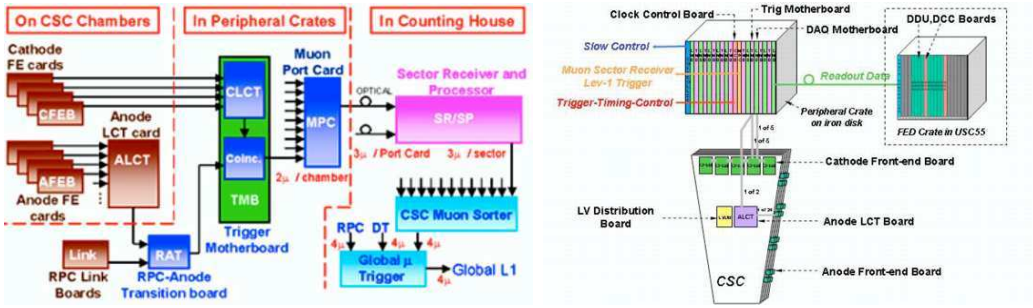

(a) The CSC Trigger electronics readout path.

(b) The CSC data electronics readout path.

Figure 6.4: The CSC off-chamber electronics.

Commissioning Process

Commissioning of these electronics has been done in several stages. Electronics boards were commissioned in full crates at a work station either in the ISR tunnel at CERN's Meyrin site or the B-904 electronics test facility at the Prevessin site. Shipments of boards arrived from UCLA, OSU and Rice. They were then put together into complete crates and connected to a test computer. This tested for damage during shipping and possible communication bugs between boards, the backplane and computers. Over the course of commissioning, several systemic issues emerged which were all taken care of. At the end of the electronics commissioning, there were ample boards declared to be functional.

Once commissioned in an electronics test facility, complete crates were brought to the CMS site and installed. Installation and cabling was a time consuming process. Previous to electronics installation, the electronics and the chambers had been tested separately. Once crates were installed, the full readout path could be tested. The most common problem this revealed was that cables had been accidentally swapped during the installation process, an issue easily rectified.

An outstanding problem remained of an easily blown fuse that was difficult to access. This had occurred in a small percentage of the chambers.

Most of the CMS detector was built on the surface at CERN's Point 5 and then lowered piece-by-piece. After being commissioned on the surface, the final commissioning was done in the cavern at Point 5. The CSC system is considered to be ready for data taking.

6.3 Calibrating the CSCs

The work of calibrating the CSCs is done by researchers at Northeastern University and Ohio State University. A number of tests were devised to calibrate the chambers in order to ensure the best understanding of the chambers and provide the finest resolutions possible. Most of these tests make use of a controlled pulse. During the CMS commissioning phase, programs have been developed that will be able to run in detector downtime, such as between stores. Data will be written to a database and stored. Some of the tests, their implementation and automation are described here. After tests were taken, they were processed using C++ based code in the CMS software framework. Macros in ROOT created a visual representation of the data. A set of perl scripts then made these ROOT plots accessible via the web. This system will eventually be a part of the CMS DAQ software.

6.3.1 Calibration Tests

While there are about 20 calibration tests in all, a few crucial examples are explained here. For each of these tests, data are taken using a designated program and written out as raw files. These raw files are used as input to programs written using the CMS software framework.

Pedestal Noise

The simplest of the calibration tests involved a readout of the system with no pulsed data in it. This provides a standard so that a baseline can be established. In a given chamber, RMS noise is calculated. Chamber to chamber fluctuations can be seen for a given data set. Using the database, drifts over time can be seen. A sample pedestal noise plot is shown in Fig. 6.5.

Noise Correlation

To ensure that the channels are responding properly a simple test of the noise was used. This involved reading out channels with no pulse in them, inputing the data to a channel-by-channel matrix and searching for correspondence.

Crosstalk

Given the environment at hand, minimizing wire crosstalk is essential. Having channels closer together increases spatial resolution but it also increases crosstalk. This induced charge has been measured, and a general trend has been established. It is seen to be within design range. The RMS and mean are calculated for each strip, layer and entire chamber. An example is shown in the browser image in Fig. 6.7.

Gain

The response of the CSCs to different input pulses has been found. Plotting the input pulse (in ADC) versus the charge essentially yields the gain of the system. This is expected to be linear

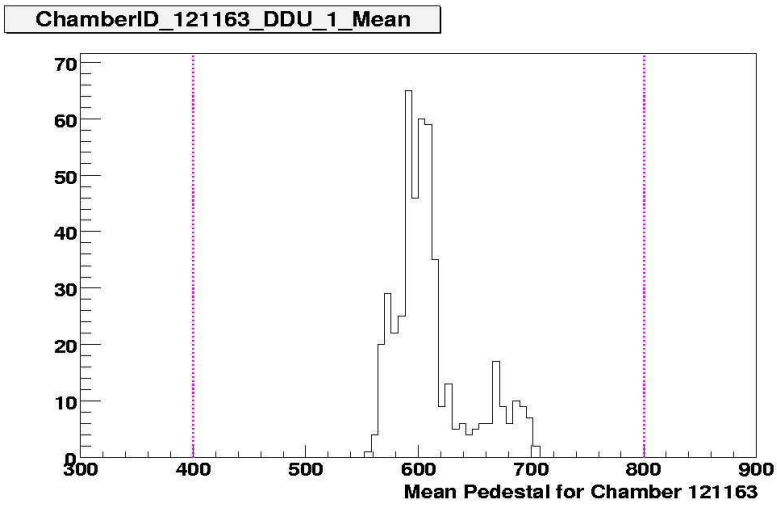

Figure 6.5: Mean pedestal values for all strips and layers in a given chamber. Values within the lines mark acceptable values.

and it has generally found to be so. As shown in Fig. 6.6, higher ADC values give a "thicker" line, which is to say that there is a small variation in the output versus the input. This variation is proportional to the input.

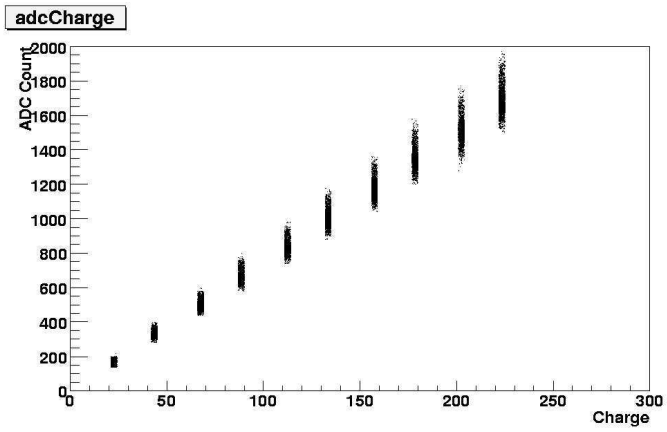

Figure 6.6: A plot of charge versus ADC count, whose slope corresponds to gain.

6.3.2 Automating the Tests

Much had to be done in order to bring these tests to a state where they could be used for the large scale of all 468 chambers. In order to test the entire system at once it was necessary to generalize the software. In the CMS software, configuration files which specify the input data file and other specifics are used. For each test, each chamber is pulsed separately and thus has a unique file. These filenames contain the information necessary to begin processing the data; namely each has the name of the test that the file corresponds to and a name which corresponds to the chamber used. Thus the file input name is enough to begin processing the data.

Configuration Files

A perl script has been used to parse the file name. It is still necessary for the user to specify which calibration test a given file corresponds to. This is due to irregularities in the file nomenclature coming from the pulse program. The perl scripts - one for each test - creates a configuration file which uses the filename as an input. This configuration file can then be used to execute the main data analysis program. This program outputs a ROOT file containing a tree with all of the information which needs to be displayed. In principle, a single shell script can be run and check for new files, which could then be automatically processed. This has not yet been implemented.

ROOT Macros

Once the C++ program, written using the CMS software, is run, a ROOT macro processes the output. The output of the C++ program contains a ROOT tree with all necessary information. The ROOT macro does no processing but only creates plots from the tree. In addition to plotting the relevant information at all scales possible, a "flag" system has been created so that a user can quickly see if there are any abnormalities. Flags are displayed on a chamber, layer and strip basis. This flag system uses levels of normal (black), warning (green) and trouble (red). A user can then check if there are any problems. One can look at the chamber-by-chamber plot and identify a problem in a given chamber. They can then look at the layer-by-layer plot for a problematic chamber and identify a troubled layer. They can then look at the strip-by-strip plot of that layer and precisely identify the problem. A directory hierarchy is created which reflects the levels of the system. For a test, such as crosstalk, which contains information down to the strip level, a folder is created for the overall plots - such as the flags - which also contains one subfolder for each chamber. These chamber subfolders contain sub-subfolders, one for each layer. These then contain strip-by-strip plots. This hierarchy is evident in the left hand side of the browser shown in Fig. 6.7.

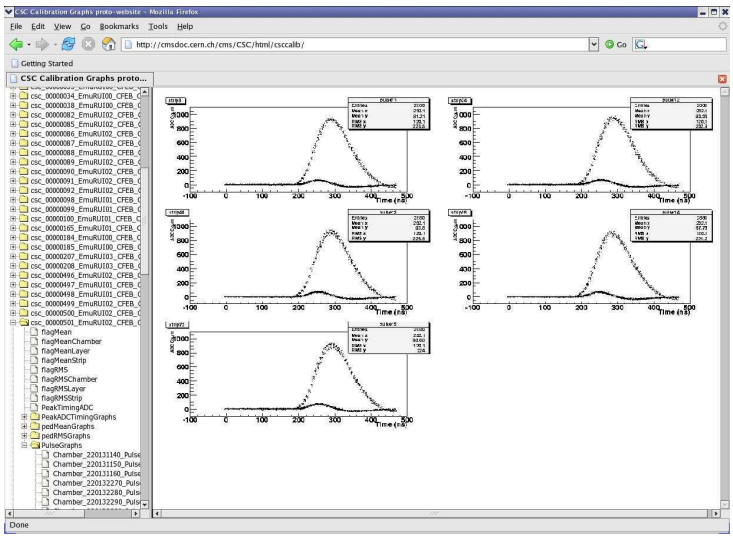

Figure 6.7: The Web browser used. The linux hierarchy which is mimicked by the browser can be seen at left. The specific plots shown here reflect the crosstalk of every strip on each layer of a given chamber. The large curve is the main pulse and the small curve is the crosstalk.

Web Display

A website is directly linked to the output of the ROOT macro. A simple freeware code, Tigra Tree Menu [30], is used to display the directory hierarchy. One can then click through the names of plots and quickly judge if there are problems. A perl script updates a "tree" file which is used as an input to the website's html files. The website is thus updated dynamically when there is a new directory hierarchy, representing a new set of calibration data. A screenshot of the website is shown in Fig. 6.7.

6.4 Reconstructing Cosmic Muons with the CSCs

In the autumn of 2006, CMS undertook the Magnet Test and Cosmic Challenge (MTCC), an ambitious period of running components of the still under construction CMS detector with full magnetic field on the surface of Point 5. The goals of the MTCC were to run the 4 T solenoid at nominal running strength and map the field at full strength, as well as to test hardware, triggers and DAQ by taking data with cosmic muons. During the first phase of running, MTCC 1, cosmic muon data were taken with a full magnetic field using the Tracker, ECAL, HCAL and both the DT and CSC components of the muon system. In a few thousand of the $\sim$ 25 million events recorded in MTCC 1 a cosmic muon is seen by all subsystems. In the second phase, the tracker and ECAL were removed and replaced by a field mapper and the field was successfully mapped. Around 250 million events were recorded in MTCC II. A first analysis of this data is presented here, taken with a magnetic field of B=3.8 T. This work relied on offline reconstruction software, a cosmic Monte Carlo (MC) generator and trigger emulation software. The ultimate goals of this analysis were to test the complete offline analysis chain, including generating MC, and to provide basic checks of the functionality of the CMS Endcap Muon System Cathode Strip Chambers.

6.4.1 Event Reconstruction

A number of software programs, developed independently by members of the group, are used together to form the event reconstruction chain. The chain for reconstructing data begins by reconstructing the hits a muon left in a chamber. This is followed by building *segments* inside each chamber that includes a direction built from hits. Track reconstruction is done using these segements. Momentum reconstruction follows segment building; this had not yet been successfully implemented for cosmic muons at the time of analysis.

Running Conditions

The CMS detector was built in several slices on the surface of Point 5, and has since been lowered $\sim$ 100 m down to the cavern where it ultimately resides. At the time of the MTCC in the middle of 2006, a small portion of the then not yet complete detector was able to receive services such as cooling and sufficient electricity. This section was commissioned on the surface and used in the MTCC. In terms of the CSC endcap system, a 60° "slice" of the system on stations ME+1, ME+2 and ME+3 was used on the +z side of CMS. The CSC contribution to the MTCC was comprised of 36 chambers, about 8% of the complete system. There were nine chambers on each ME2 and ME3 - all chambers in both rings in the slice - and the 18 ME+1/1a, ME+1/1b, ME+1/2 and ME+1/3 chambers in the slice. A graphical representation of the chambers used is shown in Fig. 6.8. The figure implies the possibility of using ME+4/1 chambers, none of which were used in this analysis.

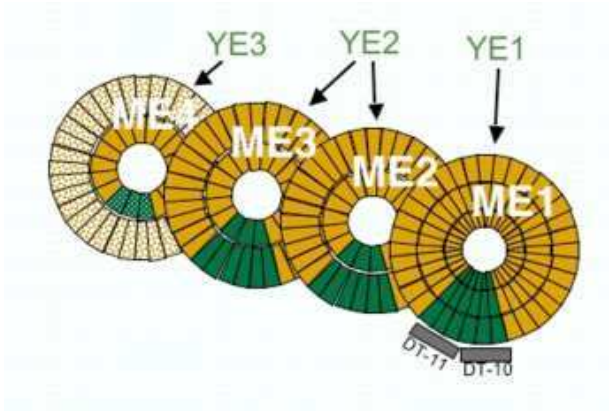

Figure 6.8: The CSCs used in the MTCC. Chambers on ME+4 where not used in this analysis.

Monte Carlo Generation

A cosmic muon Monte Carlo generator, CMSCGEN, has been developed to allow for comparisons of data to MC [31]. It is based largely on the cosmic generator from the L3+Cosmics experiment at LEP, L3CGEN [32], allowing for a precise parameterization based on earlier experience. In CMSCGEN simulated cosmic muons are generated on a flat surface, propagated onto a cylinder and then into the CMS detector. Work has been done by the CMS DT team, providing a basic check of momentum distribution reconstructions of data and MC [31]. This followed earlier studies regarding triggering on and reconstruction of muons of both beam halo and cosmic origin [33]. The CMSCGEN software allows for run-specific flags, most notably for simulation with and without a magnetic field and for data taking at the surface of Point 5 (altitude $\sim$ 400 m above sea level) versus in the cavern under $\sim$ 100 m of rock.

The CMSCGEN software is the source of all cosmic muon MC in the analysis presented here. Large-scale cosmic muon MC production had not yet been undertaken by CMS, necessitating it to be done on an individual basis. A sample of 300k events were generated "by hand" using the batch computing system, lxbatch, at CERN. This sample simulates a magnetic field of B=3.8T and data taken at the surface of Point 5 to match the running conditions of the data used here.

Trigger Emulation

The event reconstruction chain is in principle the same for MC generated events as for recorded data. The main difference is in triggering: while the trigger that is used for each event in data is stored, a trigger must be simulated for MC events. An emulator which has been verified to mimic the trigger boards' firmware is used.

The CSC trigger system uses Local Charged Tracks (LCTs) to generate local triggers. Cathode Local Charged Tracks (CLCTs) and Anode Cathode Charged Tracks (ALCTs) are calculated using a threshold to cut on the charge buildup. This is done by the on-chamber electronics, in the CFEB and AFEB respectively. A logical "AND" yields an LCT which is a basic trigger object. The cathode planes, and therefore CLCTs, have finer position resolution while the anode planes, and therefore ALCTs, have better time resolution. In order to compare Monte Carlo and data, a trigger emulator has been devised by members of the group. This software mimics the firmware run on the electronics. It has been shown, on a bit-by-bit level, to do so accurately. Despite this accuracy, inconsistencies were found which will be returned to later.

6.4.2 Reconstructed Cosmic Muons and MC Comparisons

A single run, MTCC II run 4386, deemed by all measures to be "good", is used in this analysis. Early tests showed that this run was consistent with others in general distributions. A total of eleven data files containing $\sim$ 141k events with 1.8 million reconstructed hits are used here. The trigger used here was a single LCT in a single chamber. Later runs tested more complicated trigger configurations, not analyzed here.

Due to the fact that only a portion of the entire CMS detector was used in the MTCC, MC generation was highly inefficient. A total of 240k generated events have only 814 events passing two cuts, that 1) there is an LCT in one of the MTCC chambers and 2) there are at least four reconstructed hits. This second condition stems from a feature of trigger emulation: in the algorithms used 4 reconstructed hits are needed. In the MC sample, however, a small number of events with fewer than 4 reconstructed hits were seen. This second condition was applied to remove them from the analysis sample.

In all of the plots shown here, the MC is normalized to the data. In these plots, MC is blue and data is red.

Reconstructed Hit Position

Using the standard hit reconstruction package, CSCRecHitB, muon hits are found. In the data an LCT requires at least four hits while in the MC this had to be explicitly imposed. The spatial distribution of hits in the x- and y-planes for the chambers are plotted for all events in data and MC, shown in Fig. 6.9. In these coordinates, the origin (0,0) is taken to be the center of each chamber in both dimensions, with the narrow parallel side of the trapezoid representing the bottom. Given the trapezoidal shape of a chamber, a symmetric distribution in x is expected, flat until the edge of the base and dropping off for values of x that are wider than the base of the trapezoid. An asymmetric distribution in y is expected due the the trapezoidal shape as well because a narrower base implies less hits for the lowest values of y and more for the highest. Indeed, these distributions are seen in both the MC and data.

The data shown in Fig. 6.9 reflect the geometry of the CSCs. Relatively flat distributions of hits in the x- and y-planes are expected *a priori* without taking geometry into account. The geometry of the active region of each type of CSC is summarized in Table 6.1. All types of chambers except for ME+4/1 are represented in the MTCC. It should be noted that due to low statistics in the MC sample a chamber-by-chamber analysis has not been available. These data are therefore for all chambers, representing, an overall distribution.

In the x-direction, the active region bases of chambers (excluding ME+1/1 type) average to $\sim$ 600mm while the active region tops extend to $\sim$ 800-1300 mm, as shown in Table 6.1. Thus a steady falloff is expected beginning around $\pm$ 30 cm and ending at $\sim$ 65 cm is expected. This behavior is observed in data. The ME+1/1 chambers have an active region approximately covering the range (-10,10) cm. Due to this a small excess is expected in this region; this too is observed in data. The general features are present in the MC sample, namely it is symmetric about the x-origin and begins to fall off around $|x| \approx 30$ cm.

In the y-direction, the dimensions of the chambers would imply a steady rise from the lowest edge (where x is the smallest) to the uppermost edge. The chambers used, except for the ME+2/2 and ME+3/2, are around 170 cm long. There are therefore more events expected in the range of approximately (-85,85) cm. This is not expected to be symmetric but rather rise linearly with the distance from the bottom. The ME+2/2 and ME+3/2 chambers set the scale of the plot with a length of 321 cm. A similar asymmetry is expected. The confluence of these distributions is shown in Fig. 6.9.

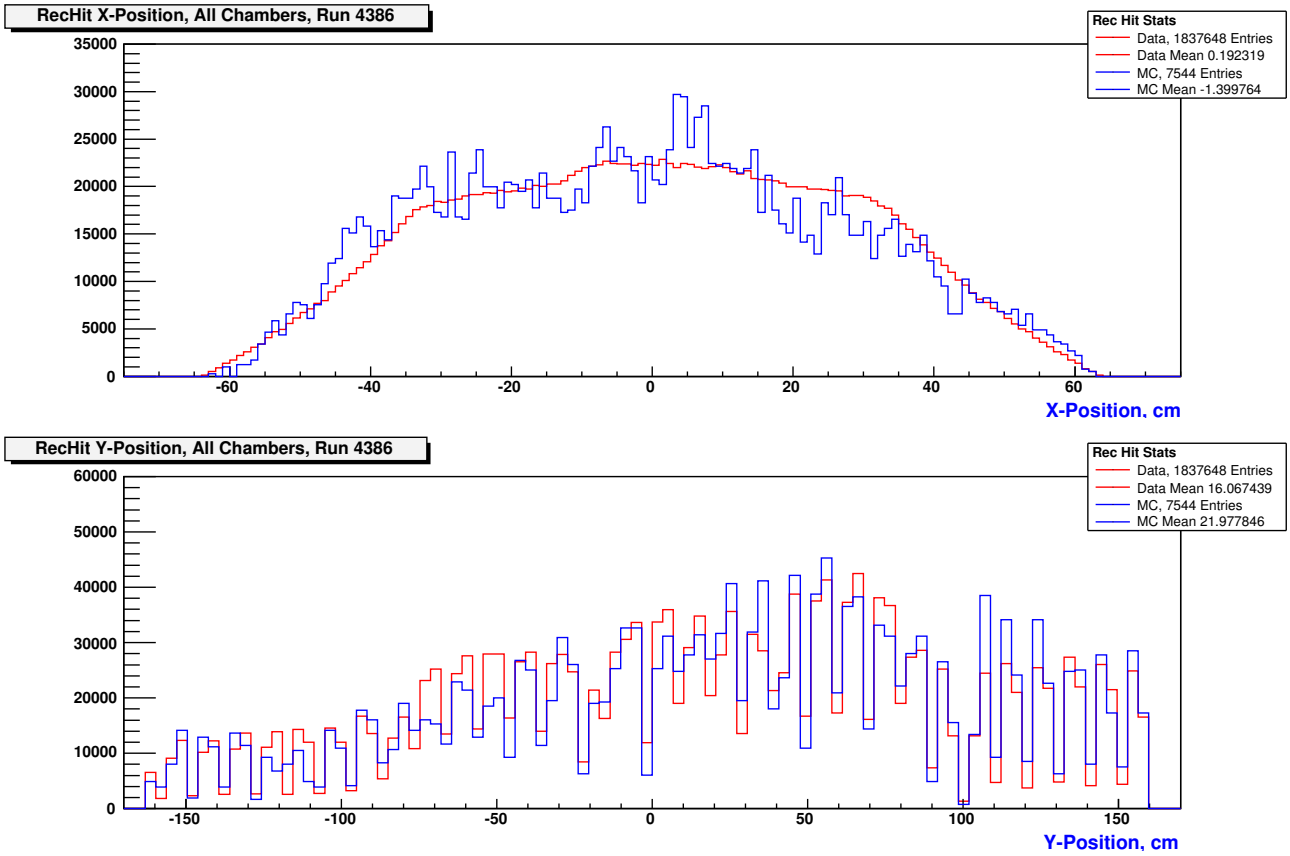

Figure 6.9: The reconstructed position of muon hits in the x- and y-planes for all chambers, plotted for both MC and MTCC data.

Reconstructed Hit Multiplicity

Any chamber with four or more hits is expected to generate an LCT. In turn, it is expected that any event in the data sample will have four or more reconstructed hits. Plotting the number of reconstructed hits in each event an asymmetric distribution is found, as shown in the bottom plot of Fig. 6.10. The anticipated sharp cutoff at four is found. The median is six hits per event while the mean is 12.25 for data and 9.03 for MC. The distribution is likely to reflect the implicit low momentum cutoff enforced by a minimum of 4 hits in a chamber. It could also be an effect of the relatively low rate of high energy cosmic muons, yielding a Poisson but not Gaussian distribution in momentum and thus hit distribution. The difference in mean between data and MC is not entirely understood; given that the median overlaps it is possible that this difference in mean reflects low statistics for MC, especially in the higher energy regime. It is also possible that this arises from triggering differences between data and MC, further discussed below.

Local Charged Tracks (LCTs)

As described earlier, the basic trigger object of the CSC system is the LCT. If a CLCT and ALCT are found, a general LCT is issued. This is sent to the trigger and data acquisition boards and the trigger and readout chains are begun. The Trigger Emulator used for MC has been verified bitwise to mimic the firmware of the trigger electronics at all levels.

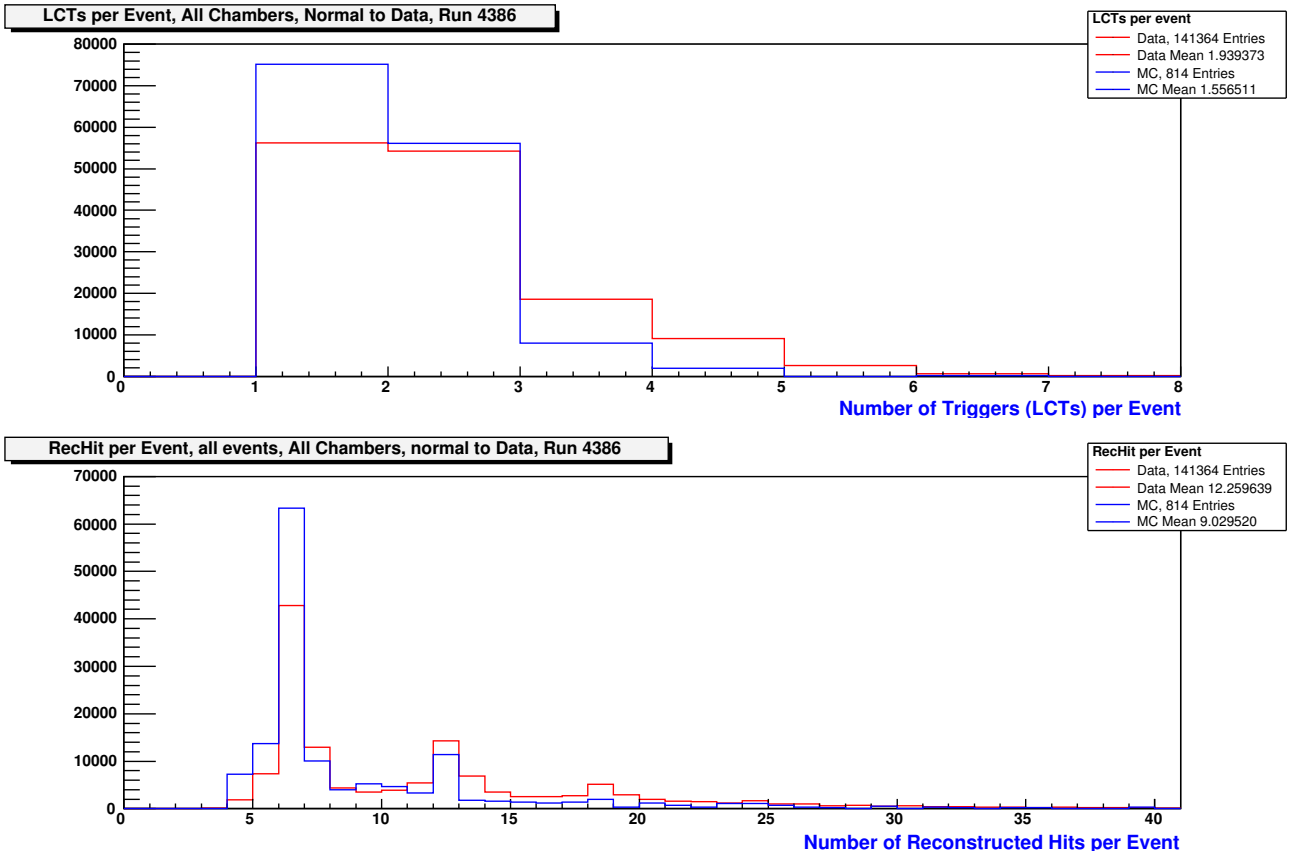

Figure 6.10: The reconstructed position of muon hits in the x- and y-planes for all chambers, plotted for both MC and MTCC data.

A comparison of LCTs found for the MC and data sample are shown in Fig. 6.10. While the same general trend is seen an exact overlap is not found. It should be noted that the scale of this plot is based on the data while the MC only reflects the 814 events used. A breakdown into the number of ALCTs and CLCTs per event is shown in Fig. 6.11. As expected, the general form of the overall LCT plot follows that of the CLCT. While the general features of the distributions are similar, the trigger experts in the group expected a far better agreement.

There are multiple possibilities that could contribute to the deviations seen. The simplest explanation is that the statistics are too low in the MC. This would be plausible for the bins of several LCTs were there are very low statistics, however the rate of one or two LCTs is not that low. Notably, the ratio of events with two LCTs to one in data is $\simeq 96\%$ while in MC it is $\approx 75\%$.

Given that a bitwise verification of the L1 trigger emulator has been carried out - the emulator returns the same information as the trigger system itself when the same input is given to both - one could conclude that the input is different for MC and data. While colleagues working with the CMS barrel DT system have verified the MC generator [31], CMSCGEN, it is still possible that the information passed from CMSCGEN to the emulator is too different at the byte level. It should be noted here that when the emulator was used, LCTs were found that corresponded to MC events containing 3 or fewer reconstructed hits. This was not anticipated. The triggers used in MTCC running were all either standard triggers or a simple variation of one. In this run, the simple trigger of having at least one LCT in a single chamber was used. Accordingly,

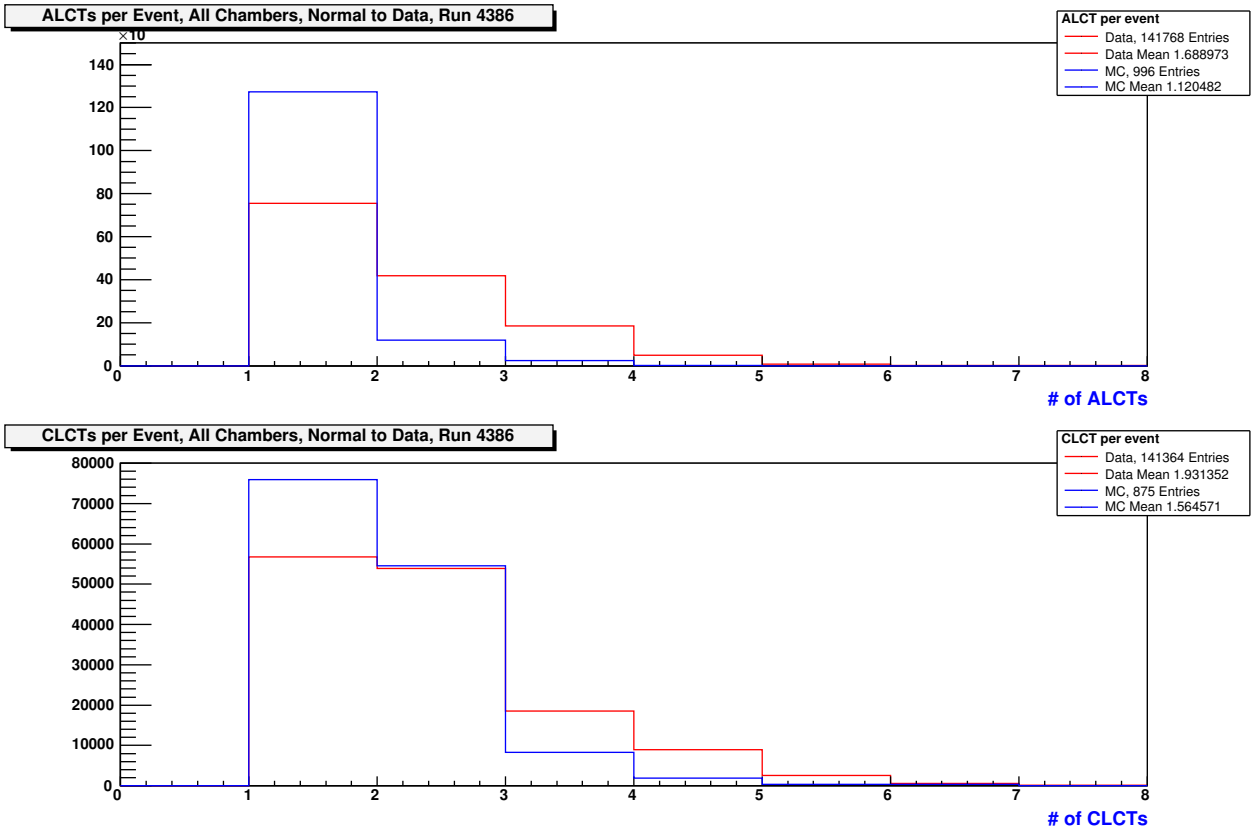

Figure 6.11: The number of ALCTs and CLCTs per event for both MC and MTCC data.

any event with at least one emulated LCT was used at first; the second condition of having at least four reconstructed hits was later imposed.

Another possible contribution to these effects is the implicit directionality of the trigger. The trigger is designed for muons coming from the interaction point. Given that the endcap system covers $1.2 < |\eta| < 2.4$, in beam (i.e. "normal") running conditions muons found in the CSCs would be expected to have an angle in the approximate range of $-45° < \theta < 45°$ with respect to the beamline. Furthermore they would all be coming from inside moving out, thus always hitting the innermost layer of a CSC first. These facts are taken into account in the design of the trigger. Taking these angular and directional considerations into account would have rendered the MC sample far too small to use, as the rate of incoming cosmic muons falls to zero as the angle approaches zero.

6.5 Outlook

The analysis presented here relied on the basic functionality of a very large number of components that had been more-or-less untested up to that point except for in stand-alone modes. The CMS collaboration took data in the Magnet Test and Cosmic Challenge (MTCC) with a fully functional magnetic field and all key detector sub-components participating. About 8% of the CSC system was used to take data with cosmic muons. In general the Cosmic Challenge has been deemed a "cosmic success" by its organizers.

The analysis presented here demonstrates that a full analysis is possible using the infrastructure available at CMS. Cosmic muons were triggered on, recorded and then reconstructed at the

level of hit and segment. Cosmic muon Monte Carlo data have been generated and the offline reconstruction chain shown to work on this. Basic checks have shown a first-order agreement in these data sets, representing that the fundamentals are functional. This analysis has been limited by time and external constraints, though it demonstrates that a full analysis will be possible at all necessary levels. Resolving the features identified here would require first and foremost a larger MC sample for better comparisons.

7 Determining the Precision of the ALFA Detector's Luminosity Measurement at ATLAS

A precise knowledge of luminosity is essential to physics analyses at collider experiments. The cross section of a physical process is one of the fundamental quantities that can be measured by such an experiment and its precise determination is a primary goal. To do this, precise knowledge of luminosity is needed. The total interaction rate, as defined earlier, is

$$R = \sigma_{tot} \times L, \tag{7.1}$$

where σ_{tot} is the total cross section of the collision and L is the luminosity provided by the accelerator. In order to determine the cross section of a specific process, σ_i, the rate at which such events occur is counted and the luminosity must be measured. Due to the partonic nature of the proton, measuring luminosity is not a trivial task at hadron colliders. The ALFA (Absolute Luminosity For ATLAS) subdetector at the ATLAS experiment is designed to measure the elastic scattering rate in dedicated runs. This will then provide an absolute calibration for LUCID (Luminosity measurement using a Cherenkov Integrating Detector), the main luminosity measurement subdetector [34]. In turn, this will allow for precise determinations of the luminosity provided to ATLAS by the LHC.

7.1 Luminosity Measurement at Hadron Colliders

The process of measuring luminosity via independent determination of the total cross section is a process which has been used before and is thus well described and studied [35]. The principle is to measure the elastic and inelastic collision rates separately, as well as to measure the elastic rate as a function of momentum transfer t. Using the optical theorem, the luminosity can be found.

Splitting the rate into elastic and inelastic parts, the total cross section is

$$\sigma_{tot} = \frac{1}{L}(R_{elas} + R_{inel}). \tag{7.2}$$

In scattering theory, the optical theorem relates σ_{tot} to the imaginary part of the forward scattering amplitude [36]. In terms of the quantities at hand, it states that

$$\sigma_{tot}^2 = \frac{16\pi(\hbar c)^2}{(1+\rho^2)} \frac{1}{L} (\frac{dR_{elas}}{dt})|_{t=0}, \tag{7.3}$$

where ρ is the ratio of the real to imaginary parts of the forward scattering amplitude. One can then obtain a relation between the measured rates and the total cross section, independent of luminosity:

$$\sigma_{tot} = \frac{16\pi(\hbar c)^2}{(1+\rho^2)} \frac{1}{(R_{elas} + R_{inel})} (\frac{dR_{elas}}{dt})|_{t=0}. \tag{7.4}$$

Having the total event rate and the total cross section thus yields the luminosity.

The measurement of elastic and inelastic rates is done separately. The inelastic rate is measured by the main detector, ATLAS, while the elastic rate is measured by an extremely forward detector, ALFA. ALFA is situated about 240 m from the ATLAS interaction region, surrounding the beamline. The ALFA detector uses position-sensitive scintillating fibers situated inside "Roman Pots", a device which has been used several times before [34] and has been shown to allow a precise approach to the beam. These Roman Pot devices were pioneered by the CERN-Rome Collaboration when working to study elastic scattering at the ISR [37], from which they derive their name. They have since been used by several other experiments at various accelerators.

The ALFA detector relies on using the LHC dipole magnets as a spectrometer. When there are elastic collisions with low momentum transfer (t), the interacting protons will loose only a small amount of their energy. As the dipoles and RF cavities of the LHC are precisely tuned to the ideal proton momentum, this slight change will send the proton off course. The ALFA detector detects these protons. In the small t approximation, the relationship between momentum transfer, beam momentum and angular deflection is approximated to be

$$-t = (p\theta)^2. \tag{7.5}$$

A 7 TeV proton that scatters at $-t = 0.005$, a typical low-t value, is deflected by ≈ 10 μrad, corresponding to a deviation of ≈ 2.4 mm from the beamline after a distance of 240 m. The Roman Pot housing allows a precise approach to the beam within this range.

The measurement of σ_{tot} with this method relies on extracting the differential of elastic scattering rate with respect to momentum transfer down to a momentum transfer of 0. Thus proton momentum must be measured over a range of t values. Doing this requires an extremely precise determination of the proton position in the ALFA detector. To resolve appreciable differences in values of t in this range it is anticipated that the ALFA detector needs an ultimate precision of $\sigma_x = \sigma_y = 30$ μm. This precision is one of the most important factors in the ALFA detector; indeed it is the main driving force in the design and implementation of the subdetector. Determination of this precision will be discussed in this thesis.

7.2 The ALFA Detector

The ALFA detector will be briefly described here. The entire detector is described in [34, 38].

The Roman Pots

The purpse of a Roman Pot is to allow a detector - housed within it - to be brought to within 1 mm of the beam. To do this, the pot must be integrated into the beam vacuum. The ALFA detector has components within the pot on either side of the beam. One of the benefits of the pot design is that it allows for independent control of each of the two sides for precision placement. A schematic of a Roman Pot can be seen in Fig. 7.1 (a) as well as a visualization of it integrated with the LHC beam vacuum in Fig. 7.1 (b). The positions of the pot with respect to the beam line is shown in Fig. 7.2 in both a retracted and working state. As the beam spot of the LHC will change from store to store, a precise dynamic positioning of the pots is necessary. The measurement of the pots' position is done using the overlap detector, which uses the same concept and a similar geometry as the one that is is used for the main detector.

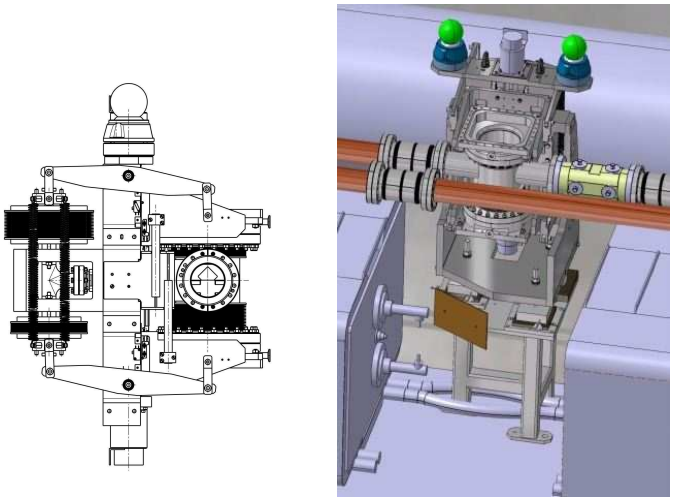

(a) A detailed sketch of a single pot.

(b) A simulation of a pot after integration with the LHC vacuum.

Figure 7.1: The ALFA Roman Pots.

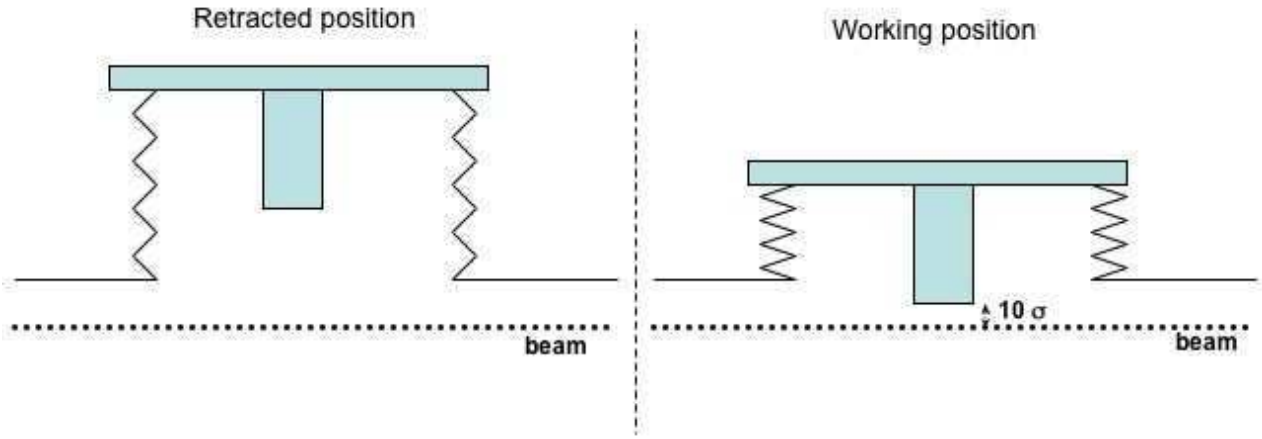

Figure 7.2: A schematic of the position of a Roman Pot with respect to the beam line in both close and near states.

Scintillating Fibers and Plate Geometry

Each side of the main detector within each Roman Pot contains a structure with a "UV" geometry comprised of ten plates which each contain two layers of 64 scintillating plastic fibers, visualized in Fig. 7.3. The fibers are glued to either side of the plate at 90° to one another. The ten plates are staggered with respect to one another at 0.5 mm / 10 $\times\sqrt{2}$ = 70.7 μm, yielding an effective fiber pitch of 50 μm. In principle this yields a spatial resolution of $\sigma_x = \sigma_y = 50$ μm / $\sqrt{12} = 14.4$ μm. Each layer on each plate is connected to a photo multiplier tube and an 8 $\times$ 8 pixel sensor, hence 64 fibers per layer.

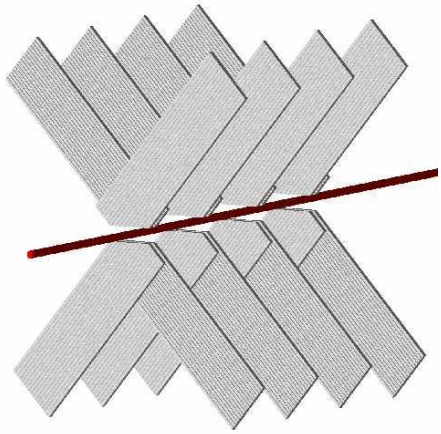

Figure 7.3: The geometry of a portion the ALFA plates on each side of every pot.

7.3 Determining Precision

In order to resolve proton position - and therefore momentum and momentum exchange - from the elastic collision a precision within 30 μm is needed. As stated earlier, the spatial resolution achievable using the ALFA detector with an ideal strip geometry is $\sigma = 14.4$ μm. It is therefore necessary to determine what the deviation from the ideal is and ensure that it is acceptable. Deviations from the ideal geometry have been found and are thought to arise from gaps between fibers which have an imprecise width [38]. The fibers are precisely 500 μm wide, however they are glued to the plates as well as to each other. The gap is on the order of 1/10 the width of the strips, ~ 50 μm. Variations in this gap size are believed to lead to a difference between ideal and real geometry. It is also possible that imprecise machining of the plates themselves could lead to shifted positions, notably if the angle between layers deviates from 90°. A detailed sketch of a single plate is given in Fig. 7.4.

To achieve the stated resolution, it is absolutely necessary that the position of each fiber is precise to within 70.7 μm, otherwise the staggering of plates becomes useless and the resolution sharply degrades. Much work on this has already been done [38] and a few points are considered here.

7.3.1 Measurement Devices and Method

There are two separate optical measurement systems being used to measure the strip position and plate geometry, one at CERN and one at DESY. The CERN microscope is out of the box, including software, while the DESY microscope is a custom design with software being written in Labview. Both microscopes use the same principle but are implemented differently. Each has variable light sources which can introduce an uncertainty to the measurements, as changing the light source changes the apparent size of the object.

Microscope at DESY

The FEC electronics clean room at DESY has provided a microscope. The lenses used in this setup are a Navitar 12x zoom lens, a Mitutoyo M Plan Apo 10x object lens and a 2x C-Mount adapter. The image is recorded using a Nikon DS-5M Digital Camera with 2560 × 1920 pixels (5 megapixels) and a 2/3" CCD. The apparatus is mounted on a 3-axis positioning table

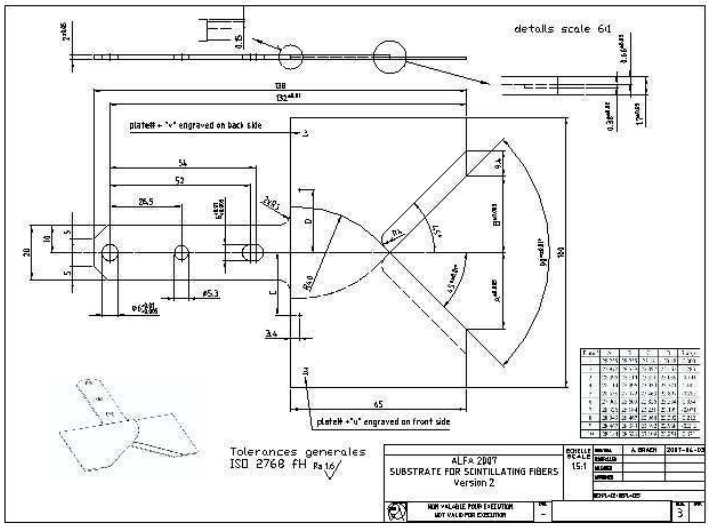

Figure 7.4: A detailed drawing of one side of a single plate used in the ALFA detector.

(x-y-z table). The table consists of two separate parts: an x-y-table (GT9-NSMA) and a z-table (LT8-LBMA) from Walter Uhl. The tables are driven by a step motor allowing for a positioning precision of 2-3 μm. The step motor is connected to a PCI-card in the PC and can be controlled through the LSTEP-PCI API by LANG [38].

Two light sources are available here: a moveable gooseneck light source and a ring light mounted to the microscope around the final objective. For the sake of consistency, now only the ring light is used and it is always used at full power.

Microscope at CERN

The team measuring the geometry of ALFA at CERN is using the SmartScope CNC 250 from OGP. According to the company, a precision of better than $(3 + 6L/1000)$ μm, where L is the measured length in mm, is achieved. The machine is equipped with a 12x zoom and variable light sources.

Measurement, Disagreement and Problem Resolution

In order to measure the fiber positions it is necessary to first determine a coordinate system. The measurements taken at both the DESY and CERN sites define the center of the coordinate system to be at the center of the lowest peg hole on the plate, visible in Fig. 7.5. For the same plates using the same coordinate system the measurements at CERN and DESY displayed a systematic shift of ≈ 70 μm. Given that this is the same as the design staggering, one possible explanation was a mislabeling of data. This was not thought to be the case, and another cause of error has recently been identified.

At the CERN site, a low-zoom image of the peg hole is taken which the provided software uses to find the center and thus the origin of coordinates. In this method, the user then zooms in to measure the fiber positions. At the DESY site, eight points around the edge of the circle are taken, then a circle is fit and the center found. This is done at the same zoom (the least possible) as is used to take the fiber position measurement.

It is a well known phenomenon in photography that changing the zoom of a lens changes

Figure 7.5: A photograph of an ALFA prototype 2 plate.

the center point, a form of optical aberration. In a zoom lens there are many internal lenses (generally 10-20) which move around to project the object image to a different distance before a final lens focuses that image to the recording plane. This causes a small shift in the recorded image.

The company which built the CERN microscope, OGP, claims that it yields a sensitivity of better than $(3 + 6L/1000)$ μm. It is notable that the website also claims that the device uses proprietary mechanisms to reach such a fine precision and thus does not explain the method used. It was therefore not possible to tell if the microscope corrects for this optical aberration.

The possibility of this shift as a cause of the 70 μm disagreement was tested at CERN. Indeed the shift seen between the two zooms used is ≈ 70 μm. This is now assumed to be the source of the disagreement. A technician from OGP is soon visiting CERN to examine the issue at hand.

7.3.2 Precision of ALFA

In order to determine the precision of ALFA, the measured positions of the scintillating fibers must be compared to the ideal position. A simple geometric comparison is used to ensure that the measured position differs from the ideal by less than 70.7 μm. A program has previously been written that takes into account the actual effect of the geometry on the precision of the result of ALFA, and that program has been used with the measured geometry.

Measurements and Deviations

Larger than anticipated deviations have been found on certain layers. Each strip on each layer has been fit to a line. These line fits for ideal and measured geometry are compared on a strip-by-strip basis for each layer. The position visualization plots show the line fits of the two data sets, measured and ideal. The variations plot shows the differences between the a- and b-parameters from these line fits for the two layers, "U" and "V".

An example of the reconstructed fiber positions on a "good" plate is shown in Fig. 7.6. The two sets of lines, blue corresponding to ideal and red corresponding to measured are in decent agreement. Accordingly, the difference in position, shown in Fig. 7.7, between measured and ideal is close to the expected Gaussian distribution, centered around 0 with a standard deviation below 30 μm.

An example of a plate with poor agreement is shown in Fig. 7.8. The lines are shown to be at an angle with respect to one another, with a select group of lines having a deviation in their b-parameter that is very large, shown in Fig. 7.9.

Work is being done to verify the measurement methods used. Namely, different methods are being used at the two test sites, and various methods are being tested at the DESY site as well.

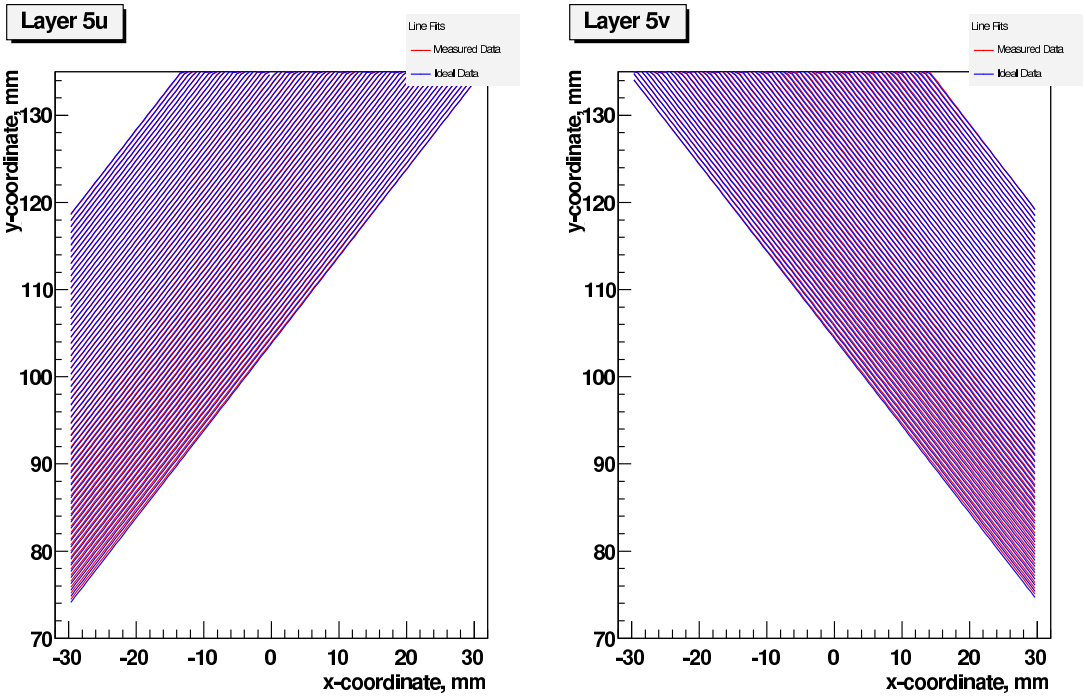

Figure 7.6: The reconstructed position of fibers in an ideal and measured geometry for one plate with relatively good agreement.

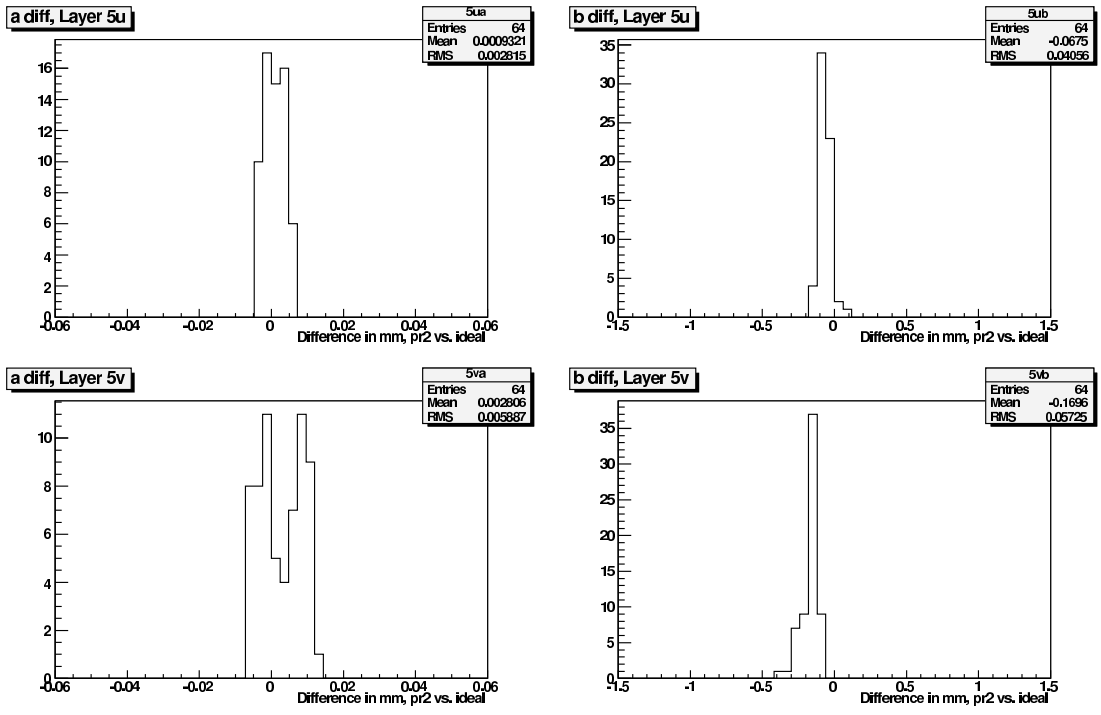

Figure 7.7: The differences between reconstructed positions of fibers in an ideal and measured geometry for one plate with relatively good agreement.

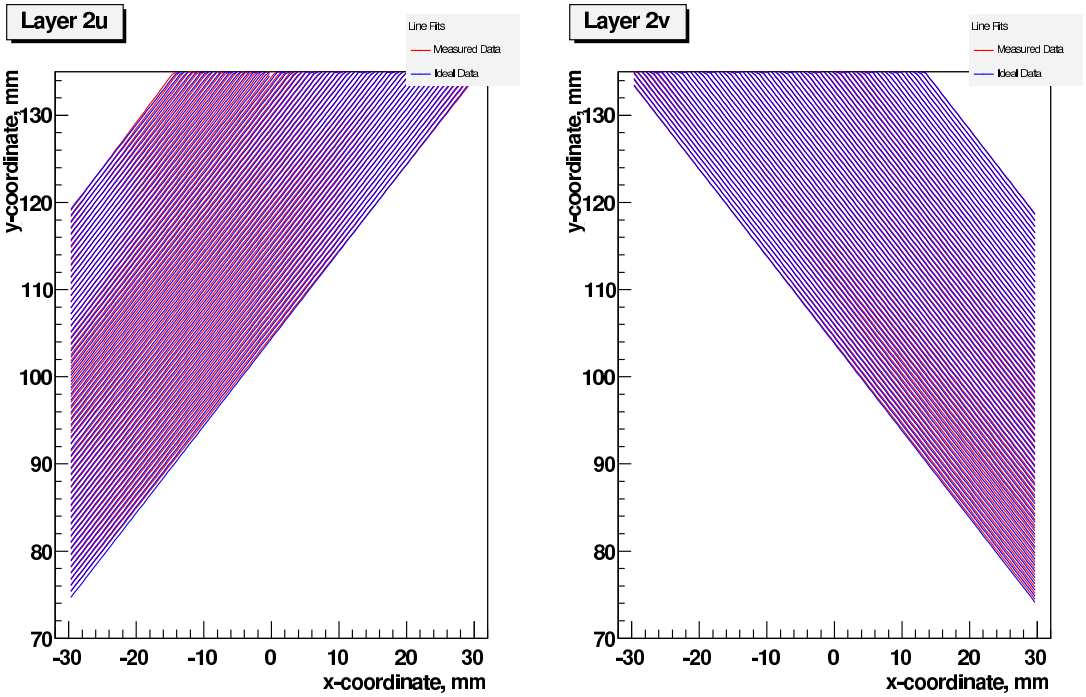

Figure 7.8: The reconstructed position of fibers in an ideal and measured geometry for one plate with relatively poor agreement.

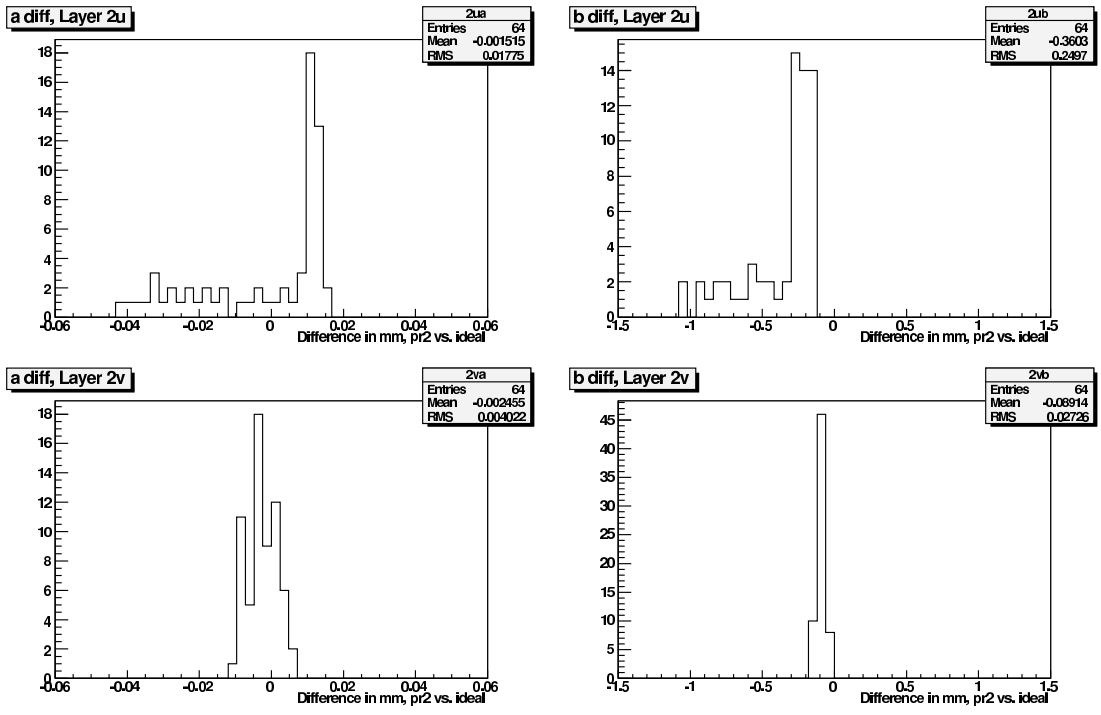

Figure 7.9: The differences between reconstructed positions of fibers in an ideal and measured geometry for one plate with relatively poor agreement.

This will help to determine whether the effects seen result from a poor measurement technique or are an accurate representation of an imprecise geometry.

Effect on Precision

A previously written program is used to estimate the effect of the fiber positions on the ultimate position resolution of a particle passing through. A plot of the precision using ideal geometry is shown in Fig. 7.10. The left and right hand sides correspond to the x- and y-measurements. These have the same precision as one would expect. The top row shows the precision of position measurements in the main detector, which is $\sigma_m = 14$ μm. The lower row shows the precision of the position measurement of the overlap detector, which is used to place the entire pot. In effect, the top plots show the precision of a measurement with respect to its pot and the second shows the precision of the placement of the pot with respect to the beam. The overlap detector's precision is $\sigma_o = 11$ μm. Adding in quadrature, the total uncertainty of a given position measurement using ideal geometry is $\sigma = \sqrt{\sigma_m^2 + \sigma_o^2} = 18$ μm.

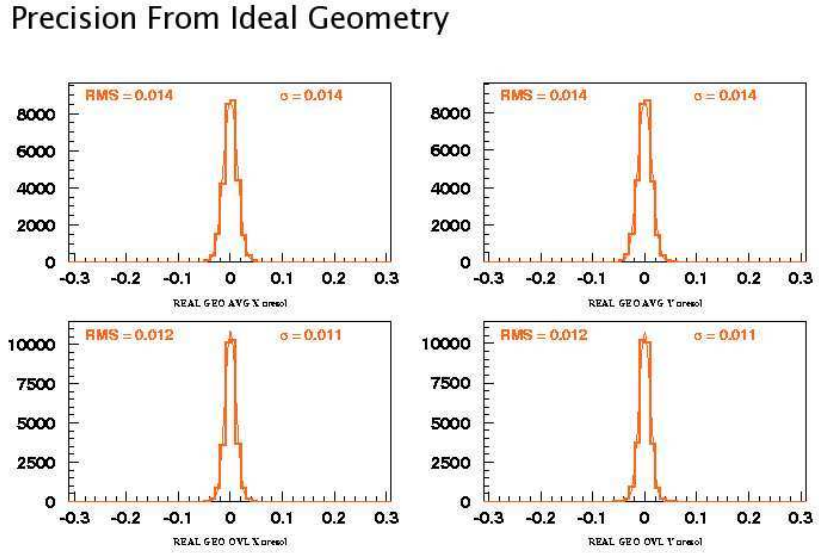

Figure 7.10: The precision of the ALFA fiber detector with ideal geometry, in mm. The top row shows the precision of a measured point with respect to the Roman Pot and the lower row shows the precision of the pot with respect to the beam.

While the measured position of the fibers deviates to large degrees as described earlier, the available precision is not degraded to such a dramatic degree. This has been simulated using the same code as above. The results are shown in Fig. 7.11. Just as in the plots shown above, the top row corresponds to the precision of an individual measurement with respect to the pot and the lower to the precision of the pot's placement. A symmetry between x- and y-values is seen here again. In this case however, the main detector has an uncertainty of $\sigma_m = 36$ μm and the overlap detector of $\sigma_o = 26$ μm. Adding in quadrature, $\sigma = \sqrt{\sigma_m^2 + \sigma_o^2} = 44$ μm.

It should be noted that while $\sigma = 44$ μm is out of the acceptable range, there is room for improvement which can bring this value to within the desired 30 μm. This can be done by using more precise measurement techniques and perhaps an improved gluing technique. Improvement of this precision continues to be a central goal of the project.

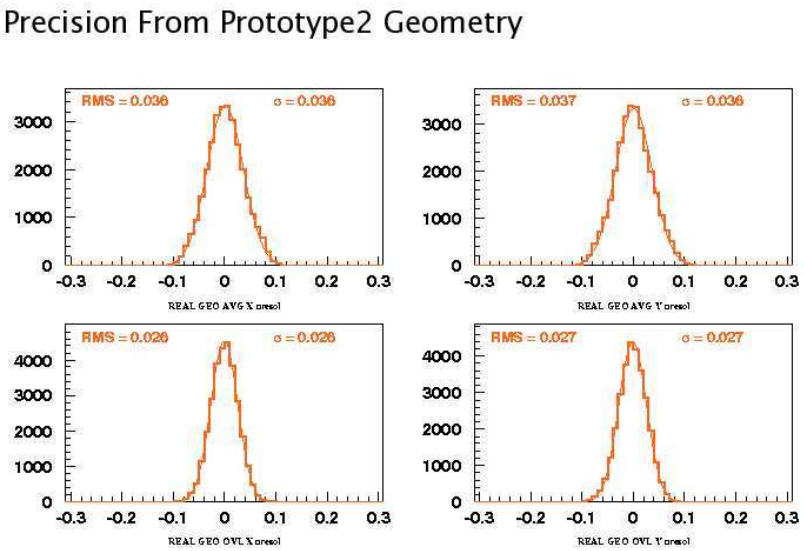

Figure 7.11: The precision of the ALFA fiber detector with the measured geometry of the Prototype 2 plates, in *mm*. The top row shows the precision of a measured point with respect to the Roman Pot and the lower row shows the precision of the pot with respect to the beam.

7.4 Towards Data Taking

It has been shown that some of the measured positions deviate from the ideal geometry to an unacceptable degree. The next thing that must be done is to verify the position measurements made to determine whether the deviations stem from the actual geometry or from measurement error. The only previously identified disagreement between measured data sets has been resolved. Now that the 70 μm difference between CERN and DESY measurements has been shown to stem from the use of different zoom lengths at CERN, a new set of measurements must be taken and the new CERN set compared to the DESY set. At the same time, work is being done to further automate the DESY measurement setup so that operator introduced inaccuracy can be reduced.

The stated goal of the ALFA group is to place one Roman Pot at its location, 240m on one side of the ATLAS interaction point, during the LHC shutdown in late 2008 or early 2009. This leaves a brief amount of time to resolve the differences. With a large and competent team comprised of many people from several institutes working to ensure that this goal is met, confidence is high.

8 Conclusion and Outlook

A number of the methods used in experimental high energy physics have been motivated and discussed here. The need for experiments to run is determined by the open questions in the Standard Model, as described earlier. The experiments which run at collider experiments are large and multi-faceted; they are built and used by a large number of collaborators, a number running into the thousands for each of the experiments at the LHC. It is thus important that each physicist has a cursory understanding of the entire system while the working groups of each subdetector ensure the highest quality possible to their fellow collaborators.

The importance of triggering to reduce data flow while maintaining physically interesting events has been reviewed. The system used at the *DØ* experiment for the first level of triggering on calorimeter data, described here, has proved to be extremely successful and thus will be used at the main LHC experiments, CMS and ATLAS. The monitoring methods, too, have proven successful.

The detection and reconstruction of muons at collider experiments is essential. Due to the long lifetime of muons and the large range of energies for which they have small ionization energy losses, muons are one of the most precisely measured fundamental particles at such an experiment. Many signatures of "new" physical processes involve high energy muons, and thus reconstructing their momentum and having unambiguous charge determination is essential. The recording of cosmic muons at CMS has been billed as a huge success and the work of analyzing the data is ongoing by other members of the working group.

Precise determination of the delivered luminosity is one of the two fundamental parameters used in determining the cross section of a physical process. Cross section is one of the most important parameters to be measured of any process because it is a constant parameter in nature and does not depend on the measuring device. The ALFA team intends to provide the ATLAS collaboration with a sufficiently precise determination of the luminosity. It will still take much work in the coming months to make ALFA operational to an acceptable degree.

When the LHC turns on and begins delivering collision data to experiments the first thing to be done is to understand the efficiencies and shortcomings of the entire detector. Doing this requires every subsystem, those discussed here and all others, to work quickly and effectively. It is only after a period of commissioning, testing and verifying that analyses of physical processes - the goal of these experiments - can begin.

Bibliography

[1] D. Perkins, *Introduction to High Energy Physics*, Cambridge U. Press, (2000).

[2] F. Halzen and A. Martin, *Quarks and Leptons: An Introductory Course in Modern Particle Physics*, John Wiley and Sons, Inc., (1984).

[3] M. Peskin and D. Schroeder, *An Introduction to Quantum Field Theory*, Westview Press, (1995).

[4] D. Griffiths, *Introduction to Elementary Particles*, John Wiley and Sons, Inc., (1987).

[5] S. Glashow, *Partial Symmetries of Weak Interactions*, Nucl. Phys. **22** (1961) 579.

[6] S. Weinberg, *A Model of Leptons*, Phys. Rev. Lett. **19** (1967) 1264.

[7] A. Salam, *Weak and Electromagnetic Interactions*, in: *Elementary Particle Theory*, Stockholm, (1968), Almquist and Wiksell, 1968 p. 367.

[8] P. Higgs, *Broken Symmetries, Massless Particles and Gauge Fields*, Phys. Lett. **12** (1964) 132.

[9] P. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, Phys. Rev. Lett. **13** (1964) 508.

[10] The Super-Kamiokande Collaboration, Y. Ashie *et al.*, *Evidence for an oscillatory signature in atmospheric neutrino oscillation*, Phys. Rev. Lett. **93** (2004) 101801, `hep-ex/0404034`.

[11] D. Politzer, *Reliable Perturbative Results for Strong Interactions?*, Phys. Rev. Lett. **30** (1973) 1346.

[12] D. Gross and F. Wilczek, *Ultraviolet Behaviour of Non-Abelian Gauge Theories*, Phys. Rev. Lett. **30** (1973) 1343.

[13] C. Quigg, *The Double Simplex*, `hep-ph/0509037`, 2005.

[14] M. Blaskiewicz and J. Brennan, *Bunched Beam Stochastic Cooling in a Collider*, Phys. Rev. Special Topics - Accelerators and Beams **10** (2007).

[15] S. van der Meer, *Stochastic Cooling and the Accumulation of Antiprotons*, Rev. Mod. Phys. **57(3)** (Jul 1985) 689.

[16] The Tevatron Luminosity Upgrade Group, *Tevatron Run II Handbook*, (1998), `http://www-bd.fnal.gov/lug/`.

[17] The Fermilab Accelerator Division, *Tevatron Luminosity*, (2008), `http://www.fnal.gov/pub/now/tevlum.html`.

[18] *LHC Design Report*, 2008, `http://ab-div.web.cern.ch/ab-div/Publications/LHC-DesignReport.html`.

[19] The *DØ* Collaboration, V. M. Abazov *et al.*, *The Upgraded DØ Detector*, Nucl. Instrum. Meth. A **565** (2006) 463, `physics/0507191v1`.

[20] *CMS Technical Design Report, Vol. 1: Detector Performance and Software*, 2000, `http://cmsdoc.cern.ch/cms/cpt/tdr/ptdr1_final_colour.pdf`.

[21] *CMS Technical Design Report, Vol. 1: The Trigger Systems*, 2000, `http://cmsdoc.cern.ch/cms/TDR/TRIGGER-public/trigger.html`.

[22] *ATLAS Detector and Physics Performance Technical Design Report, Vol. 1*, 1999.

[23] *ATLAS Detector and Physics Performance Technical Design Report, Vol. 2*, 1999.

[24] The *DØ* Collaboration, M. Abolins *et al.*, *Design and Implementation of the New DØ Level-1 Calorimeter Trigger*, Nucl. Instrum. Meth. A **584** (2008) 75, `arXiv:0709.3750v2`.

[25] The D0 Collaboration, V. M. Abazov *et al.*, *Evidence for production of single top quarks and first direct measurement of $|V(tb)|$*, Phys. Rev. Lett. **98** (2007) 181802, `hep-ex/0612052`.

[26] The *DØ* Collaboration, V. Abazov *et al.*, *Observation of ZZ production in ppbar collisions at sqrt(s) = 1.96 TeV*, `hep-ex/0808.0703`, Submitted to Phys. Rev. Lett., 2008.

[27] The *DØ* Collaboration, V. M. Abazov *et al.*, *$ZZ \to \ell^+\ell^-\nu\bar{\nu}$ production in $p\bar{p}$ collisions at $\sqrt{s}$= 1.96 TeV*, `hep-ex/0808.0269`, Submitted to Phys. Rev. D, 2008.

[28] The TEVNPH Working Group, *Combined CDF and DØ Upper Limits on Standard Model Higgs- Boson Production with up to 2.4 fb-1 of data*, (2008), 2008.

[29] *CMS Technical Design Report: The Muon Report*, 2000.

[30] *Tigra Tree Menu*, 2008, `http://www.softcomplex.com/products/tigra_tree_menu/`.

[31] P. Biallass, T. Hebbeker, and K. Hoepfner, *Simulation of Cosmic Muons and Comparison with Data from the Cosmic Challenge using Drift Tube Chambers*, CMS-NOTE-2007-024.

[32] A. Korn and T. Hebbeker, *Simulation Programs for the L3+Cosmics Experiment*, CERN-CMS-NOTE-2007-024.

[33] E. Barberis, D. R. Wood, P. Biallass, K. Hoepfner, and V. Drollinger, *Trigger and reconstruction studies with beam halo and cosmic muons*, CERN-CMS-NOTE-2006-012.

[34] *ATLAS Forward Detectors for Measurement of Elastic Scattering and Luminosity Determination*, Letter of Intent, 2007.

[35] F. Abe *et al.*, *Measurement of the antiproton-proton total cross section at $\sqrt{s}$ =546 and 1800 GeV*, Phys. Rev. D **50(9)** (Nov 1994) 5550.

[36] J. D. Jackson, *Classical Electrodynamics*, John Wiley and Sons, Inc., (1975).

[37] The CERN-Rome Collaboration, Phys. Lett. B **66(390)** (1977).

[38] D. Petschull, *Precision Measurement of the Fibre detector of ALFA*, Diplomarbeit, Uni. Hamburg, 2008.

Acknowledgements

Thanks are due to a number of people. Prof. Arnulf Quadt at Universiät Göttingen has offered fine guidance in the past few months, and he along with Prof. Markus Klute have agreed to be the Referent and Korreferent of this Diplomarbeit. Dr. Kevin Kröninger at Universiät Göttingen has been a great help in this process, including editing. The work described here on *DØ* and CMS was done with Prof. Darien Wood and Prof. Ela Barberis, both of Northeastern University. I have learned immensely from them. In my time at CMS I worked closely with a Dr. Oana Boeriu at Northeastern University who showed me the ropes there. While working on ALFA as a Summer Student at DESY in Hamburg, my supervisor Dr. Tobias Haas and his PhD student, Dennis Petschull, have helped me to understand the subtleties of the project. My friends and family have always been supportive of my work as I traipse around the globe which has been essential to this endeavor.

Appendix A

"Design and Implementation of the new $D\emptyset$ Level-1 Calorimeter Trigger"

Originally printed in 2008 in Nuclear Instrumentations and Methods In Physics Research Section A: accelerators, spectrometers, detectors and associated equipment [24].

Contributed to by many people at several institutes including Adam Roe.

Reprinted here in this thesis from the author's personal copy.

ELSEVIER

Nuclear Instruments and Methods in Physics Research A 584 (2008) 75–97

NUCLEAR INSTRUMENTS & METHODS IN PHYSICS RESEARCH Section A

www.elsevier.com/locate/nima

Design and implementation of the new D0 level-1 calorimeter trigger

M. Abolins[g], M. Adams[m], T. Adams[e], E. Aguilo[l,n], J. Anderson[d], L. Bagby[d], J. Ban[a], E. Barberis[h], S. Beale[n], J. Benitez[g], J. Biehl[g], M. Bowden[d], R. Brock[g], J. Bystricky[b], M. Cwiok[k], D. Calvet[b], S. Cihangir[d], D. Edmunds[g], H. Evans[f,*], C. Fantasia[h], J. Foglesong[d], J. Green[d], C. Johnson[a], R. Kehoe[j], S. Lammers[a], P. Laurens[g], P. Le Dû[b], P.S. Mangeard[f,p], J. Mitrevski[a], M. Mulhearn[a], M. Mur[b], Md. Naimuddin[c,o], J. Parsons[a], G. Pawloski[i], E. Perez[b], P. Renkel[j], A. Roe[h], W. Sippach[a], A. Stone[m,o], W. Taylor[n], R. Unalan[g], N. Varelas[m], M. Verzocchi[d], H. Weerts[g], D.R. Wood[h], L. Zhang[a], T. Zmuda[d]

[a]*Columbia University, New York, NY 10027, USA*
[b]*DAPNIA/Service de Physique des Particules, CEA, Saclay, France*
[c]*Delhi University, Delhi, India*
[d]*Fermi National Accelerator Laboratory, Batavia, IL 60510, USA*
[e]*Florida State University, Tallahassee, FL 32306, USA*
[f]*Indiana University, Bloomington, IN 47405, USA*
[g]*Michigan State University, East Lansing, MI 48824, USA*
[h]*Northeastern University, Boston, MA 02215, USA*
[i]*Rice University, Houston, TX 77005, USA*
[j]*Southern Methodist University, Dallas, TX 75275, USA*
[k]*University College Dublin, Dublin, Ireland*
[l]*University of Alberta, Edmonton, Alta., Canada*
[m]*University of Illinois at Chicago, Chicago, IL 60607, USA*
[n]*York University, Toronto, Ont., Canada*
[o]*Now at Fermi National Accelerator Laboratory, Batavia, IL 60510, USA*
[p]*Now at Université d'Aix, Centre de Physique des Particules de Marseille, Marseille, France*

Received 24 September 2007; received in revised form 1 October 2007; accepted 8 October 2007
Available online 13 October 2007

Abstract

Increasing luminosity at the Fermilab Tevatron collider has led the D0 collaboration to make improvements to its detector beyond those already in place for Run IIa, which began in March 2001. One of the cornerstones of this Run IIb upgrade is a completely redesigned level-1 calorimeter trigger system. The new system employs novel architecture and algorithms to retain high efficiency for interesting events while substantially increasing rejection of background. We describe the design and implementation of the new level-1 calorimeter trigger hardware and discuss its performance during Run IIb data taking. In addition to strengthening the physics capabilities of D0, this trigger system will provide valuable insight into the operation of analogous devices to be used at LHC experiments.

PACS: 29.40.Vj; 07.05.Hd

Keywords: Fermilab; DZero; D0; Trigger; Calorimeter

*Corresponding author. Tel.: +1 812 856 3828; fax: +1 812 855 0440.
E-mail address: hgevans@indiana.edu (H. Evans).

0168-9002/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.nima.2007.10.014

1. Introduction

During the five year period between the end of Run I in 1996 and the beginning of Run IIa in 2001, the Fermilab Tevatron accelerator implemented an ambitious upgrade program [1] in which the proton–antiproton center of mass energy was increased from 1.8 to 1.96 TeV and the instantaneous luminosity was boosted by an order of magnitude. To take advantage of the new accelerator conditions, the collider experiments, CDF and D0, also embarked on major upgrades to their detectors.

The D0 upgrade, described fully in Ref. [2], involved a complete replacement of the Run I tracking system with a new set of silicon micro-strip and scintillating fiber trackers as well as the addition of a 2 T solenoid magnet. Although the uranium and liquid argon calorimeter was left unchanged, its electronics were overhauled to match the new Tevatron bunch structure, and a series of preshower detectors was added outside of the solenoid to help measure energy of electrons, photons, and jets. Muon detection was improved with the addition of new detectors and shielding. Finally, the trigger and data acquisition (DAQ) systems were almost completely redesigned.

As originally proposed [1], approximately 20 times the integrated luminosity delivered in Run I was scheduled to be accumulated during Run II, for a total of $2\,\mathrm{fb}^{-1}$. To accomplish this goal major improvements were made to all aspects of the Tevatron, particularly in the areas of antiproton production. The bunch structure of the machine was also changed to accommodate 36×36 bunches of protons × antiprotons, with an inter-bunch spacing of 396 ns, which is an improvement over the 6×6 mode of operation in Run I. Future enhancements to 132 ns inter-bunch spacing were also foreseen, motivating a Tevatron RF structure with 159 potential bunch crossings (separated by 132 ns) during the time it takes a proton or antiproton to make a single revolution, or *turn* around the Tevatron. Of these potential crossings, only 36 contain actual proton–antiproton collisions.

Driven by ambitious physics goals of the experiments, a series of continued Tevatron improvements was also planned [3], beyond the Run II baseline, with the aim of increasing the total integrated luminosity collected to the 4–8 fb^{-1} level. To achieve this performance, instantaneous luminosities in excess of $2 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ are required. Tevatron upgrades for this period include fully commissioning the Recycler as a second stage of antiproton storage and implementing electron cooling in the Recycler. The majority of these improvements were successfully completed during a Tevatron shutdown lasting from February to May 2006, which marks the beginning of Run IIb.

The long-term effects of the Run IIb Tevatron upgrade on the D0 experiment are threefold. First, the additional integrated luminosity to be delivered to D0 during the course of Run IIb will also increase the total radiation dose accumulated by the silicon detector. Best estimates indicate that such a dose will compromise the performance of the inner layer of the detector, affecting the ability of D0 to tag *b*-quarks—a necessary ingredient in much of the experiment's physics program. Second, the increased instantaneous luminosity stresses the trigger system, decreasing the ability to reject background while maintaining high efficiency for signal events. And finally, the plan of having real bunch crossings separated by 132 ns, although not realized in the final Run IIb configuration, would have created problems matching calorimeter signals with their correct bunch crossing in the Run IIa calorimeter trigger system.

The first of the effects mentioned above led D0 to propose the addition of a radiation-hard inner silicon layer (*Layer*-0) to the tracking system [4]. The second and third effects required changes to various aspects of the trigger system [5]. These additions and modifications, collectively referred to as the *D0 Run IIb Upgrade*, were designed and implemented between 2002 and 2006 and were installed in the experiment during the 2006 Tevatron shutdown.

In the following we describe the Level-1 calorimeter trigger system (L1Cal) designed for operation during Run IIb. Section 2 contains a brief description of the Run IIa D0 calorimeter and the three-level trigger system. Section 3 discusses the motivation for replacing the L1Cal trigger, which was used in Run I and Run IIa. Algorithms used in the new system and their simulation are described in Sections 4 and 5, while the hardware designed to implement these algorithms is detailed in Sections 6–9. Mechanisms for online control and monitoring of the new L1Cal are outlined in Sections 10 and 11. This article then concludes with a discussion of early calibration and performance results in Sections 12 and 13, with a summary presented in Section 14.

2. Existing framework

2.1. The D0 calorimeter

The basis of the Run IIb L1Cal trigger is the D0 calorimeter, described in more detail in Refs. [2,6]. This detector, shown schematically in Fig. 1, consists of three sampling calorimeters (a barrel and two endcaps), in three separate cryostats, using liquid argon as the active medium and depleted uranium, uranium–niobium alloy, copper or stainless steel as the absorber. It also includes detectors in the intercryostat region (ICR), where the barrel and endcaps meet, consisting of scintillating tiles, as well as instrumented regions of the liquid argon without absorbers. The calorimeter has three longitudinal sections—electromagnetic (EM), fine hadronic (FH) and coarse hadronic (CH)—each themselves divided into several layers. It is segmented laterally into cells of size $\Delta\eta \times \Delta\phi \sim 0.1 \times 0.1$ [7] arranged in pseudo-projective towers (except for one layer in the EM section, which has $\Delta\eta \times \Delta\phi \sim 0.05 \times 0.05$). The calorimeter system provides coverage out to $|\eta| \sim 4$.

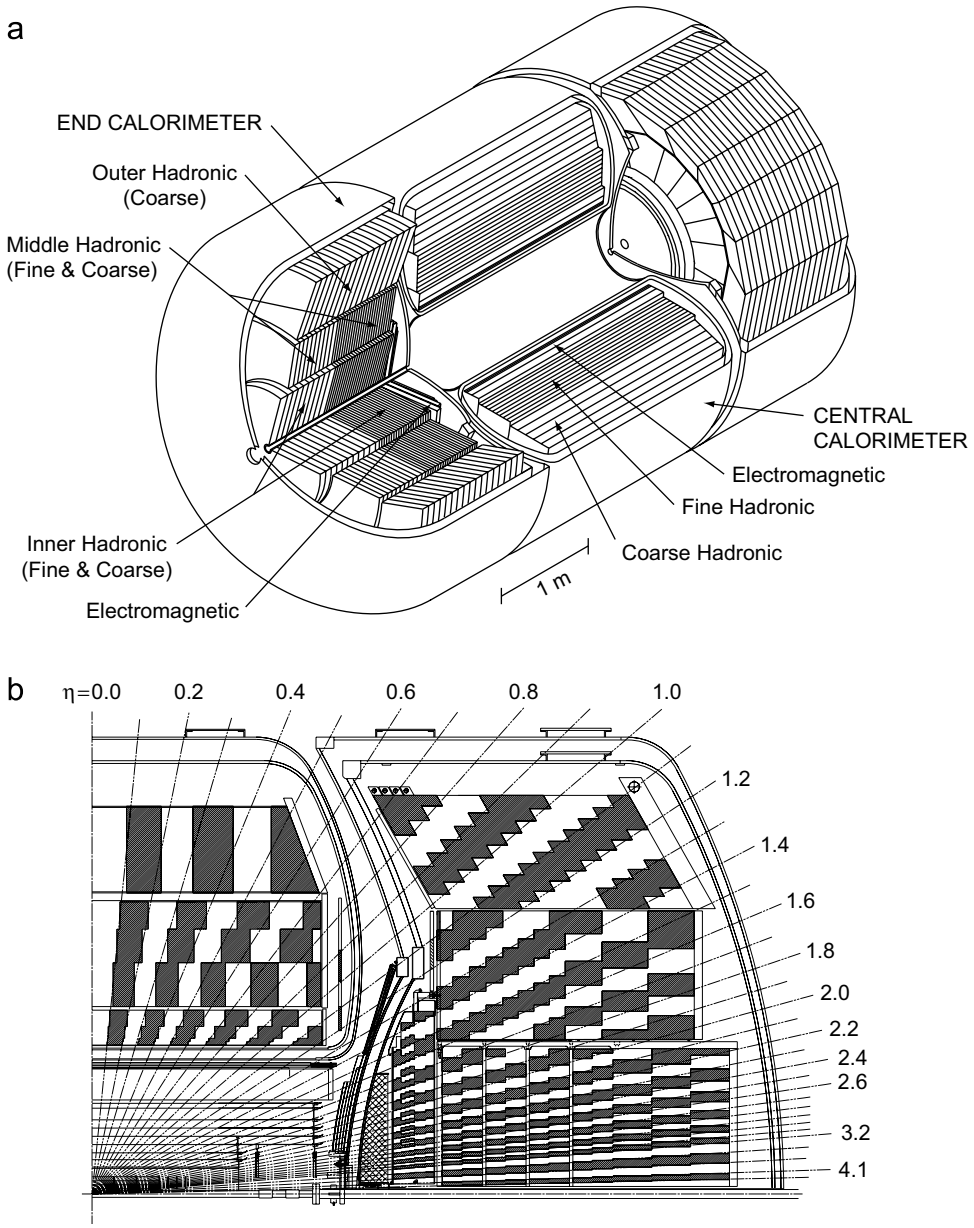

Fig. 1. An isometric view of the central and two endcap calorimeters (a) and a schematic view of a portion of the calorimeter showing the transverse and longitudinal segmentation pattern (b).

Charge collected in the calorimeter is transmitted via impedance-matched coaxial cables of ~10 m length to charge sensitive preamplifiers located on the detector. The charge integrated output of these preamplifiers has a rise time of ~450 ns, corresponding to the electron drift time across a liquid-argon gap, and a fall time of ~15 μs. The single-ended preamplifier signals are sent over ~25 m of twisted pair cable to Baseline Subtractor (BLS) cards.

On the 1152 BLS cards, the preamplifier signals are split into two paths: the *precision readout* and the *trigger sum pickoff*. Precision readout path signals for each calorimeter cell are shaped, baseline subtracted, and stored in a set of switched capacitor arrays awaiting Level-1 and Level-2 trigger decisions. Signals on the trigger sum pickoff path, shown in Fig. 2 are shaped to a triangular pulse with a fast rise and a linear fall over 400 ns. They are then passed to analog summers that add signals in different cells, weighted appropriately for the sampling fraction and capacitance of each cell to form EM and HD trigger towers (TTs). EM TTs contain all cells (typically 28) in $0.2 \times 0.2\,\Delta\eta \times \Delta\phi$ regions of the EM section of the calorimeter, while HD TTs use (typically 12) cells in the FH section of the calorimeter to form 0.2×0.2 regions. This granularity leads to 1280 EM and 1280 HD TTs forming a 40×32 grid in $\eta \times \phi$ space, which covers the entire azimuthal region for $|\eta| < 4.0$. Due mainly to overlapping collisions, which complicate the forward environment, however, only the region $|\eta| < 3.2$ is used for triggering.

The EM and HD TT signals are transmitted differentially to the Level-1 calorimeter trigger electronics on two

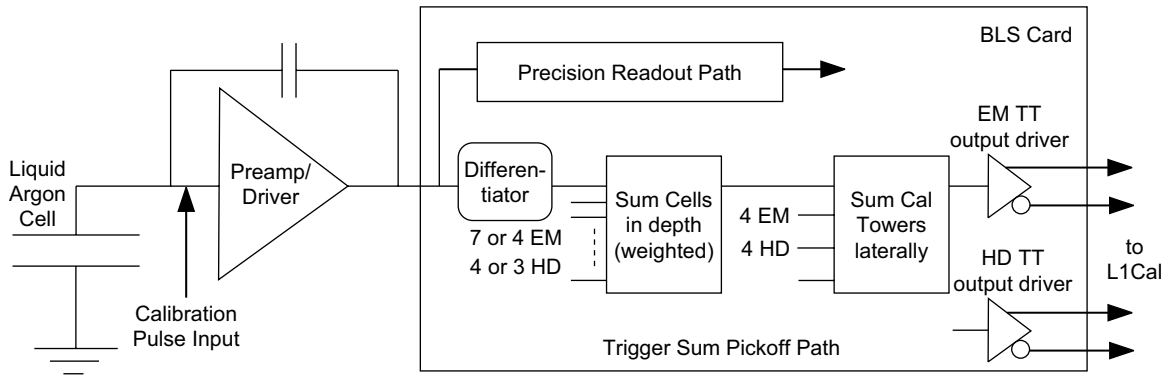

Fig. 2. The calorimeter readout chain, including a preamplifier and Baseline Subtractor (BLS) card with emphasis on the elements of the trigger sum pickoff path.

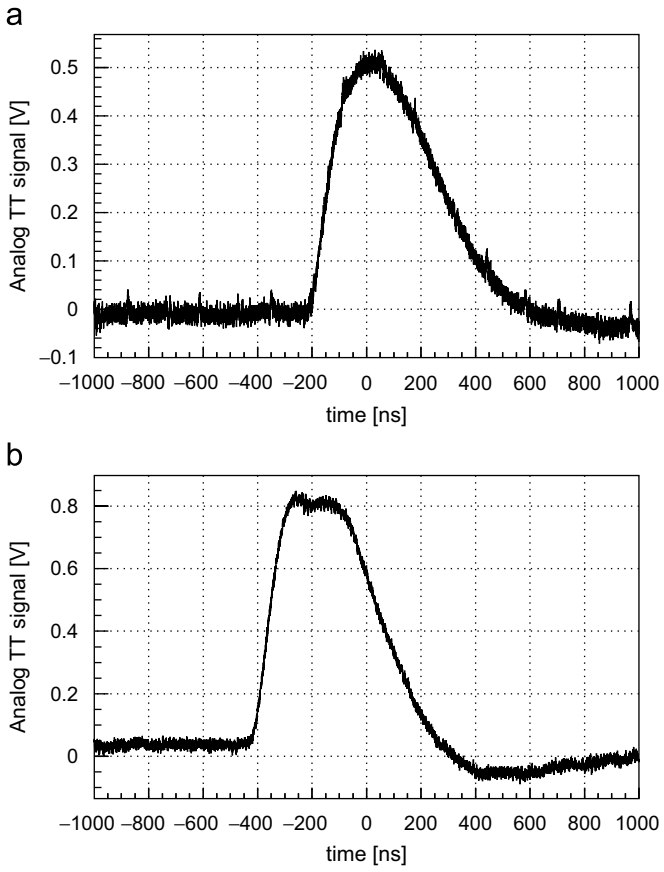

Fig. 3. Typical EM (a) and HD (b) analog signals. In both plots, the non-inverted minus the inverted differential signals are shown.

separate miniature coaxial cables. Although the signal characteristics of these cables are quite good, some degradation occurs in the transmission, yielding L1Cal input signals with a rise time of ~250 ns and a total duration of up to 700 ns. Typical EM and HD TT signals are shown in Fig. 3.

2.2. Overview of the D0 trigger system

The D0 experiment uses a three-level trigger system, shown schematically in Fig. 4 and described in more detail in Ref. [2], to select interesting events from the 1.7 MHz of bunch crossings seen in the detector. Individual elements contributing to the Level-1 (L1) and Level-2 (L2) systems, as used in Run IIb, are shown in Fig. 5.

The L1 trigger system, implemented in custom hardware, examines data from the detector for every bunch crossing. It consists of separate elements for calorimeter (L1Cal), scintillating fiber tracking (L1CTT), muon (L1Muon), and forward proton (L1FPD) data. New for Run IIb is an element that matches tracks and calorimeter clusters at L1 (L1CalTrk), which is functionally similar to L1Muon. Each L1 trigger element sends its decisions on a set of criteria (for example, the presence of two jets with transverse energy above a threshold) to the Trigger Framework (TFW). The TFW uses these decisions, referred to as the *and/or* terms to decide whether the event should be accepted for further processing or rejected. Because of the depth of data pipelines in the detector's front end electronics, L1 decisions from each of the trigger elements must arrive at the TFW within 3.7 μs of the bunch crossing producing their data. This pipeline depth was increased from its Run IIa value of 3.3 μs in order to accommodate the extra latency induced by the L1CalTrk system. After an L1 accept, data are transferred off of the pipelines, inducing deadtime in the system. The maximum allowable L1 accept rate, generally around 2 kHz, is set by the desire to limit this deadtime to the 5% level.

The L2 system receives data from the detector and from the L1 trigger elements on each L1 accept. It consists of detector-specific pre-processor engines for calorimeter (L2Cal); preshower (L2PS); scintillating fiber (L2CTT) and silicon (L2STT) tracking; and muon (L2Muon) data. Processed data from each of these elements are transmitted to a global processor (L2Global) that selects events based on detector-wide correlations between its input elements. The L2 trigger operates at a maximum input rate of 2 kHz and provides L2 accepts at a rate of up to 1 kHz.

The final stage in the D0 trigger system, Level-3 (L3), consists of a farm of PCs that have access to the full detector readout on L2 accepts. These processors run a simplified version of the offline event reconstruction and make decisions based on physics objects and the

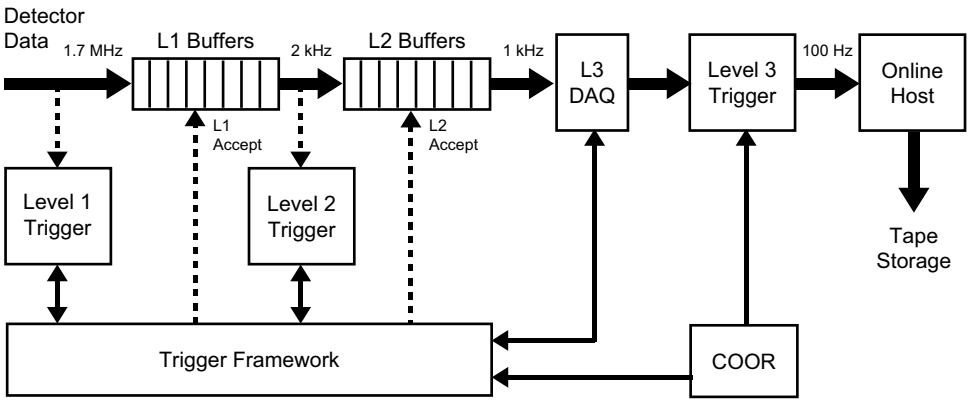

Fig. 4. An overview of the D0 trigger and data acquisition system.

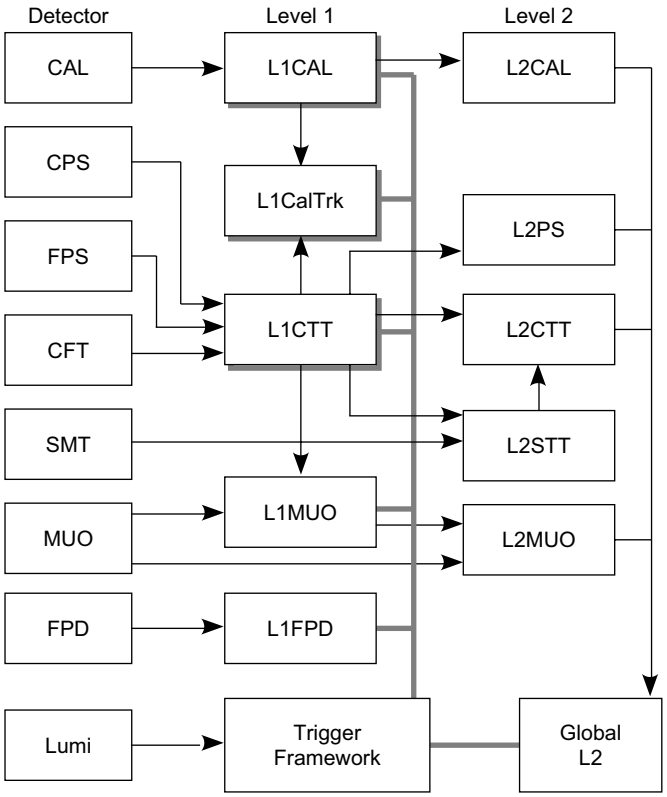

Fig. 5. A block diagram of the D0 L1 and L2 trigger systems.

relationships between them. L3 accepts events for permanent storage at a rate of up to 150 Hz (typically, 100 Hz).

The configuration of the entire D0 trigger system is accomplished under the direction of the central coordination program (COOR), which is also used for detector configuration and run control.

3. Motivation for the L1Cal upgrade

By the time of the start of Run IIa in 2001, there was already a tentative plan in place for an extension to the run with accompanying upgrades to the accelerator complex [3], leading to an additional 2–6 fb^{-1} of integrated luminosity beyond the original goal of 2 fb^{-1}. This large increase in statistical power opens new possibilities for physics at the Tevatron such as greater precision in critical measurements like the top quark mass and W boson mass, the possibility of detecting or excluding very rare Standard Model processes (including production of the Higgs boson), and greater sensitivity for beyond the Standard Model processes like supersymmetry.

At a hadron collider like the Tevatron, however, only a small fraction of the collisions can be recorded, and it is the trigger that dictates what physics processes can be studied and what is left unexplored. The trigger for the D0 experiment in Run IIa had been designed for a maximum luminosity of $1 \times 10^{32}\,cm^{-2}\,s^{-1}$, while the peak luminosities in Run IIb are expected to go as high as $3 \times 10^{32}\,cm^{-2}\,s^{-1}$. In the three-level trigger system employed by D0, only the L3 trigger can be modified to increase its throughput; the maximum output rates at L1 and L2 are imposed by fundamental features of the subdetector electronics. Thus, fitting L1 and L2 triggers into the bandwidth limitations of the system can only be accomplished by increasing their rejection power. While an increase in the transverse energy thresholds at L1 would have been a simple way to achieve higher rejection, such a threshold increase would be too costly in efficiency for the physics processes of interest. For example, raising the thresholds on the two jet triggers used in the search for $p\bar{p} \to ZH \to \nu\bar{\nu}b\bar{b}$ events to yield acceptable rates would have resulted in a decrease in signal efficiency of more than 20%. The D0 Run IIb Trigger Upgrade [5] was designed to achieve the necessary rate reduction through greater selectivity, particularly at the level of individual L1 trigger elements.

The L1Cal trigger used in Run I and in Run IIa [8] was based on counting individual TTs above thresholds in transverse energy (E_T). Because the energy from electrons/photons and especially from jets tends to spread over multiple TTs, the thresholds on tower E_T had to be set low relative to the desired electron or jet E_T. For example, an EM TT threshold of 5 GeV is fully efficient only for electrons with E_T greater than about 10 GeV, and a

5 GeV threshold for EM + HD tower E_T only becomes 90% efficient for jet transverse energies above 50 GeV.

The primary strategy of the Run IIb upgrade of L1Cal is therefore to improve the sharpness of the thresholds for electrons, photons and jets by forming clusters of TTs and comparing the transverse energies of these clusters, rather than individual tower E_T's, to thresholds.

The design of clustering using sliding windows (SW) (see Section 4) in Field Programmable Gate Arrays (FPGAs) meets the requirements of this strategy, and also opens new possibilities for L1Cal, including sophisticated use of shower shape and isolation; algorithms to find hadronic decays of tau leptons through their characteristic transverse profile; and requirements on the topology of the electrons, jets, taus, and missing transverse energy in an event.

4. Algorithms for the Run IIb L1Cal

Clustering of individual TTs into EM and Jet objects is accomplished in the Run IIb L1Cal by the use of a sliding windows algorithm (SWA). This algorithm performs a highly parallel cluster search in which groups of contiguous TTs are compared to nearby groups to determine the location of local maxima (LM) in E_T deposition. Variants of the SWA have been studied extensively at different HEP experiments [9], and have been found to be highly efficient at triggering on EM and Jet objects, while not having the latency drawbacks of iterative clustering algorithms. For a full discussion of the merits of the sliding windows algorithm, see Ref. [10].

The implementation of the SWA in the D0 calorimeter trigger occurs in three phases. In the first phase, the digitized transverse energies of several TTs are summed into Trigger Tower Clusters (TTCL). These TTCL sums, based on the size of the EM or Jet SW, are constructed for every point in TT space, and are indexed by the η, ϕ coordinate of one of the contributing TTs, with different conventions being used for different algorithms (see Sections 4.1 and 4.2). This process, which yields a grid of TTCLs that share energy with their close neighbors, is shown in the first and second panels of Fig. 6.

In the second phase, the TTCL are analyzed to determine locations of large energy deposits called LM. These LM are chosen based on a comparison of the magnitude of the E_T of a TTCL with that of its adjacent TTCL. Multiple counting of jet or EM objects is avoided by requiring a spatial separation between adjacent LM as illustrated in the third panel of Fig. 6.

In the third phase, additional information is added to define an output object. In the case of Jet objects, shown in the fourth panel of Fig. 6, energy of surrounding TTs is added to the TTCL energy to give the total Jet object energy. EM and Tau objects are also refined in this phase using isolation information (see Sections 4.2 and 4.3).

Results for the entire calorimeter can be obtained very quickly using this type of algorithm by performing the LM finding and object refinement phases of the algorithm in parallel for each TTCL.

4.1. Jets

Jets at the Tevatron have lateral sizes of order one unit in η, ϕ space and deposit energy in both the EM and hadronic portions of the calorimeter. Therefore, Jet objects in the D0 L1Cal are defined using the sum of the EM and HD energies as the input to the TTCL-sums. The TTCL are 2×2 in TT units, corresponding to a region 0.4×0.4 in $\eta \times \phi$ space on the inner face of the calorimeter. LM are required to be separated by one TT and the final energy sums are 4×4 in TT space, corresponding to a region 0.8×0.8 in $\eta \times \phi$ space.

The values of these clustering parameters were determined by optimizing Jet object energy and position resolution.

4.2. EM objects

EM objects (electrons or photons) have lateral shower profiles that are much smaller than the TT size, and tend not to deposit energy in the hadronic calorimeter. For this reason, EM TTs are input directly to the local maximum finding algorithm (the TTCL size is 1×1 in TT units). Because electrons or photons may deposit energy close to the boundary between TTs, the final EM object, as shown in Fig. 7, is composed of two adjacent TTs, oriented horizontally (containing two TTs in η) or vertically

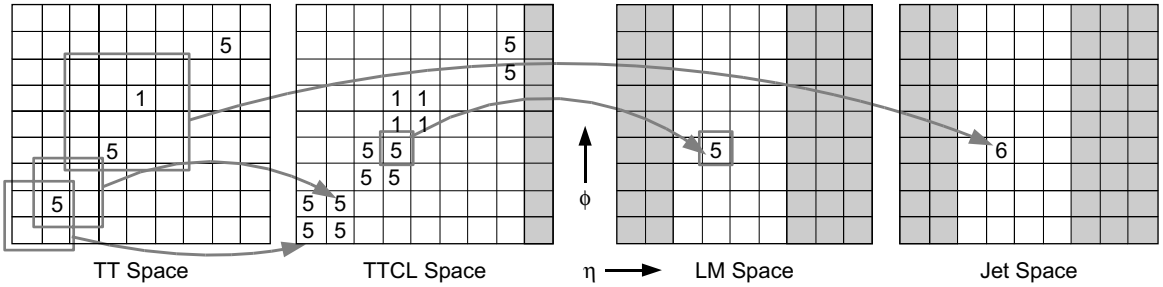

Fig. 6. The stages of algorithm flow for the sliding windows algorithm. In this example, which corresponds to the Run IIb Jet algorithm, a 2×2 TT TTCL is used, indexed by the position of its smallest η, ϕ TT. Baseline subtracted TT energies are indicated by numbers, and local maxima are required to be separated by at least 1 TT. Jet objects are defined as the E_T sum of the 4×4 TTs centered on the TTCL. Light gray regions in the diagrams indicate areas for which the object in question cannot be constructed because of boundary effects.

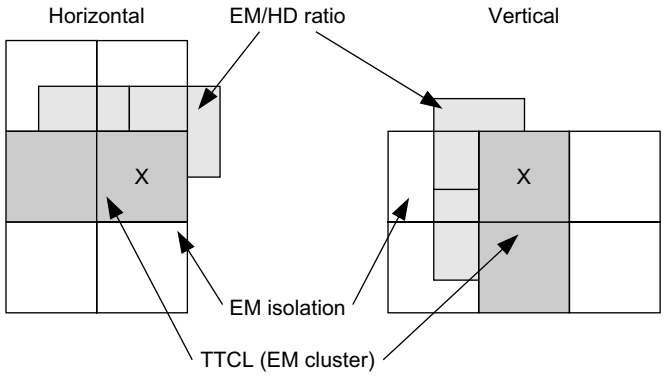

Fig. 7. Definition of EM trigger objects.

(containing two TTs in ϕ), where the first tower is the LM and the second is the neighboring tower with the highest E_T. Cuts can also be applied on the electromagnetic fraction (EM/HD) and isolation of the candidate EM object. The former is determined using the ratio of the EM TT energies making up the EM object and the corresponding two HD TTs directly behind it. The isolation region is composed of the four EM TTs adjacent to the EM object; cuts are placed on the ratio of the total E_T in the EM-isolation region and the EM object E_T. In both cases, the ratio cut value is constrained to be a power of two in order to reduce latency in the divide operation as implemented in digital logic.

This algorithm was chosen based on an optimization of the efficiency for triggering on electrons from $W \to e\nu$ and $J/\psi \to e^+e^-$ decays.

4.3. Taus

Tau leptons that decay hadronically look similar to jets, but have narrow, energetic cores. This allows extra efficiency for processes containing taus to be obtained by relaxing E_T threshold requirements on these objects (compared to Jet thresholds) but additionally requiring that only small amounts of energy surround the Tau candidate. The Run IIb L1Cal uses the results of the Jet algorithm as a basis for Tau objects but also calculates the ratio of the 2×2 TT TTCL to the 4×4 total Jet object E_T. Large values of this isolation ratio, as well as large E_T, are required in the definition of a Tau object. Because of data transfer constraints in the system, however, the E_T associated with the Tau object is taken from the Jet object closest in ϕ to the LM passing the Tau isolation cut.

4.4. Sum E_T and missing E_T

Scalar and vector E_T sums are computed for the EM + HD TTs. In constructing these sums, the η range of the contributing TTs can be restricted and an E_T threshold can be applied to the TTs entering the sums to avoid noise contamination.

4.5. Use of the intercryostat detectors

Object and sum energies in the Run IIb L1Cal can be configured to include energies seen in the ICR. Because of complicated calibrations and relatively poor resolution in these regions, however, this option is currently not in use.

4.6. Topological triggers

Because of its increased processing capabilities, the Run IIb L1Cal can require spatial correlations between some of its objects to create *topological* trigger terms. These triggers can be used to distinguish signals that have numbers of objects identical to those observed in large backgrounds but whose event topologies are much rarer. An example of such a topology occurs in associated Higgs production in which the decay $ZH \to \nu\bar{\nu}b\bar{b}$ yields two jets acoplanar with respect to the beam axis, and large missing transverse energy. Since the only visible energy in such an event is reflected in the jets, it is difficult to distinguish this process from the overwhelming dijet QCD background. The Run IIb L1Cal contains a trigger that specifically selects dijet events in which the two jets are required to be acolinear in the transverse plane. Other topological triggers that have been studied are back-to-back (in the transverse plane) EM object triggers to select events containing J/ψ mesons, and triggers that select events with jet-free regions of the calorimeter containing small energy deposits, for triggering on mono-jet events.

5. Simulation and predictions

Two independent methods of simulating the performance of the L1Cal algorithms have been developed: a module included in the overall D0 trigger simulation for use with Monte Carlo or real data events (*TrigSim*), and a tool developed to estimate and extrapolate trigger rates based on real data accumulated during special low-bias runs (*Trigger Rate Tool*). Both of these methods were used to develop a new Run IIb trigger list that will collect data efficiently up to the highest luminosities foreseen in Run IIb.

5.1. Monte Carlo based simulation

A C++ simulation of the Run IIb L1Cal trigger has been developed, as part of the full D0 trigger simulation (TrigSim)—a single executable program that provides a standard framework for including code that simulates each individual D0 trigger element. This framework allows the specification of the format of the data transferred between trigger elements, the simulation of the time ordering of the trigger levels and the simulation of the data transfers. The L1Cal portion of TrigSim emulates all aspects of the L1Cal algorithms. It can be run either as part of the full D0 trigger simulation or in a stand-alone mode on both Monte

Carlo simulated data and real D0 data, allowing checks on hardware performance, as well as estimates of signal efficiencies and background rates, as part of algorithm optimization.

5.2. Trigger Rate Tool

A great benefit in designing and testing the algorithms for L1Cal in Run IIb was the availability of real collision data from Run IIa. In every event recorded in Run IIa, the transverse energy of every trigger tower was saved. These energies serve as input to a stand-alone emulation of the Run IIb algorithms (the Trigger Rate Tool) used to estimate rates and object-level efficiencies from actual data. Special data runs were taken with low tower thresholds, and the Trigger Rate Tool was applied to these runs to predict the rates for any list of emulated triggers with a proper treatment of the correlations among triggers in the list. The Trigger Rate Tool was also used to compare the Run IIa and Run IIb trigger lists and to extrapolate rates from the relatively low luminosities existing when the Run IIa data were taken to the much higher values anticipated in Run IIb. Predictions based on results obtained from this tool indicated that the upgraded trigger would reduce the overall Level 1 rates by about a factor of two while maintaining equal or improved efficiency for signal processes at the highest instantaneous luminosities foreseen in Run IIb.

5.3. Predictions

Predictions of the impact of the new L1Cal SWAs on the L1 trigger rates and efficiencies were obtained using simulations of dijet events and various physics processes of interest in Run IIb. After trying different configurations that gave the same rate as those experienced during Run IIa, the most efficient configurations were chosen and put in an overall trigger list to check the total rate.

Fig. 8 shows the predicted rates at a luminosity of $2 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, estimated using the Trigger Rate Tool, for trigger lists based on Run IIa algorithms (v14) and their Run IIb equivalents (v15). Both trigger lists were designed to give similar efficiencies for physics objects of interest in Run IIb. However, the Run IIb trigger list yields a rate approximately a factor of two smaller than that achievable using Run IIa algorithms.

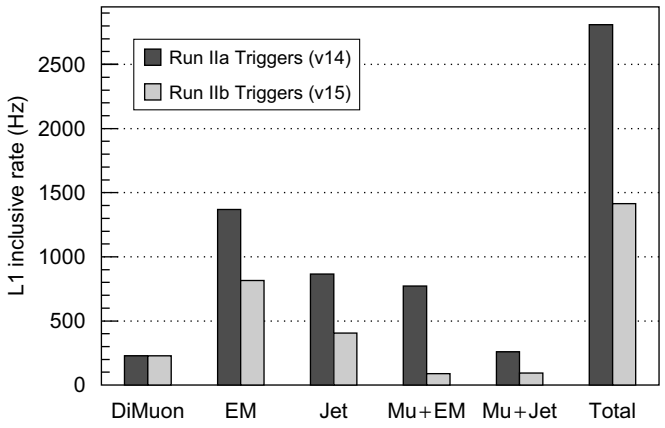

Fig. 8. Predicted rates for Run IIa (v14) and Run IIb (v15) trigger lists, extrapolated to a luminosity of $2 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ from trigger-unbiased data collected at lower luminosity.

6. Hardware overview

The algorithms described previously are implemented in several custom electronics boards designed for the new L1Cal. An overview of the main hardware elements of the Run IIb L1Cal system is given in Fig. 9. Broadly, these elements are divided into three groups.

(1) The *Analog and Digital Filter* (*ADF*) *System*, containing those elements that receive and digitize analog TT signals from the BLS cards, and perform TT-based signal processing.
(2) The *Trigger Algorithm Broads* (*TABs*)/*Global Algorithm Board* (*GAB*) System, where algorithms are run on the digitized TT signals to produce trigger terms.
(3) The *Readout System*, which inserts L1Cal information into the D0 data path for permanent storage.

The L1Cal also communicates with other elements of the D0 trigger and DAQ system, including the following:

- The TFW, which delivers trigger decisions and synchronizes the entire D0 DAQ. From the L1Cal point of view, the TFW sends global timing and control signals (see Table 1) to the system over Serial Command Links (SCL) and receives the L1Cal and/or terms.
- The L1Cal Trigger Control Computer (L1Cal TCC), which configures and monitors the system.
- The Level-1 Cal-Track Match trigger system (L1CalTrk), another L1 trigger system that performs azimuthal matching between L1CTT tracks and L1Cal EM and Jet objects.

Within the L1Cal, the ADF system consists of the *Transition System*, the ADF cards, and the Serial Command Link Distributor (SCLD). The Transition System, consisting of *Patch Panels*, Patch Panel Cards (PPCs), ADF Transition Cards (ATC), and connecting cables, adapts the incoming BLS signal cables to the higher density required by the ADFs. These ADF cards, which reside in four 6U VME-64x crates [11], filter, digitize and process individual TT signals, forming the building blocks of all further algorithms. They receive timing and control signals from the SCL via an SCLD card.

Trigger algorithms are implemented in the L1Cal in two sets of cards: the TABs and the GAB, which are housed in a single 9U crate with a custom backplane. The TABs identify EM, Jet and Tau objects in specific regions of the calorimeter using the algorithms described in Section 4 and

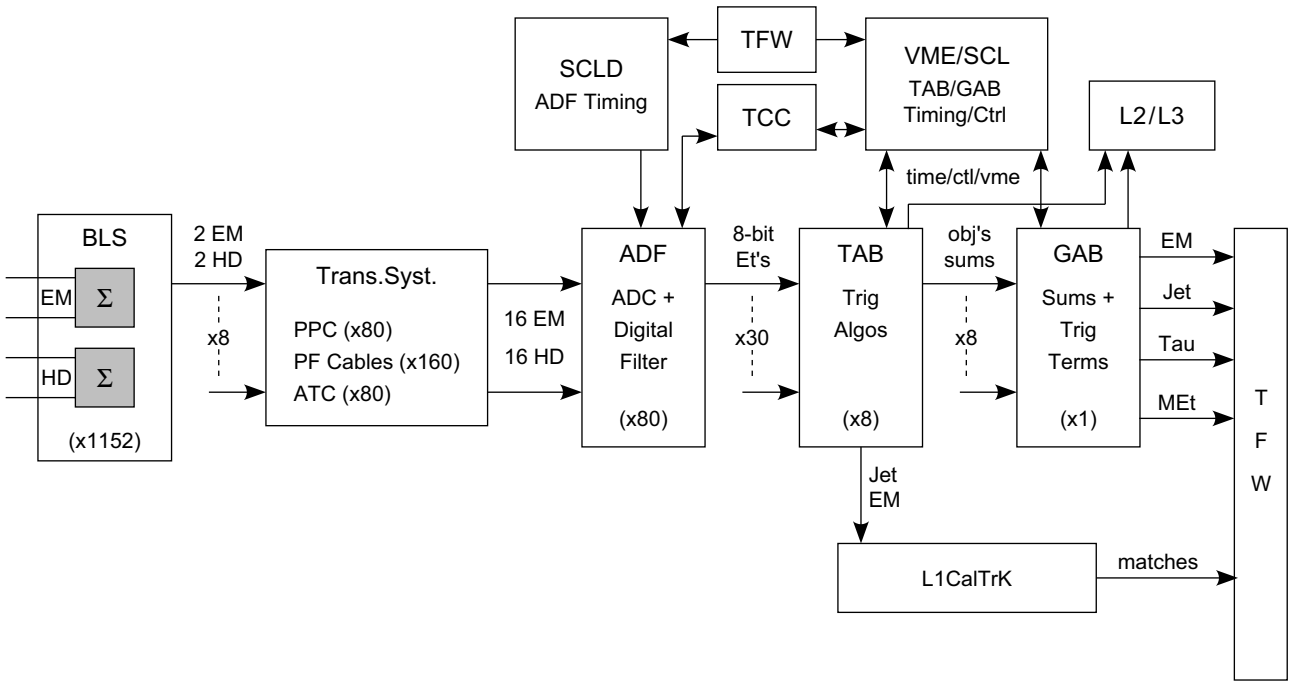

Fig. 9. A block diagram of the main hardware elements of the Run IIb L1Cal system and their interconnections.

Table 1
Timing and control signals used in the L1Cal system

SCL	ADF	TAB/GAB	Description
INIT	–	Yes	Initialize the system
CLK7	Yes	Yes	132 ns Tevatron RF clock
TURN	Yes	Yes	Marks the first crossing of an accelerator turn
REALBX	Yes	–	Flags clock periods containing real beam crossings
BX_NO	–	Yes	Counts the 159 bunch crossings in a turn
L1ACCEPT	Yes	Yes	Indicates that an L1 Accept has been issued by the TFW
MONITOR	Yes	–	Initiates collection of ADF monitoring data
L1ERROR	–	Yes	A TAB/GAB error condition transmitted to the SCL hub
L1BUSY	–	Yes	Asserted by the TABs/GAB until an observed error is cleared
–	ADF_MON	–	Allows TCC to freeze ADF circular buffers
–	ADF_TRIG	–	Allows TCC to fake a MONITOR signal on the next L1 Accept
–	–	TAB_RUN	TAB/GAB data path synchronization signal
–	–	TAB_TRIG	Pulse to force writing to TAB/GAB diagnostic memories
–	–	TAB_FRM	Used for synchronization of TAB/GAB VME data under VME/SCL control
–	–	TAB_ADDR	Internal address for TAB/GAB VME read/write operations
–	–	TAB_DATA	Data for TAB/GAB VME read/write operations

Included are D0 global timing and control signals (SCL) used by the ADFs and the TAB/GAB system, as well as intra-system communication and synchronization flags described later in the text.

also calculate partial global energy sums. The GAB uses these objects and energy sums to calculate and/or terms, which the TFW uses to make trigger decisions. Finally, the *VME/SCL card*, located in the L1Cal Control Crate, distributes timing and control signals to the TABs and GAB and provides a communication path for their readout.

The architecture of the L1Cal system and the number of custom elements required, summarized in Table 2, is driven by the large amount of overlapping data required by the SWA. In total, more than 700 Gbits of data per second are transmitted within the system. Of this, each local maximum calculation requires 4.4 Gbits/s from 72 separate TTs. The most cost-effective solution to this problem, which still results in acceptable trigger decision latency, is to deal with all data as serial bit streams. Thus, all intra-system data transmission is done bit serially using the Low Voltage Differential Signal (LVDS) protocol and nearly all algorithm arithmetic is performed bit serially as well, at clock speeds such that all bits of a data word are examined in the

132 ns Tevatron bunch crossing interval. Examples of a bit-serial adder and comparator are shown in Fig. 10. The only exception to this bit-serial arithmetic rule is in the calculation of Tau object isolation, which requires a true divide operation (see Section 4) and thus introduces an extra 132 ns of latency to the trigger term calculation. Even with this extra latency, the L1Cal results arrive at the TFW well within the global L1 decision time budget.

Table 2
A summary of the main custom electronics elements of the L1Cal system

Board	Input TT region	Output TT region	Total number
PPC	4×4	4×4	80
ATC	4×4	4×4	80
ADF	4×4	4×4	80
SCLD	All	All	1
TAB	40×12	31×4	8
GAB	All	All	1
VME/SCL	All	All	1

For each board, the TT region (in $\eta \times \phi$) that the board receives as input and sends on as output is given as well as the total number of each board type required in the system.

7. The ADF system

7.1. Transition System

Trigger pick-off signals from the BLS cards of the EM and HD calorimeters are transmitted to the L1Cal trigger system, located in the Movable Counting House (MCH), through 40–50 m long coaxial ribbon cables. Four adjacent coaxial cables in a ribbon carry the differential signals from the EM and HD components of a single TT. Since there are 1280 BLS trigger cables distributed among 10 racks of the original L1Cal trigger electronics, the L1Cal upgrade was constrained to reuse these cables. However, because the ADF input signal density is much larger than that in the old system (only four crates are used to house the ADFs as opposed to 10 racks for the old system's electronics) the cables could not be plugged directly into

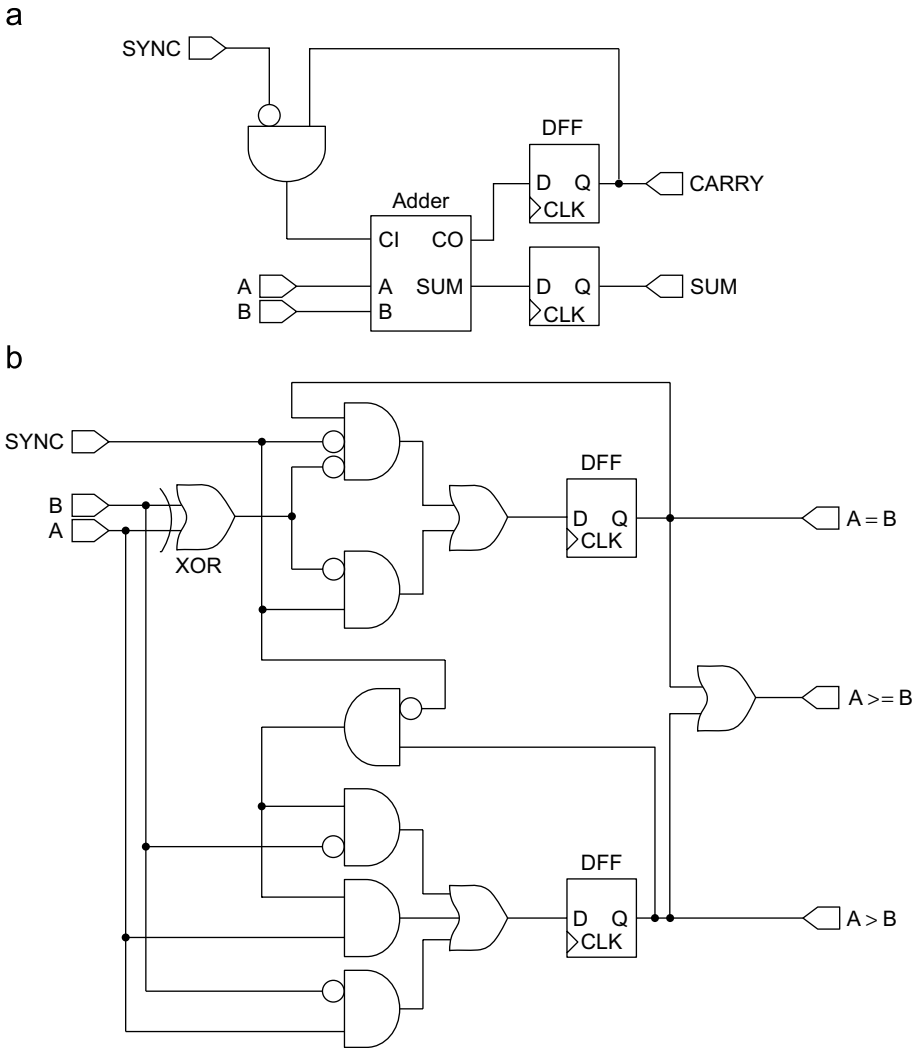

Fig. 10. Logic diagrams for a bit-serial adder (a) and a bit-serial comparator (b).

the upgraded L1Cal trigger electronics; a Transition System was needed.

The Transition System is composed of passive electronics cards and cables that route signals from the BLS trigger cables to the backplane of the ADF crates (see Section 7.2). It was designed to allow the trigger cables to remain within the same Run I/IIa rack locations. It consists of the following elements.

- Patch Panels and PPCs: A PPC receives the input signals from 16 BLS trigger cables and transmits the output through a pair of Pleated Foil Cables. A PPC also contains four connectors which allow the monitoring of the signals. Eight PPCs are mounted two to a Patch Panel in each of the 10 racks originally used for Run I/IIa L1Cal electronics.
- Pleated Foil Cables: Three meter long Pleated Foil Shielded Cables (PFC), made by the 3M corporation [12], are used to transfer the analog TT output signals from the PPC to the ADF cards via the ADF Transition Card. There are two PFCs for each PPC for a total of 160 cables. The unbalanced characteristic impedance specification of the PFC is 72 Ω, which provides a good impedance match to the BLS trigger cables.
- ADF Transition Card (ATC): The ATCs are passive cards connected to the ADF crate backplane. These cards receive the analog TT signals from two PFCs and transmit them to the ADF card. There are 80 ATCs that correspond to the 80 ADF cards. Each ATC also transmits the three output LVDS cables of an ADF card to the TAB crate—a total of 240 LVDS cables.

7.2. ADF cards

The ADF cards are responsible for sending the best estimate of the transverse energy (E_T) in the EM and HD sections of each of the 1280 TTs to the eight TAB cards for each Tevatron beam crossing. The calculation of these E_T values by the 80 ADF cards is based upon the 2560 analog trigger signals that the ADF cards receive from the BLS cards, and upon the timing and control signals that are distributed throughout the D0 DAQ system by the SCL. The ADF cards themselves are 6U × 160 mm, 12-layer boards designed to connect to a VME64x backplane using P0, P1 and P2 connectors. The ADF system is set up and monitored, over VME, by a TCC, described in Section 10.

7.3. Signal processing in the ADFs

Each ADF card, as shown schematically in Fig. 11, uses 32 analog trigger signals corresponding to the EM and HD components of a 4 × 4 array of TTs. Each differential, AC coupled analog trigger signal is received by a passive circuit that terminates and compensates for some of the characteristics of the long cable that brought the signal out of the collision hall. Following this passive circuit the active part of the analog receiver circuit rejects common mode noise on the differential trigger signal, provides filtering to select the frequency range of the signal caused by a real Tevatron energy deposit in the calorimeter, and provides additional scaling and a level shift to match the subsequent ADC circuit.

The analog level shift in the trigger signal receiver circuit is controlled, separately for each of the 32 channels on an ADF card, by a 12 bit pedestal control DAC, which can swing the output of the ADC that follows it from slightly below zero to approximately the middle of its full range. This DAC is used both to set the pedestal of the signal coming out of the ADC that follows the receiver circuit and as an independent way to test the full signal path on the ADF card. During normal operation, we set the pedestal at the ADC output to 50 counts which is a little less than 5% of its full scale range. This offset allows us to accommodate negative fluctuations in the response of the BLS circuit to a zero-energy signal.

The 10 bit sampling ADCs [13] that follow the receiver circuit make conversions every 33 ns—four times faster than the Tevatron BX period of 132 ns. This conversion rate is used to reduce the latency going through the pipeline ADCs and to provide the raw data necessary to associate the rather slow rise-time trigger signals (250 ns typical rise-time) with the correct Tevatron beam crossing. Although associating energy deposits in the calorimeter with the correct beam crossing is not currently an issue since actual proton–antiproton collisions only occur every 396 ns, rather than every 132 ns as originally planned, the over-sampling feature has been retained for the flexibility it provides in digital filtering.

On each ADF card the 10 bit outputs from the 32 ADCs flow into a pair of FPGAs [14], called the *Data Path FPGAs*, where the bulk of the signal processing takes place. This signal processing task, shown schematically in Fig. 12, is split over two FPGAs with each FPGA handling all of the steps in the signal processing for 16 channels. Two FPGAs were used because it simplified the circuit board layout and provided an economical way to obtain the required number of I/O pins.

The first step in the signal processing is to align in time all of the 2560 trigger signals. The peak of the trigger signals from a given beam crossing arrive at the L1Cal at different times because of different cable lengths and different channel capacitances. These signals are made isochronous using variable length shift registers that can be set individually for each channel by the TCC. Once the trigger signals have been aligned in time, they are sent to both the Raw ADC Data Circular Buffers where monitoring data are recorded and to the input of the Digital Filter stage.

The Raw ADC Data Circular Buffers are typically set up to record all 636 of the ADC samples registered in a full turn of the accelerator. This writing operation can be stopped by a signal from the TCC, when an L1 Accept flagged with a special *Collect Status* flag is received by the

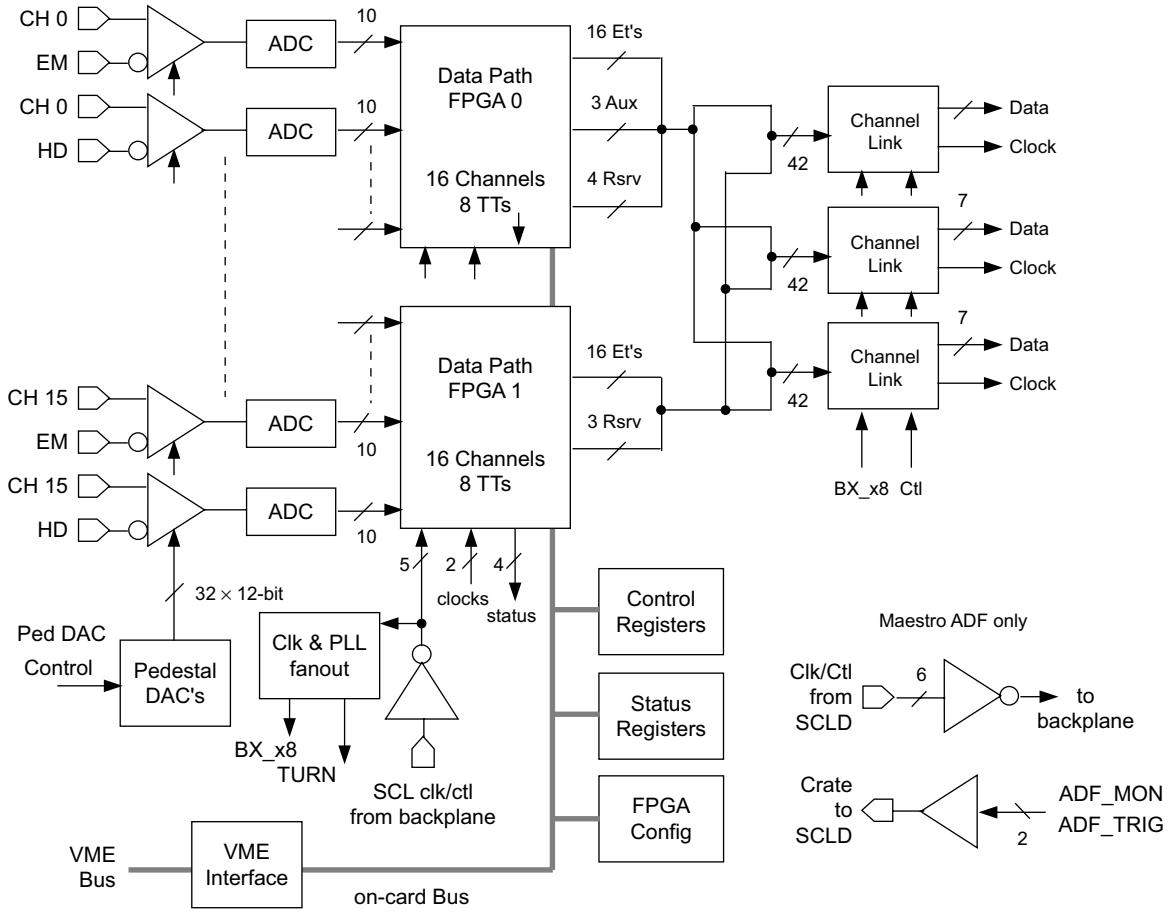

Fig. 11. ADF card block diagram.

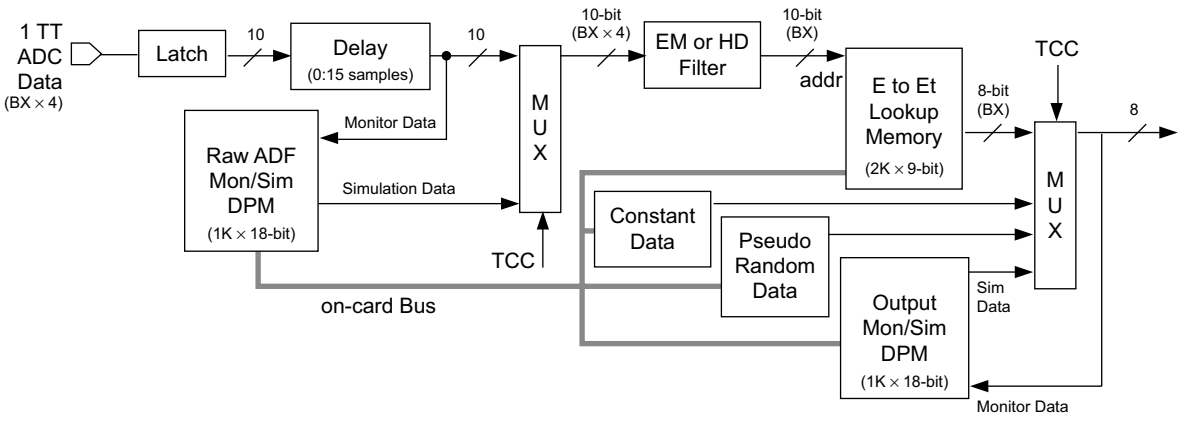

Fig. 12. Block diagram of signal processing for a single TT in the ADF.

system on the SCL, or in a *self-trigger* mode where any TT above a programmable threshold causes writing of all Circular Buffers to stop. Once writing has stopped, all data in the buffers can be read out using the TCC, providing valuable monitoring information on the system's input signals. The Raw ADC Data Circular Buffers can also be loaded by the TCC with simulated data, which can be inserted into the ADF data path instead of real signals for testing purposes.

The Digital Filter in the signal processing path can be used to remove high frequency noise from the trigger signals and to remove low frequency shifts in the baseline. This filter is currently configured to select the ADC sample at the peak of each analog TT signal. This mode of operation allows the most direct comparison with data taken with the previous L1Cal and appears to be adequate for the physics goals of the experiment.

The 10 bit output from the Digital Filter stage has the same scale and offset as the output from the ADCs. It is used as an address to an *E to* E_T *Lookup Memory*, the output of which is an eight-bit data word corresponding to the E_T seen in that TT. This E to E_T conversion is normally programmed such that one output count corresponds to 0.25 GeV of E_T and includes an eight count pedestal, corresponding to zero E_T from that TT.

The eight-bit TT E_T is one of four sources of data that can be sent from the ADF to the TABs under control of a multiplexer (on a channel by channel and cycle by cycle basis). The other three multiplexer inputs are a fixed eight-bit value read from a programmable register, simulation data from the Output Data Circular Buffer, and data from a pseudo-random number generator.

The latter two of these sources are used for system testing purposes. During normal operation, the multiplexers are set up such that TT E_T data are sent to the TABs on those bunch crossing corresponding to real proton–antiproton collisions, while the fixed pedestal value (eight counts) is sent on all other accelerator clock periods. If noise on a channel reaches a level where it significantly impacts the D0 trigger rate, then this channel can be disabled, until the problem can be resolved, by forcing it to send the fixed pedestal on all accelerator clock periods, regardless of whether they contain a real crossing or not. Typically, less than 10 (of 2560) TTs are excluded in this manner at any time.

Data are sent from the ADF system to the TAB cards using a National Semiconductor Channel-Link chipset with LVDS signal levels between the transmitter and receiver [15]. Each Channel-Link output from an ADF card carries the E_T data for all 32 channels serviced by that card. A new frame of E_T data is sent every 132 ns. All 80 ADF cards begin sending their frame of data for a given Tevatron beam crossing at the same point in time. Each ADF card sends out three identical copies of its data to three separate TABs, accommodating the data sharing requirements of the SWA.

7.4. Timing and control in the ADF system

The ADF system receives timing and control signals listed in Table 1 over one of the SCLs [2]. Distribution of these signals from the SCL to the 80 ADF cards is accomplished by the *SCLD* card. The SCLD card receives a copy of the SCL information using a D0-standard SCL Receiver mezzanine card and fans out the signals mentioned in Table 1 to the four VME-64x crates that hold the 80 ADF cards using LVDS level signals. In addition, each ADF crate sends two LVDS level signals (`ADF_MON` and `ADF_TRIG`) back to the SCLD card, allowing the TCC to cause synchronous readout of the ADFs.

Within an ADF crate, the ADF card at the mid-point of the backplane (referred to as the *Maestro*) receives the SCLD signals and places them onto spare, bused VME-64x backplane lines at TTL open collector signal levels. All 20 of the ADF cards in a crate pick up their timing and control signals from these backplane lines. To ensure a clean clock, the `CLK7` signal is sent differentially across the backplane and is used as the reference for a PLL on the ADFs. This PLL provides the jitter-free clock signal needed for LVDS data transmission to the TABs and for ADC sampling timing.

7.5. Configuring and programming the ADF system

The ADF cards are controlled over a VME bus using a VME-slave interface implemented in a PAL that is automatically configured at power-up. Once the VME interface is running, the TCC simultaneously loads identical logic files into the two data path FPGAs on each card. Since each data path FPGA uses slightly different logic (e.g., the output check sum generation), the FPGA *flavor* is chosen by a single ID pin. After TCC has configured all of the data path FPGAs, it then programs all control-status registers and memory blocks in the ADFs. Information that is held on the ADF cards that is critical to their physics triggering operation is protected by making those programmable features "read only" during normal operation. TCC must explicitly unlock the write access to these features to change their control values. In this way no single failed or mis-addressed VME cycle can overwrite these critical data.

8. ADF to TAB data transfer

Digitized TT data from each ADF's 4×4, $\eta \times \phi$ region are sent to the TABs for further processing, as shown in Fig. 13. To accommodate the high density of input on the TABs, the eight-bit serial trigger-tower data are transmitted using the Channel-Link LVDS chipset [15], which serializes 48 CMOS/TTL inputs and the transmission clock onto seven LVDS channels plus a clock channel. In the L1Cal system, the input to the transmitter is 60 MHz TTL (eight times the bunch crossing rate), which is stepped up to 420 MHz for LVDS transmission.

Each ADF sends three identical copies of 36 eight-bit words to three different TABs on each bunch crossing. This data transmission uses eight LVDS channels—seven data channels containing six serialized data words each, and one clock—on Gore cables with 2 mm HM connectors [16]. The 36 data words consist of the digitized E_T of 16 EM and 16 HD TTs and four control words. The bunch crossing number control word indicates which accelerator crossing produced the ADF data being transmitted, and is used throughout the system for synchronization. The frame-bit control word is used to help align the least significant bits of the other data words. The parity control word is the logical XOR of every other word and is used to check the integrity of the data transmission. Finally, one control word is reserved for future use.

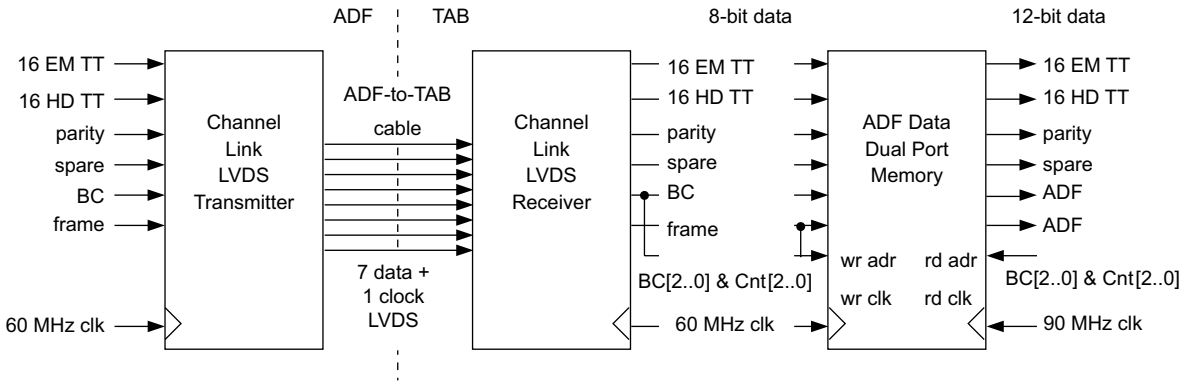

Fig. 13. ADF to TAB data transmission, reception and the dual-port-memory transition from 8-bit to 12-bit data.

While the ADF logic is eight-bit serial (60 MHz) the TAB logic is 12-bit serial (90 MHz). To cross the clock domains, the data passes through a dual-port memory with the upper four bits padded with zeros. The additional bit space is required to accommodate the SWA sums.

The dual port memory write address is calculated from the frame and bunch crossing words of the ADF data. The least significant address bits are a data word bit count, which is reset by the frame signal, while the most significant address bits are the first three-bits of the bunch crossing number. This means that the memory is large enough to contain eight events of eight-bit serial data.

By calculating the read address in the same fashion, but from the TAB frame and bunch crossing words, the dual-port memory crosses 60 MHz/90 MHz clock domains, maintains the correct phase of the data, and synchronizes the data to within eight crossings all at the same time. This means the TAB timing can range between a minimal latency setting where the data are retrieved just after they are written and a maximal latency setting where the data are retrieved just before it is overwritten. If the TAB timing is outside this range, they are the data from eight previous or following crossings will be retrieved.

Although off-the-shelf components were used within their specifications, operating 240 such links reliably was found to be challenging. Several techniques were employed to stabilize the data transmission. Different cable lengths (between 2.5 and 5.0 m) were used to match the different distances between ADF crates and the TAB/GAB crate. The DC-balance and pre-emphasis features of the channel-link chipset [15] were also used, but deskewing, which was found to be unreliable, was not.

9. The TAB/GAB system

9.1. Trigger Algorithm Board

The TABs find EM, Jet and Tau candidates using the SWA and perform preliminary sums for total and missing E_T calculations. Each TAB is a double-wide 9U × 400 mm, 12-layer card designed for a custom backplane. The main functional elements of the TAB are shown in Fig. 14.

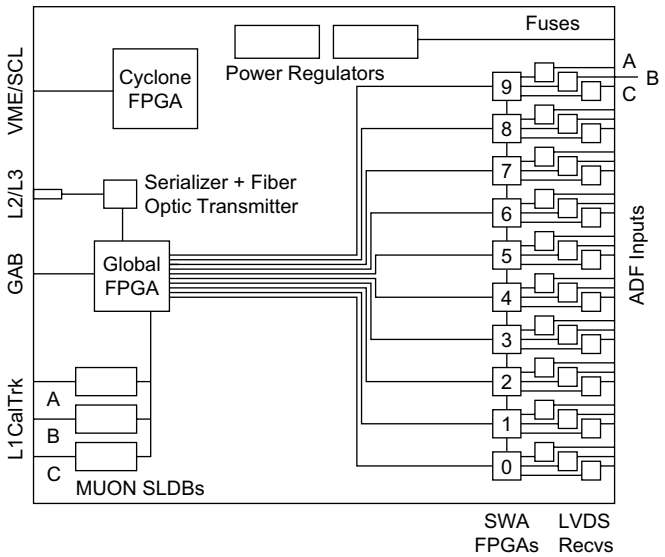

Fig. 14. Block diagram of the TAB.

In the TAB's main trigger data path, LVDS cables from 30 ADFs are received at the back of the card using feedthrough connectors on the backplane. The data from these cables are extracted using Channel-Link receivers [15] and sent, as individual bit streams for each TT, to 10 TAB SWA FPGAs [17] for processing. These chips also pass some of their data to their nearest neighbors to accommodate the data sharing requirements of the SWA. The algorithm output from each SWA is sent to a single TAB global FPGA [17]. The global FPGA calculates regional sums and sends the results out the front of the board to the GAB, over the same type of cables used for ADF to TAB data transmission (see Section 8) using embedded LVDS functionality in the FPGA. This data transmission occurs at a clock rate of 636 MHz.

The global FPGA also sends three copies of Jet and EM object information for each bunch crossing to the L1CalTrk system for processing using Gbit/s serial link transmitter daughter cards (MUON SLDB) [2]. Upon receiving an L1 accept from the D0 TFW, the TAB global chip also sends data out on a serial fiber-optic link [18] for

use by the L2 trigger and for inclusion in the D0 event data written to permanent storage on an L3 accept.

Low-level board services are provided by the TAB Cyclone chip [19], which is configured by an on-board serial configuration device [20] on TAB power-up. These services include providing the path for power-up and configuration of the other FPGAs on the board, under the direction of the VME/SCL card; communicating with VME and the D0 SCL over the specialized VME/SCL serial link; and fanning out the 132 ns detector clock using an on-board clock distribution device [21].

9.2. Global Algorithm Board

The GAB receives data containing regional counts of Jet, EM, and Tau physics objects calculated by the TABs and produces a menu of and/or terms, which is sent to the D0 TFW. Like the TAB, the GAB is a double-wide 9U × 400 mm, 12-layer circuit board designed for a custom backplane. Its main functional elements are shown in Fig. 15.

LVDS receivers, embedded in four Altera Stratix FPGAs (LVDS FPGAs) [17] each receive the output of two TABs, synchronizing the data to the GAB 90 MHz clock using a dual-port memory. The synchronized TAB data from all four LVDS FPGAs are sent to a single GAB S30 FPGA [17], which calculates and/or terms, and sends them to the TFW through TTL-to-ECL converters [22]. There are five 16-bit outputs on the GAB, although only four are used by the framework.

Much like the TABs, upon receiving an L1 accept, the GAB S30 sends data to L2 and L3 on a serialized fiber-optic link [18]. Also as on the TABs, a Cyclone FPGA [19] provides low-level board services.

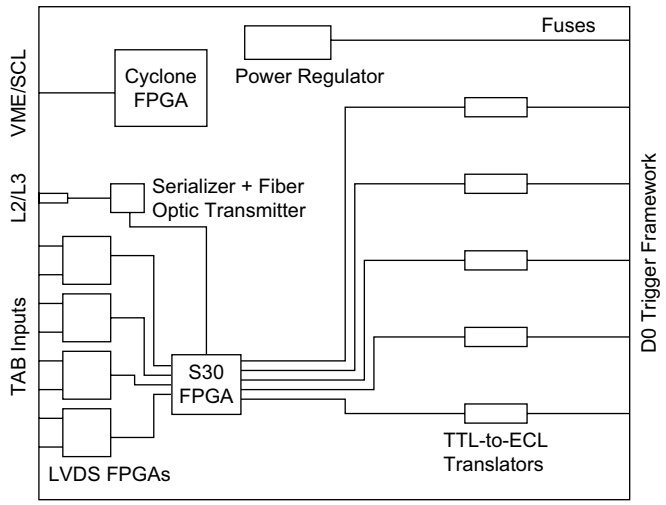

Fig. 15. Block diagram of the GAB.

9.3. VME/SCL board and the TAB/GAB control path

Because of the high-density of inputs to the TAB and GAB modules, direct connections of these cards to a VME bus is impossible. A custom control path for these boards is provided by the VME/SCL module, a double-wide 9U × 400 mm, 8-layer board. A block diagram of the main elements of this card can be found in Fig. 16. SCL signals arrive at the VME/SCL board via an SCL receiver daughter card and those signals used by the TAB/GAB system are selected for fanout by the SCL FPGA [25], which also handles transmission/reception of serialized VME communications with the TABs and GAB. Any VME communication, directed to (from) a card in the TAB/GAB system, is received by (transmitted from) the VME bus PLD, which implements the VME protocol. Those commands whose destination (source) is one of the TAB or GAB boards are translated to (from) the custom serial protocol listed in Table 1 by the Serial VME PLD, which connects to the SCL FPGA for signal transmission (reception). Serial communications between the VME/SCL card and the TABs and GAB is

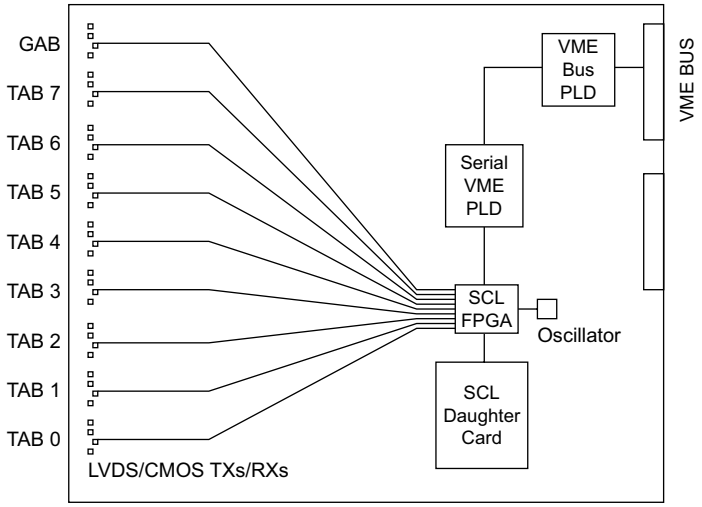

Fig. 16. Block diagram of the VME/SCL board.

accomplished using LVDS protocol [23], on nine cables—one for each TAB and GAB.

9.4. TAB/GAB trigger data path

The path of *trigger data* through the TAB/GAB system is shown in Fig. 17. Each of the eight TABs receives data from 30 ADF cards, covering a 40×12 region in $\eta \times \phi$ space. Eight-bit TT E_T data are translated to 12-bit words in the ADF Data DPM and are transmitted serially to the SWA FPGAs where EM, Jet and Tau objects are found. Each of the 10 TAB SWA chips finds objects in a 4×4, $\eta \times \phi$ grid, for which it requires a 9×9 region of input TTs. This TT data comes from the three LVDS receivers (A,B,C in Fig. 14) attached directly to the chip and also, indirectly, from its nearest neighbor SWA chips. A map of the TT inputs to a single SWA chip is given in Table 3. In this table and the following discussion, we use *global indices* ($\eta[0,39]$ and $\phi[0,31]$) when referring to the entire grid but switch to *local indices* ($\Delta\eta[-2,6]$ and $\Delta\phi[-2,6]$) for single SWA chips. The translation between the two systems is given below:

$$\phi = 4 \times (\text{TAB No.}) + \Delta\phi$$
$$\eta = 4 \times (\text{SWA chip No.}) + \Delta\eta. \quad (1)$$

Note that data for η indices 0, 1, 38, and 39, at all ϕ positions, correspond to signals from the ICR detectors, which can be added to the relevant calorimeter TTs if desired.

Each SWA chip sends the results of its algorithms to the Global Chip as 12-bit serial data on 25 lines. The data transmitted consists of the following:

- The highest of seven possible E_T thresholds passed by EM and Jet objects at each of the 4×4, $\eta \times \phi$ positions considered by this chip, or zero if the object E_T is below all thresholds. This information (three bits for each position and object) is packed into a total of eight, 12-bit words, with each word containing data from the four η locations at a specific ϕ for one object type.
- The highest of seven possible Tau isolation ratio thresholds (see Section 4.3) passed by Tau objects at each of the 4×4, $\eta \times \phi$ positions considered by this chip, or zero if the ratio is below all thresholds. This information is packed in the same way as the EM and Jet object data above.
- The results of the EM isolation and EM fraction calculations (see Section 4.2) for each of the 4×4 locations considered in this chip. A single bit, corresponding to a specific $\Delta\eta, \Delta\phi$ location is set if the EM

Table 3
TT input to a single TAB SWA chip

Cable	*C*		*B*				*A*			Chip
$\Delta\eta/\Delta\phi$	−2	−1	0	1	2	3	4	5	6	
6	72	73	74	75	76	77	78	79	80	SWA $i+1$
5	63	64	65	66	67	68	69	70	71	
4	54	55	56	57	58	59	60	61	62	
3	45	46	47	48	49	50	51	52	53	SWA i
2	36	37	38	39	40	41	42	43	44	
1	27	28	29	30	31	32	33	34	35	
0	18	19	20	21	22	23	24	25	26	
−1	9	10	11	12	13	14	15	16	17	SWA $i-1$
−2	0	1	2	3	4	5	6	7	8	

The TT grid is labeled in the SWA chip local coordinates, $\Delta\eta$ (row), $\Delta\phi$ (column), while individual TTs are labeled, 0–81, as they are used in the firmware.

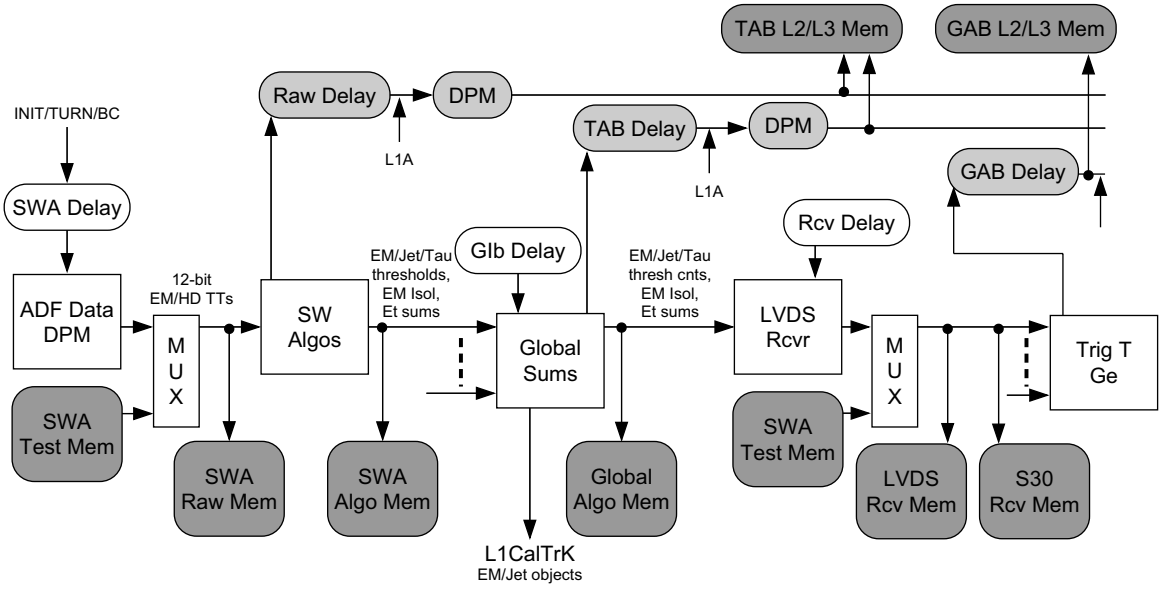

Fig. 17. Data and timing in the TAB/GAB system. The *trigger path* consists of those elements with white background, while the *readout path* is highlighted in light gray and the *test memories* have darker gray backgrounds.

object at that location passes both the EM isolation and EM fraction cuts.
- Sums over four η locations of EM + HD E_T for each ϕ position considered in this chip.
- Four-bit counts of the number of TTs in the chip with EM + HD E_T greater than three programmable thresholds. This information is used to aid in the identification of noisy channels.
- Raw TT E_T's for transmission to the L2 and L3 systems (only transmitted on those BCs marked as L1ACCEPT).
- The bunch crossing number and a flag indicating if there was a bunch crossing number mismatch between the ADF data and the TAB's local BX.
- Status information.
- Two spare lines.

The Global Chip receives these data from the 10 SWA Chips and constructs object counts in a 31×4 region, as well as E_T sums. The reduced number of positions available for TAB object output comes from edge effects in the SWA and from the use of TTs at η indices 0, 1, 38, and 39 for ICR energies. The TABs further concentrate their data by summing object counts in three η ranges—North (N), Central (C), and South (S) [24]—before sending their results to the GAB.

A total of 48 12-bit data words are transmitted from each TAB to the GAB. These data include the following:

- Two-bit counts of the number of EM and Jet objects over each of six possible E_T thresholds in the N, C, and S regions for each of the four ϕ positions considered by the TAB. Each 12-bit word contains counts for all six thresholds for a specific object in an η region and ϕ position.
- Two-bit counts of the number of Tau objects over each of six possible Tau ratio thresholds in the same format as the EM and Jet information above.
- Single bits indicating that at least one EM object passed the isolation criteria in an η region (S,C,N) at a specific ϕ position. Since not enough data lines were available to transmit isolation information for each possible EM object, this grouping represents a compromise that allows the GAB to construct isolated EM triggers if *any* EM object in an η region is found to be isolated.
- Sums of EM + HD E_x, E_y, and scalar E_T over the 40×4 region belonging to the TAB. E_x and E_y are calculated using sine and cosine look-ups appropriate for each TT's ϕ position.
- Eight-bit counts of the number of TTs with EM + HD E_T greater than three thresholds.
- Bunch crossing, status, synchronization and parity information.

At the GAB, TAB data are received and transmitted unchanged to the S30 Chip where and/or trigger terms are constructed as described in Section 13.1. A total of 64 and/or terms are sent from the GAB to the TFW.

9.5. TAB/GAB timing and readout

The timing and readout of the TAB and GAB modules, shown in Fig. 17, are interrelated. Both data traveling on the *trigger path* and on the *readout path* to the L2 and L3 systems on L1ACCEPT must be synchronized so that they correspond to a single, known bunch crossing number. This synchronization is accomplished by setting adjustable *Delay* FIFOs in the TABs and GAB such that the BX_NO stamp on the data at each stage in processing corresponds to the BX_NO being transmitted to the TAB/GAB system by the VME/SCL card. Errors are stored in status registers if a mismatch between these numbers is detected at any point in the chain.

Readout of TAB/GAB data for further processing in the L2 and L3 trigger systems is accomplished by storing data, at various stages of the processing, in pipelines (*Raw Delay*, *TAB Delay*, and *GAB Delay*), whose depth is adjusted so that the data appears at the end of the pipeline when the L1 trigger decision arrives at the boards. If the decision is L1ACCEPT, then the relevant data are moved to Dual Port Memory buffers for transmission, via optical fiber, to the L2 and L3 systems.

Identical data are sent to L2 and L3 by optically splitting the output signals. These data consist of the following:

- The raw eight-bit EM and EM + HD E_T values for each TT (*Raw*).
- A bit-mask with each bit corresponding to a possible EM, Jet or Tau object either set or not depending on whether the object has passed a L2 E_T threshold (*TAB*).
- The set of 64 and/or terms transmitted from the GAB and the total E_T, E_x, and E_y sums (*GAB*).
- A set of control, status and data integrity checksum words.

9.6. TAB/GAB data to L1CalTrk

The L1CalTrk system receives EM and Jet object data for each ϕ position from the TAB Global Chips [26]. Each TAB sends three identical copies of its data (to eliminate cracks in the acceptance) to the L1CalTrk system using three Muon Serial Link Daughter cards. These daughter cards serialize seven 16-bit words per bunch crossing period and transmit them to Muon Serial Link Receiver Daughter cards in the L1CalTrk electronics. Four of these words contain EM and Jet information for each of the ϕ regions considered by the TAB. Each word is broken into seven-bit EM and Jet parts, where bit i is set in each part if any object of that type above threshold i is found in the full η range. A parity word and two spare words are also transmitted.

9.7. TAB/GAB diagnostic memories

The TAB and GAB modules have a series of VME-readable diagnostic memories (see Fig. 17) designed to

capture data from each step of the algorithm calculation. Their contents are snapshots of data transferred between elements of the TAB/GAB system and are generally capable of holding data for 32 consecutive bunch crossing periods, although the *L2/L3 Memories* and the *S30 Trig Memory* are limited to one event's worth of data. These memories are normally written when a `TAB_TRIG` signal is sent from the VME/SCL board under user control. Both the TABs and GAB also have VME-writable test input memories, which allow arbitrary patterns to be used in the place of the incoming data from the ADF or TAB cards.

10. Online control

Most components of the D0 trigger and DAQ system are programmable. The Online System allows this large set of resources and parameters to be configured to support diverse operational modes—broadly speaking, those used during proton–antiproton collisions in the Tevatron (*physics modes*) and those used in the absence of colliding beams (*calibration/testing modes*), forming a large set of resources and parameters needing to be configured before collecting data.

The L1Cal fits seamlessly into this Online System, with its online control software hiding the complexity of the underlying hardware, while making the run time programming of the L1Cal Trigger accessible to all D0 users in simple and logical terms. A diagram of the L1Cal, from an online data and control point of view, is shown in Fig. 18. The main elements of L1Cal online control are listed below, with those aspects specific to L1Cal described in more detail in the following sections. For more information on D0-wide components see Ref. [2].

- The TFW delivers global D0 timing and control signals to the L1Cal and collects and/or terms from the GAB as described in Section 6.
- COOR [2], a central D0 application, coordinates all trigger configuration and programming requests. Global trigger lists, containing requirements and parameters for all triggers used by the experiment, are specified using this application as are more specific trigger configurations (several of which may operate simultaneously) used for calibration and testing.
- The L1Cal TCC, a PC running the Linux operating system, provides a high level interface between COOR and the L1Cal hardware and allows independent expert control of the system.
- The Communication Crate contains cards that provide an interface between the L1Cal custom hardware in the ADF and TAB/GAB crates, and the L1Cal TCC and SCL.
- The L1Cal Readout Crate allows transmission of L1Cal data to the L3 trigger system.
- Monitoring Clients, consisting of software that may run on a number of local or remote computers, display

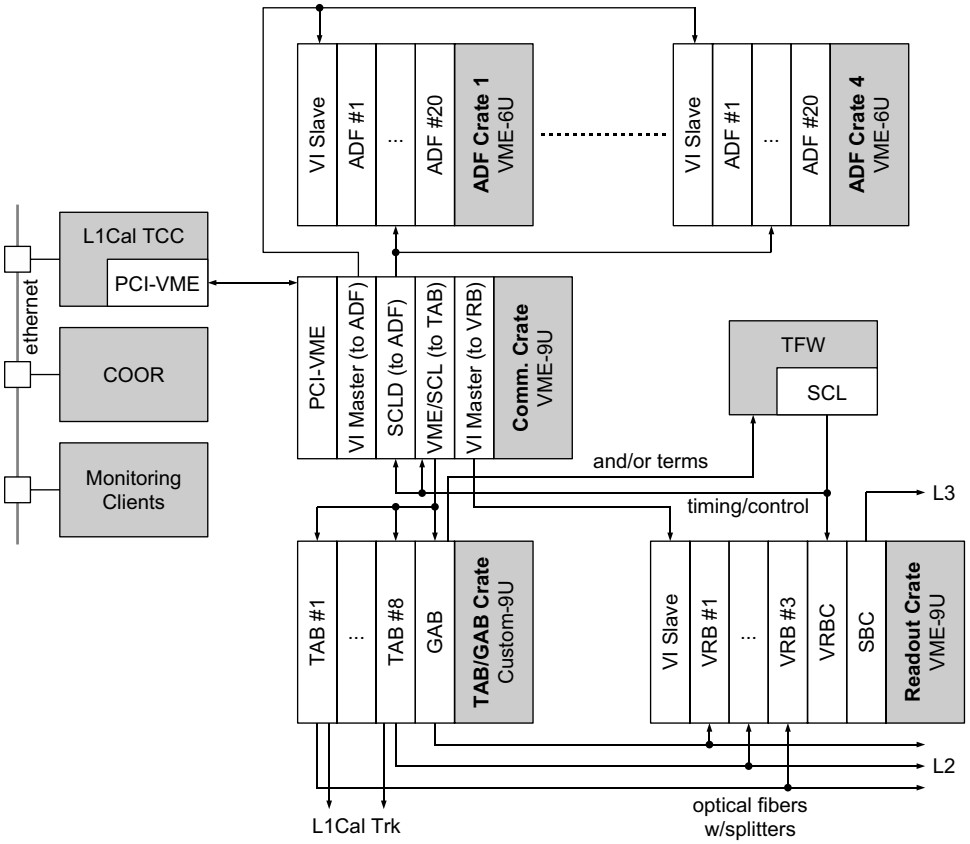

Fig. 18. Communications in the L1Cal system.

information useful for tracking L1Cal operational status.

10.1. L1Cal control path

The L1Cal TCC needs to access the 80 ADF cards in their four 6U VME crates; the eight TAB and one GAB cards in one 9U custom crate; and the readout support cards in one 9U VME crate. It uses a commercial interface to the VME bus architecture—the model 618 PCI-VME bus adapter [27]. This adapter consists of one PCI module located in the TCC PC and one VME card located in the Communication Crate, linked by an optical cable pair.

To access the four ADF crates, the L1Cal system uses a set of Vertical Interconnect (VI) modules built by Fermilab [28]. One VI Master Card is located in the Communication Crate, and is connected to four VI Slaves, one in each ADF Crate. The VI Master maps the VME A24 address space of each remote ADF crate onto four contiguous segments of VME A32 addresses in the Communication Crate. User software running on the L1Cal TCC generates VME A32/D16 cycles in the Communication Crate, and A24/D16 in the ADF crate, via the VI Master–Slave interface. The Communication crate also hosts one additional VI Master to access a VI Slave located in the L1Cal Readout Crate.

As discussed in Section 9.3, VME transactions with the TAB/GAB crate are accomplished via the VME/SCL card, housed in the Communication Crate. User software running on L1Cal TCC generates VME A24/D32 cycles to the VME/SCL, which in turn generates a serialized transaction directly to the targeted TAB or GAB module.

10.2. L1Cal control software

The functionality required from the control software on the L1Cal TCC is defined by three interfaces: the COOR Interface, the L1Cal Expert Interface, and the Monitoring Interface. The first two of these are used to configure and control L1Cal operations globally (COOR) or locally when performing tests (Expert). The Monitoring Interface collects monitoring information from the hardware for use by Monitoring Clients (see Section 11).

The L1Cal online code itself is divided into two parts: the Trigger Control Software (TCS), written in C++ and C, where the main functionality of the above three interfaces is implemented; and the L1Cal Graphical User Interface (GUI), written in Python [29] with TkInter [30], which allows experts to interact directly with the TCS.

While the GUI normally runs on the L1Cal TCC computer, it can also be launched from a different computer located at D0 or at a remote institution. Since it is a non-critical part of the control software, the GUI does not need to run all the time, but several instances of it can be started and stopped as desired, independently from the TCS. Once started, an instance of the GUI communicates with the TCS by exchanging XML (Extensible Markup Language) [31] text strings.

For communications across each of its three interfaces the TCS uses the ITC (Inter Task Communication) package developed by D0 and based on the open-source ACE (Adaptive Communication Environment) software [32]. ITC provides high level management of client–server connections where communication between separate processes, which may be running on separate computers, is dynamically buffered in message queues. The TCS uses ITC to: receive text commands from COOR and send acknowledgments back with the command completion status; receive XML string commands from the GUI application and send XML strings back to the GUI; and receive fixed format binary monitoring requests from the Monitoring Clients and send the requested fixed format binary block of data.

10.3. Main control operations

Control operations in the L1Cal online software fall into three main categories: configuration, initialization and run-time programming. Configuration consists of loading pre-synthesized firmware into all the FPGAs in the system. Initialization then brings the system into a well-defined idle state. During initialization, all control registers, geometric constants, lookup tables, calibration parameters, etc., are overwritten with their desired values. It is also at this stage that problematic TTs are excluded from consideration by programming their corresponding ADF cards to always report zero E_T for the TTs in question. The most IO intensive part of the initialization is in the programming and verification of the 2560 ADF E_T Lookup Memories, which takes approximately 5 s. After initialization, COOR performs the run time programming step, where the specific meaning of each L1Cal trigger output signal (the and/or terms) is defined. This involves loading E_T threshold values and other algorithm parameters into the TABs as well as associating combinations of objects and selection criteria in the GAB with individual output and/or bits. Once these tasks have been accomplished, the system runs largely without external intervention, except for monitoring data collection.

11. Managing monitoring information

The monitoring resources available in the ADF, TAB and GAB cards are described in Sections 7.3 and 9.7. This information is collected by the TCC Control Software and is made available to Monitoring Clients via the Monitoring Interface as outlined in Section 10.2. During normal operation, monitoring data are collected approximately every 5 s when the *Collect Status* qualifier is asserted on the SCL along with `L1ACCEPT`. If data flow has stopped, monitoring data are still collected from the L1Cal, initiated by the TCS, which times out after six seconds of inactivity.

Monitored data include the following:

- The ADF output E_T of all TTs for all 36 active bunch crossings of the accelerator turn containing the

L1ACCEPT for which the *Collect Status* signal is asserted.
- The bunch crossing number within this turn that identifies the L1ACCEPT.
- The contents of all error and status registers in the TABs and GAB (associated with each SWA and Global chip on the TABs and with the LVDS and S30 chips on the GAB). These registers indicate, among other information, synchronization errors on data transfer links, parity errors on each transfer, and bunch crossing number mismatches at various points in the TAB/GAB signal processing chain.

Monitoring information is displayed in the D0 control room and remotely using Monitoring Client GUIs. This application, written in Python [29] with Tkinter [30], requests and receives data from the TCS via calls to ITC. It displays average pedestal values and RMSs for each TT, to aid in the identification of noisy or dead channels, as well as system status information.

Another tool for monitoring data quality in the control room is a suite of ROOT-based [33] software packages called *Examine*. The L1Cal Examine package receives a stream of data from L3 and displays histograms of various quantities related to L1Cal performance, including comparisons between L1Cal and calorimeter precision readout estimates of TT energies. Data distributions can be compared directly to reference curves provided on the plots, which can be obtained either from an earlier sample of data or from simulation.

12. Calibration of the L1Cal

Several methods are employed to ensure that the E_T of individual TTs, used in the system, is correctly calibrated—i.e., that one output count corresponds to 0.25 GeV of E_T and that the zero-E_T baseline is set to eight counts.

12.1. Online pedestal adjustment and noise

The most frequently used of these procedures is a tool, run as part of the TCS, which samples ADC-level data from the ADFs when no true energy is expected to be deposited in the calorimeter. Based on this data, corrections to the DAC values used to set each channel's zero-energy baseline are calculated and can be downloaded to the system.

This online pedestal adjustment is performed every few days because of periodic pedestal shifts that occur in a small number of channels—typically less than ten. These pedestal shifts arise because of synchronous noise, with a period of 132 ns, observed in the system due largely to pickup from the readout of other, nearby detector systems. Although the amplitude of this noise varies from channel to channel (it is largest in only a handful of TTs), its phase is stable over periods of several stores of particle beams in the Tevatron, which sets the timescale for pedestal readjustment.

12.2. Calorimeter pulser

The calorimeter pulser system [2], which injects carefully calibrated charge pulses onto the calorimeter preamps, is also used by the L1Cal to aid in the identification of dead and noisy channels. Special software compares E_T values observed in the ADFs with expectations based on the pattern of preamps pulsed and the pulse amplitudes used. Results are displayed graphically to allow easy identification of problematic channels. In addition to its utility in flagging bad channels, this system also provides a quick way to check that the L1Cal signal path is properly cabled.

12.3. Offline gain calibration

The desired TT response of the L1Cal, 0.25 GeV per output count, is determined by comparing offline TT E_T's to the corresponding sums of precision readout channels in the calorimeter, which have already been calibrated against physics signals. For this purpose, data taken during normal physics running of the detector are used. An example can be seen in Fig. 19. Gain calibration constants, for use in the ADF E_T Lookup Memories, are derived from the means of distributions of the ratio of TT to precision readout channel sums for each EM and HD TT.

Gain coefficients derived in this way have been determined to be stable to within ~2% over periods of months. Thus, this type of calibration is normally performed only after extended Tevatron shutdown periods.

13. Results

13.1. Run IIb trigger list

The trigger list for Run IIb was designed, with the help of the simulation tools described in Section 5, to select all physics processes of interest for the high luminosity running period, and to run unprescaled at all instantaneous luminosities below $3 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. The entire Run IIb L1 trigger menu normally produces an accept rate of up to 1800 Hz. It includes a total of 256 and/or terms, of which 64 come from L1Cal, falling into the following broad categories:

- one- two- and three-jet terms with higher jet multiplicity triggers requiring looser E_T cuts;
- single- and di-EM terms without isolation requirements capturing high energy electrons;
- single- and di-EM terms with isolation constraints (which currently consist of requiring that both the EM/HD and the EM-isolation ratios, described in Section 4.2, be greater than eight) designed for low energy electrons;
- tau terms, which select jets with three different isolation criteria;

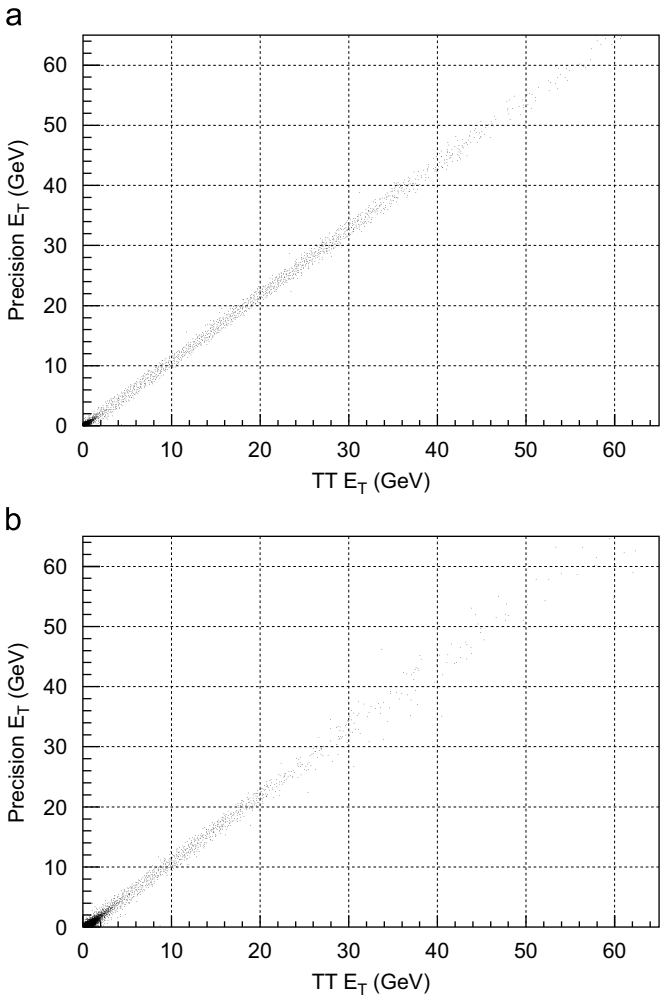

Fig. 19. Precision vs. TT E_T for one EM (a) and one HD (b) TT. The linear relationship with slope = 1 indicates the good calibration of the tower. The excursion away from an absolute correlation is an indication of the inherent noise of the system.

- topological terms, such as a jet with no other jet directly opposite to it in ϕ, targeting specific signals that are difficult to trigger using the basic Jet, EM and Tau objects; and
- missing E_T terms.

These terms can be used individually, or combined using logical ands, to form D0 L1 triggers in the TFW.

13.2. Algorithm performance and rates

L1Cal algorithm performance has been measured relative to unbiased offline reconstruction of jets, electrons, taus, and missing E_T using runs taken at luminosities greater than $1 \times 10^{32}\,cm^{-2}\,s^{-1}$, with a special "low threshold" trigger list designed to minimize trigger bias in the data. Some of these results are summarized in Fig. 20. In Fig. 20(a), turn-on curves (efficiency vs. reconstructed jet E_T) are shown for single jet triggers using Run IIb (with a Jet-object E_T threshold of 15 GeV) and Run IIa (requiring two TTs with $E_T > 5$ GeV) trigger algorithms. The significantly steeper transition between low and high efficiency for the Run IIb algorithm is evident here. The turn-on curve for the Run IIb 20 GeV threshold missing E_T trigger is shown in Fig. 20(b). The performance of this trigger is comparable to, or better than that observed in Run IIa. EM trigger performance is summarized in Fig. 20(c), for a sample of $Z \to e^+e^-$ events, collected using unbiased triggers. The plot shows the efficiency vs. reconstructed EM object E_T for the logical OR of two separate trigger terms, representative of trigger combinations used in electron-based analyses at D0 . The Run IIb terms used are a single EM trigger term with a threshold of 19 GeV, or an isolated single EM trigger with a threshold of 16 GeV; while for Run IIa, the requirements are a single TT with EM $E_T > 16.5$ GeV or two TTs with $E_T > 8.25$ GeV. Both of these triggers produce rates of 370–380 Hz at a luminosity of $3 \times 10^{32}\,cm^{-2}\,s^{-1}$. However, the Run IIb trigger combination gives a sharper turn-on and allows for a lower effective threshold, yielding a significantly higher efficiency for selecting $Z \to e^+e^-$ decays than that achievable using the Run IIa system. Finally, results using the new Run IIb Tau algorithm are summarized in Fig. 20(d). In this plot, trigger turn-on curves are shown for single tau and single jet triggers using a sample of $Z \to \tau^+\tau^-$ candidates, selected offline from events collected using unbiased (muon) triggers. The L1Cal Tau algorithm allows lower object thresholds to be used (15 GeV taus compared to 20 GeV jets) yielding higher signal selection efficiencies for the same trigger rate.

Measured trigger rates using the new algorithms and trigger list are consistent with those based on extrapolations of Run IIa data to Run IIb instantaneous luminosities, shown in Fig. 8. As can be seen, the total trigger rate observed using the new Run IIb list, to which L1Cal contributes more than 50% of the events, fits into the bandwidth limitations of the experiment. A Run IIa trigger list, designed to give the same selection efficiency as the Run IIb list above, would have exceeded these limits by a factor of two or more.

14. Conclusions

The new D0 Run IIb L1Cal trigger system was designed to cope efficiently with the highest instantaneous luminosities foreseen during the Run IIb operating period of the Tevatron at Fermilab. To accomplish this goal clustering algorithms have been developed using a novel hardware architecture that uses bit-serial data transmission and arithmetic to produce a compact, cost-effective system built using commercially available FPGAs. Although data transmission rates in the system approach one tera-bit per second, the system has been remarkably stable since it began to operate at the beginning of Run IIb.

With the Tevatron regularly producing instantaneous luminosities in excess of $2 \times 10^{32}\,cm^{-2}\,s^{-1}$, the new trigger

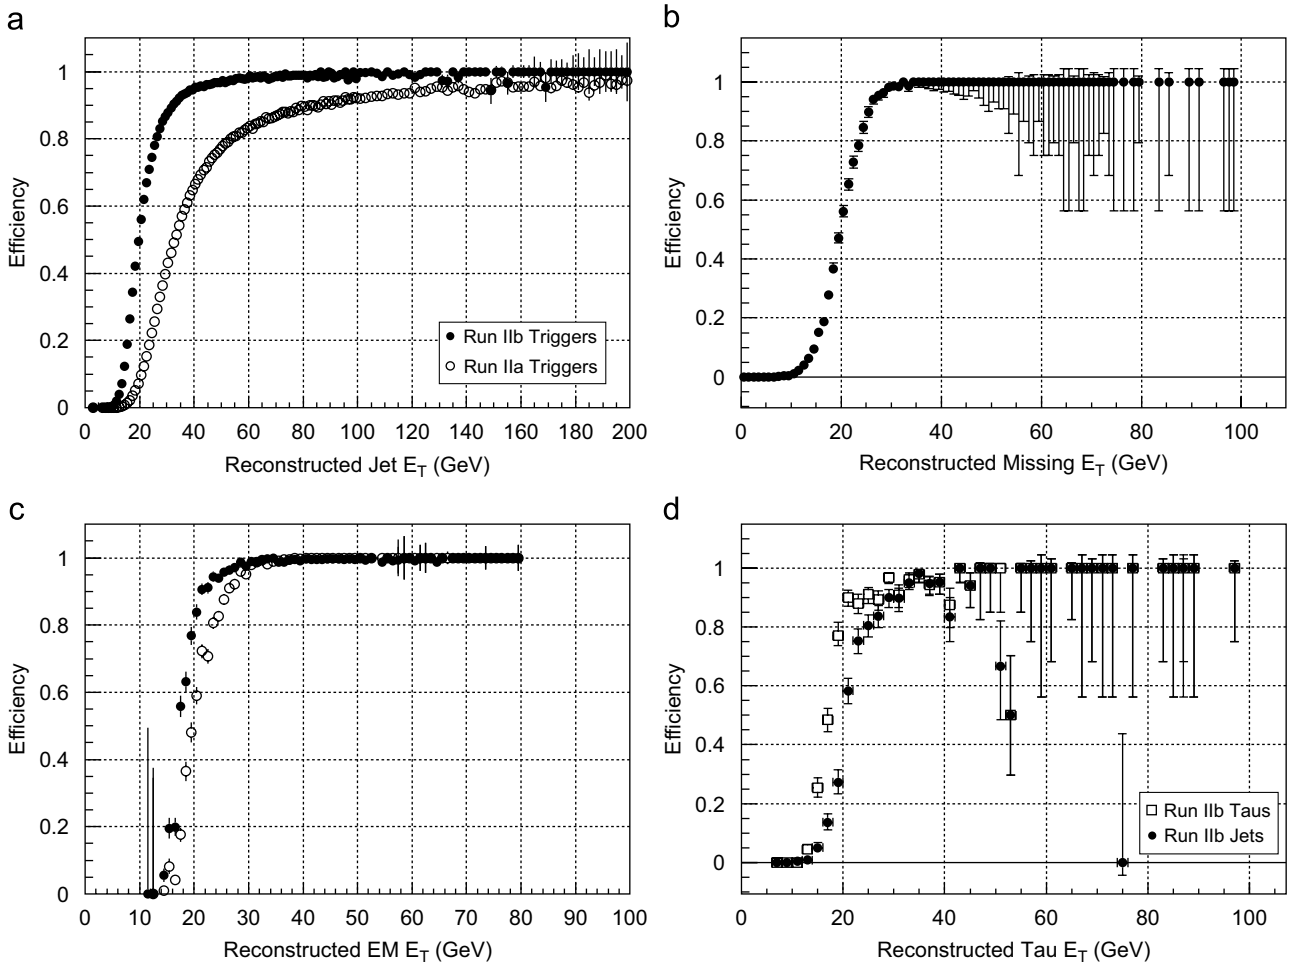

Fig. 20. Trigger turn-on curves for (a) single jet triggers using the Run IIa and Run IIb algorithms; (b) a Run IIb missing E_T trigger; (c) single EM object triggers using Run IIa and isolated Run IIb algorithms on offline selected $Z \to e^+e^-$ events; and (d) a Run IIb single tau trigger compared with a Run IIb single jet trigger that runs at the same rate, using offline selected $Z \to \tau^+\tau^-$ events.

system has been tested extensively at its design limits. So far it has performed exceptionally well, achieving background rejection factors sufficient to fit within the bandwidth limitations of the experiment while retaining the same or better efficiencies as observed in Run IIa for interesting physics processes.

Acknowledgments

We gratefully acknowledge the guidance of George Ginther, Jon Kotcher, Vivian O'Dell, and Paul Padley as managers of the Run IIb project, as well as the technical advice of Dean Schamberger. We would also like to thank Samuel Calvet, Kayle DeVaughan, Ken Herner, Marc Hohlfeld, Bertrand Martin, and Thomas Millet for analyzing L1Cal data and for producing the performance plots shown in this paper. Finally, we thank the staffs at Fermilab and the collaborating institutions, and acknowledge support from the DOE and NSF (USA); CEA (France); and the CRC Program, CFI and NSERC (Canada).

References

[1] Tevatron Run II Handbook ⟨http://www-bd.fnal.gov/lug/⟩.
[2] V.M. Abazov, et al., Nucl. Instr. and Meth. A 565 (2006) 463.
[3] The Tevatron Run II Upgrade Project ⟨http://www-bd.fnal.gov/run2upgrade/⟩.
[4] R. Lipton, Nucl. Instr. and Meth. A 566 (2006) 104; M. Weber, Nucl. Instr. and Meth. A 566 (2006) 182.
[5] M. Abolins, et al., IEEE Trans. Nucl. Sci. NS-51 (2004) 340.
[6] S. Abachi, et al., Nucl. Instr. Meth. A 338 (1994) 185.
[7] The right-handed D0 coordinate system is defined with the z-axis in the direction of the proton beam, with ϕ measuring the azimuthal angle in the plane transverse to the beam direction, with θ measuring the polar angle, and with the pseudo-rapidity, $\eta = -\ln[\tan(\theta/2)]$.
[8] M. Abolins, D. Edmunds, P. Laurens, B. Pi, Nucl. Instr. and Meth. A 289 (1990) 543.
[9] See for example: The ATLAS Level-1 Trigger Group, ATLAS Level-1 Trigger Technical Design Report ⟨http://atlas.web.cern.ch/Atlas/GROUPS/DAQTRIG/TDR/tdr.html⟩ ATLAS TDR-12, 1998; The CMS Collaboration, The TriDAS Project Technical Design Report, vol. 1: The Trigger Systems ⟨http://cmsdoc.cern.ch/cms/TDR/TRIGGER-public/trigger.html⟩, CERN/LHCC 2000-38, 2000; J. Alitti, et al., Z. Phys. C 49 (1991) 17.

[10] The D0 collaboration, Run IIb Upgrade Technical Design Report, Fermilab-Pub-02/327-E, 2002.
[11] WIENER, Plein & Bauss Ltd ⟨http://www.wiener-us.com/⟩.
[12] 3M Corporation, Pleated foil shielded cable (90211 Series) ⟨http://www.3m.com⟩.
[13] Analog Devices, 10-Bit 40 MSPS 3 V Dual A/D Converter (AD9218BST-40) ⟨http://www.analog.com⟩.
[14] Xilinx, Virtex-II FPGA (XC2V1000-4FG456C) ⟨http://www.xilinx.com/⟩.
[15] National Semiconductor, 48-bit Channel Link SER/DES(DS90CR483/4) ⟨http://www.national.com⟩.
[16] W.L. Gore & Associates, Inc., Eye-Opener Millipacs 2/Z-Pack 2 mm HM Cable Connector (2MMS02xx) ⟨http://www.gore.com⟩.
[17] Altera, Stratix Device Family (EP1Sxx) ⟨http://www.altera.com⟩ (EP1S20F780C7, EP1S20F780C6, EP1S10F780C6, and EP1S30F-1020C6 devices are used for the TAB SWA, TAB global, GAB LVDS and GAB S30 FPGAs, respectively).
[18] Agilent Technologies, Transmit/Receive Chip Set (HDMP-1022) ⟨http://www.agilent.com⟩; Stratos Optical Technologies, Optical Gbit Dual Transmitter (M2T-25-4-1-L) ⟨http://www.stratoslightwave.com⟩.
[19] Altera, Cyclone Device Family (EP1C6Q240C7) ⟨http://www.altera.com⟩.
[20] Altera, Serial Configuration Device (EPCS4SI8) ⟨http://www.altera.com⟩.
[21] Cypress Semiconductor Corporation, Clock Distribution Buffer (CY29948AI) ⟨http://www.altera.com⟩.
[22] Fairchild Semiconductor, TTL to ECL Converter (FDLL4148) ⟨http://www.fairchildsemi.com⟩.
[23] National Semiconductor, Differential Line Drivers/Receivers (DS90LV047/8 ATMTC) ⟨http://www.national.com⟩.
[24] As currently implemented the North, Central and South regions correspond to $3.2<\eta<1.0$, $1.0<\eta<-1.0$, and $-1.0<\eta<3.2$, respectively.
[25] Altera, ACEX 1K Series FPGA (EP1K50QC208-2) ⟨http://www.altera.com⟩.
[26] Because of the extra latency involved in the divide operation necessary for the calculation of the Tau algorithm isolation, this information is not available in time for transmission to the L1CalTrk system.
[27] GE Fanuc Embedded Systems, PCI to VME bus adapter (Model 618) ⟨http://www.gefanucembedded.com⟩.
[28] Vertical Interface Board Documentation ⟨http://www-linac.fnal.gov/LINAC/hardware/vmesys/boards/vi/viInfo.html⟩.
[29] Python Programming Language—Official Website ⟨http://www.python.org⟩.
[30] TkInter Documentation Website ⟨http://wiki.python.org/moin/TkInter⟩.
[31] W3C: Extensible Markup Language Website ⟨http://www.w3.org/XML/⟩.
[32] The ADAPTIVE Communication Environment Website ⟨http://www.cs.wustl.edu/~schmidt/ACE.html⟩.
[33] R. Brun, et al., ROOT: An Object-Oriented Data Analysis Framework ⟨http://root.cern.ch/root/⟩.