Forward Production of J/ψ in Hadronic Interactions
and
Calibration of a Large BGO Electromagnetic Calorimeter

Philip E. Kaaret

A dissertation presented to
the Faculty of Princeton University
in candidacy for the degree of
Doctor of Philosophy

Recommended for acceptance
by the Department of Physics

June 1989

Abstract

In the first part of this thesis, we describe an analysis of J/ψ produced in the forward direction in the reaction $\pi N \rightarrow \mu^+\mu^- X$. The data for this analysis were collected by Fermilab experiment E615. We measured the cross section for J/ψ production and the angular distribution of muons from J/ψ decay. We found evidence for longitudinal polarization of J/ψ produced in the kinematic limit where the J/ψ carries a large fraction of the incident pion's longitudinal momentum. This is the first experimental observation of longitudinal polarization of J/ψ produced in hadronic interactions.

In the second part of this thesis, we describe the construction and calibration of a large Bismuth Germanate (BGO) electromagnetic calorimeter designed to study e^+e^- collisions at center-of-mass energies near the Z_0 mass. The calorimeter is a subdetector of the L3 detector and will be installed in the Large Electron Positron collider (LEP) of the European Organization for Nuclear Research. We present the results of a calibration of the calorimeter in an electron test beam at electron energies of 2, 10, and 50 GeV. We show that the accuracy of the calibration is 0.8% at 2 GeV, improving to better than 0.5% at 10 GeV and above.

Acknowledgements

I would first like to thank Jon Bakken for the education and encourage he has given me over the past five years. Working with him, I have learned the most important lesson of my graduate career: keep at it until it works. I would next like to thank the people who helped with my L3 work: Peter Denes, Anne Heavey, Bob Clare, David Stickland, Dick Sumner, Stan Chidzik, Ken Read, Terence Boehringer, Laurent Vuilleumier, Thei Wynen, Sylvie Rosier, Yannis Karyotakis, Peter Klok, Gilles Sauvage, Dominique Boutigny, Werner Ruckstuhl, Luciano Barone, and rest of the L3-BGO group. Working with this varied group of people was a constant source of enjoyment. My thesis advisor, Pierre Piroué, has a vast store of knowledge that extends far beyond physics and has been of great assistance. My advisors on the E615 analysis have been extremely valuable. I greatly enjoyed working with Bill Louis on the analysis and owe many thanks to Kirk McDonald for greatly improving my presentation of my results and my understanding of the experiment. I would also like to thank all of the members of the E615 collaboration for collecting the data. I would particularly like to thank Marcel and Esther Merk and Toby Falk for finding something to do in Geneva while I wasn't working. Finally, I thank my parents for all they have given me. Without them, none of this would have been possible.

To Amy

Contents

List of Tables

List of Figures

Part I

Forward Production of J/ψ in Hadronic Interactions

Chapter 1

Introduction

Since Mendelejev first showed that the elements could be arranged in a periodic table, elucidation of the fundamental structure of matter has been a pursuit at the forefront of science. In the currently accepted view of the world (the Standard Model), matter is divided into two categories: leptons and hadrons. The best available probes show no internal structure in leptons; they appear to be point-like, fundamental objects. However, hadrons do have internal structure. They are believed to be bound states of more fundamental objects, quarks, which are held together by carriers of the strong interactions, gluons.

The essential difference between quarks and leptons is that quarks participate in the strong interaction. Quarks have a quantum number, referred to as color, which is the strong-interaction analog of the electromagnetic charge. Leptons have neutral color. The coupling of the strong interaction is so large that single quarks have never been observed in isolation. The best available theory describing the strong interaction is Quantum Chromodynamics (QCD). The complete Lagrangian for QCD is known, but the strength of the coupling makes calculation in the theory so difficult that it has not been used to calculate the masses of the observed hadrons or the distribution of quarks within hadrons.

The quark-parton model, in which hadrons are treated as a collection of noninteracting quarks and gluons, has been found to provide a useful first approximation to the correct theory. The success of the model is due to a property of the strong interaction known as asymptotic freedom: the QCD coupling becomes small at small distances or, equivalently, large momentum transfer. The parton model does not predict the form of the quark distribution with hadrons. However, it does provide a convenient language to discuss experimental results. In addition, since the strong-interaction coupling constant at high energies is relatively small, the parton model provides a starting point for the calculation of QCD-based corrections for high-energy processes. Predictions based on such calculations provide ways to test some aspects of QCD.

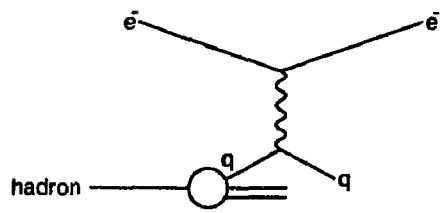

Figure 1.1: Deep-inelastic scattering

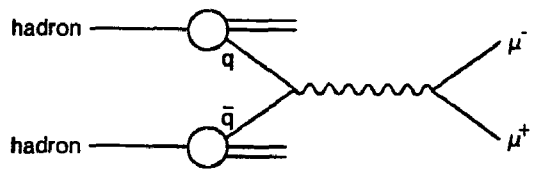

Figure 1.2: The Drell-Yan process

Because leptons do not participate in the strong interaction, it is possible to use leptons as probes to measure the quark distribution within hadrons. The first indication of point-like constituents within hadrons came from 'deep-inelastic' scattering experiments [1]. The fundamental deep-inelastic scattering process in the quark-parton model is shown in Fig. 1.1. In these experiments, a beam of leptons is focused onto a hadron target and the angular distribution of scattered leptons is measured. The experiment is similar to Rutherford's scattering experiment which revealed the presence of nuclei within atoms. At lepton energies large enough to probe subnuclear distances, the observed cross section takes on the form of a sum of cross sections from point-like objects. Each scattering event is the result of an interaction between the lepton and a single quark, therefore the properties of quarks within the hadron can be determined from the scattering cross section.

A complementary approach to the study of hadron structure is to measure the spectrum of dimuons produced in hadron-hadron collisions. This technique was pioneered by a group headed by Lederman in a study of the reaction $pN \rightarrow \mu^+\mu^- X$ [2]. Experiments of this type are often interpreted using a theoretical framework initially developed by Drell and Yan [3]. Their description uses the quark distribution functions and the assumptions of noninteracting quarks and asymptotic freedom of the parton model. They considered electromagnetic annihilation of a quark and an-

tiquark to a virtual photon which then decays to a lepton-antilepton pair, as shown in Fig. 1.2. The process is basically the deep-inelastic scattering diagram of Fig 1.1 with the time order of one lepton leg and one quark leg reversed. In addition to the $\mu^+\mu^-$ pair, a large number of other particles are produced when the two hadrons fragment. However, in the approximation that the transverse momentum of the quark and antiquark are small, the quark-antiquark initial state can be completely determined by the final-state dimuon. Therefore the physics of quarks can be studied by concentrating on the final state of the muon pair. To relate the process to the structure of hadrons the initial state of the hadrons must also be measured.

The reaction $h_1h_2 \rightarrow l^+l^-X$ is more complicated than deep-inelastic scattering because two hadrons are present. The quark distributions of both hadrons must be analysed simultaneously. However, the reaction is a very useful experimental tool because it makes possible the study of the structure of unstable, but long lived, hadrons and because it explores a different kinematic region than deep-inelastic scattering. In particular, experimental results on the Drell-Yan interaction when one quark carries a large fraction of the momentum of an incident pion have shown that the angular distribution of the decay leptons at high x_π is strikingly different from that at low x_π [4] (we will use x_π to indicate the fraction of the incident pion's longitudinal momentum carried by its constituent involved in the interaction). This change in angular distribution was predicted using a model of 'higher-twist' QCD corrections to the initial state of the Drell-Yan process [5]. The kinematic limit at large x_π, inaccessible to deep-inelastic scattering, provides a means by which some of the properties of QCD can be tested.

The process $h_1h_2 \rightarrow l^+l^-X$ can be used to study hadrons produced as intermediate states as well as the properties of the hadrons h_1 and h_2. The virtual photon in the Drell-Yan process can be replaced by any intermediate state with electromagnetic couplings which has the correct quantum numbers. The first-discovered particle containing charmed quarks, the J/ψ, was observed in the reaction $p\mathrm{Be} \rightarrow e^+e^-X$ (it was independently discovered in e^+e^- collisions) [6,7].

This thesis is an analysis of J/ψ produced in the reaction $\pi N \rightarrow \mu^+\mu^-X$ in data collected by Fermilab Experiment E615. E615 is the third in a series of experiments done by a collaboration between the University of Chicago and Princeton University, which studied the reaction $\pi N \rightarrow \mu^+\mu^-X$. The first experiment, E331, mainly studied the production of the J/ψ particle. Measurements were made of the dependence of the J/ψ production cross section on the kinematic variables of the J/ψ and on the nuclear composition of the target [8]. The second experiment, E444, was better able to study the Drell-Yan process[9] and produced the first measurements of the quark structure of the pion [10]. E444 also extended the measurements of the J/ψ made by E331 [12]. Building on the results of the previous experiments,

E615 was designed to extend the study of the Drell-Yan process for dimuons with large invariant mass in the kinematic limit where a single quark carries most of the momentum of the incident pion (high x_π). The experiment was designed to measure only the momenta and trajectories of the outgoing muons, since this provides sufficient information to reconstruct the quark interaction, as discussed above. The experiment had good acceptance for muons emitted in the forward direction, since this is necessary to study interactions of particles with high x_π.

The experiment collected a total of 40,000 dimuon events with invariant mass greater than 4 GeV/c^2. As a 'background' to the desired signal, the experiment also collected more than one million dimuon events from J/ψ decays. This thesis is an analysis of the J/ψ signal. The results of the analysis have been published previously [11]. The goals of the analysis were to determine the dependence of the cross section for J/ψ production on the transverse momentum and on the fraction (x_F) of the total longitudinal momentum of the hadrons carried by the J/ψ, and to study the angular distribution of the muons from J/ψ decays. The analysis extends the results of a separate analysis of an early sample of 80,000 J/ψ from E615 [13].

A unique contribution is made to the understanding of hadronic J/ψ production at high x_F because the design of E615 and the large number of events obtained make possible a much more detailed investigation of this kinematic extreme than was previously possible. Since this analysis was done after the change in the muon angular distribution in the Drell-Yan process at high x_π was observed, particular emphasis was placed on determining if the angular distribution of muons produced by J/ψ with large x_F differed from the isotropic distribution previously found at low x_F [13]. The presence of a change in the angular distribution would taken as an indication of QCD corrections and could provide additional evidence for higher-twist effects in QCD.

In the remainder of this chapter, we introduce the reference frame and kinematic variables used in the analysis, present a theoretical framework for the analysis, and motivate our interest in the angular distribution of the decay muons. In the second chapter we discuss the design of the experiment, emphasizing the problems inherent in the study of the Drell-Yan process and the requirements for good acceptance for forward muons. In the third chapter we describe how the physical quantities of interest are derived from the measurements made. In the last chapter we present our measurement of the cross section for J/ψ production and the angular distribution of the decay muons and discuss our results.

1.1 The Kinematic variables

Before discussing the experiment or the theoretical background in greater detail, we will first introduce the coordinate systems used to describe the interaction. Three coordinates systems are used. The first is the laboratory frame in which all of the measurements were made. The second is the pion-nucleon center-of-mass reference frame, which is used to relate the kinematics of dimuon production to the distributions of quarks (and gluons) within the hadrons. The third frame is the dimuon center-of-mass frame, which is used to analyse the muon angular distribution.

We begin by defining the variables used to describe the incident pion and nucleon and the outgoing muons as measured in the laboratory reference frame. The nucleon was taken to be at rest. No attempt was made to account for the Fermi motion of the nucleons because any correction applied would have necessarily been model dependent.

$$
\begin{aligned}
(E_\pi, \vec{P}_\pi) &= \text{the momentum four vector of the incident pion} \\
M_\pi &= \text{the mass of the pion} \\
M_N &= \text{the average mass of a nucleon} \\
s &= \text{the } \pi N \text{ center-of-mass energy squared} \\
&= 2E_\pi M_N + {M_\pi}^2 + {M_N}^2 \\
(E_\pm, \vec{P}_\pm) &= \text{the momentum four vectors of the } \mu^+ \text{ and } \mu^-
\end{aligned}
$$

Next, we present the variables used to describe the total momentum of the quark-antiquark, or equivalently, muon-antimuon pair. We can define the invariant mass, M, of the pair independent of any coordinate system. For the other variables we chose the direction of the pion momentum as our reference axis, the z-axis, and work in the pion-nucleon center-of-mass frame. We define the longitudinal momentum, P_L, of the pair as the component of the momentum parallel to the z-axis and the transverse momentum, P_T, as the momentum perpendicular to the z-axis. For the analysis, rather than use P_L, we will use x_F, which is the fraction of the maximum possible longitudinal momentum, $P_{L,max}$, carried by the pair. We have calculated $P_{L,max}$ assuming a massless recoil system. One more coordinate, Φ, is needed to completely specify the muon-pair-momentum four vector. We define Φ as the angle between the pair momentum and vertical in the plane perpendicular to the incident-pion momentum.

$$
M = \text{the invariant mass of the } \mu^+\mu^- \text{ pair}
$$

$$
\begin{aligned}
&= \sqrt{(E_+ + E_-)^2 - (\vec{P}_+ + \vec{P}_-)^2} \\
P_L &= \text{the longitudinal momentum of the pair} \\
&= \frac{E_\pi + M_N}{\sqrt{s}}(\vec{P}_+ + \vec{P}_-)\cdot\hat{P}_\pi - |\vec{P}_\pi|\frac{E_+ + E_-}{\sqrt{s}} \\
P_{L,max} &= \text{the maximum possible } P_L \text{ in the } \pi N \text{ center-of-mass frame} \\
&= \frac{\sqrt{s}}{2}\left[(1-\frac{M^2}{s})^2 - \frac{4{P_T}^2}{s}\right]^{1/2} \\
x_F &= \text{the longitudinal momentum fraction of the pair} = P_L/P_{L,max} \\
P_T &= \text{the transverse momentum of the pair} \\
&= |(\vec{P}_+ + \vec{P}_-) - (\vec{P}_+ + \vec{P}_-)\cdot\hat{P}_\pi\hat{P}_\pi| \\
\Phi &= \text{the azimuthal angle of } (\vec{P}_+ + \vec{P}_-)
\end{aligned}
$$

We now introduce coordinates to describe the quark (or gluon) momenta before the interaction. The system of coordinates used was chosen to describe the Drell-Yan process, *i.e.*, an interaction with a quark-antiquark pair in the initial state. The initial state in J/ψ production is thought to be more complex, involving gluons as well as quarks, but we will use the same coordinate system. The variables x_π and x_N are the longitudinal momentum fractions of the pion and nucleon constituents involved in the interaction. They are related to the momentum fraction of the muon pair, x_F, the center-of-mass energy squared, s, and $\tau = M^2/s$ by

$$x_\pi x_N = \tau \quad \text{and} \quad x_\pi - x_N = x_F(1-\tau) \tag{1.1}$$

The last set of coordinates we need to introduce describe the angular distribution of the muons. The choice of reference frame should be made in a way which highlights the physics of interest. One of the axes of the reference frame should be parallel to the momenta of the interacting quarks. Because of the motion of the quarks within the hadrons, we are ignorant of the exact conditions of the quark interaction and must instead define the frame in terms of the motion of the hadrons. Because the apparatus was designed with its longitudinal axis closely parallel to the beamline, the analysis is simplest when the incident pion momentum is used as a reference axis. We have chosen to use the t-channel, or Gottfried-Jackson, frame. This frame is a rest frame of the dimuon system which has its z^*-axis aligned with the pion momentum.

We define the muon angles relative to the μ^+. Since the t-channel frame is a rest frame of the muon pair the μ^- angles are the complements of those for the μ^+. The angles are shown in Fig. 1.3. In the t-channel frame the z^*-axis is aligned with momentum of the incident pion, $\vec{\pi}$, and the x^*-z^* plane is determined by $\vec{\pi}$

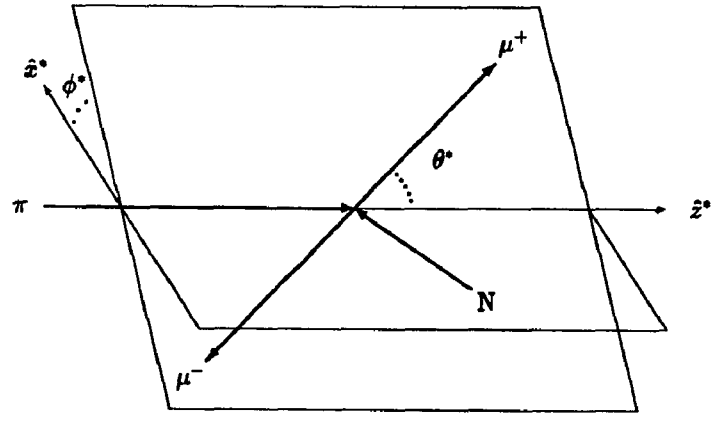

Figure 1.3: Momenta in the t-channel

and momentum of the nucleon, $\vec{N}$, so that θ^* is the angle between the μ^+ and π momenta and ϕ^* is the angle between the $\mu^+\mu^-$ plane and the πN plane. The variables used to describe the muons in the final state are

$$\begin{aligned} \phi^* &= \text{the angle of the } \vec{p}_+ \text{ relative to the plane containing } \vec{p}_\pi \text{ and } \vec{p}_N \\ \cos\theta^* &= \text{the cosine of the angle between } \vec{p}_+ \text{ and } \vec{p}_\pi \end{aligned}$$

In summary, we used as independent variables M, x_F, P_T, Φ, $\cos\theta^*$, and ϕ^*. These six variables are calculated from the six variables (the momentum three vectors of two muons) found by the track-reconstruction procedure. The determination of x_F also requires knowledge of the incident momentum of the pion. Since the target was not polarized, the cross section cannot depend on Φ and we have averaged over this variable in all of the results presented in this analysis.

1.2 Hadron Structure

Current theories of the interactions of subatomic particles deal primarily with the quark and gluon contents of particles rather than the particles themselves. Fermilab experiment E615 studied the interaction of high-energy pion beams (both π^+ and π^-) with a tungsten target. We are therefore interested in the quark and gluon composition of π^+, π^-, protons, and neutrons.

Particle	p n	π^+ π^-
Valence quarks	uud udd	$u\bar{d}$ $\bar{u}d$
β (valence)	3-4	1.2 ± 0.2
β (sea)	8 ± 1	8 ± 1
β (gluon)	4-7	2-3

Table 1.1: Constituents of p, n, π^+, and π^-

The constituents of hadrons are divided into three groups of particles, valence quarks, sea quarks, and gluons. The valence quarks determine the observed quantum numbers of the hadron. The valence quarks of the four hadrons of interest are given in Table 1.1. The sea quarks are 'virtual' particles created in $q\bar{q}$ pairs from quantum fluctuations of the fields within the hadron. The sea is thought to consist of equal parts of $u, \bar{u}, d, \bar{d}$ with the $s, \bar{s}$ component suppressed by a factor of 2. The relative proportion of a type of quark in the sea is inversely proportional to its mass, so only a small component of $c, \bar{c}$ and heavier quarks. The third group of particles are the gluons.

Most experiments on hadron structure done at high momentum transfer have concentrated on the longitudinal-momentum distribution of the constituents. This is the variable of primary importance in scattering experiments. The spin distribution is also of interest, but its measurement requires a polarized beam and target. The experiment considered here used an unpolarized beam and target.

The longitudinal-momentum distribution of the constituents is commonly described by a structure function, $f(x)$, defined as the probability that a constituent carries a fraction x of the total momentum of the hadron multiplied by the momentum fraction. The structure-functions used in this analysis were taken from a previous analysis of the E615 experiment[15]. The nucleon structure function show that approximately 35% of the total momentum is carried by the valence quarks, 15% is carried by the sea quarks, and the remaining 50% is carried by the gluons.

A few general properties of the structure functions are important for this analysis. We are interested in limit where a single constituent carries all of the longitudinal momentum of the hadron, $x \to 1$. In this limit, the structure functions of all three types of constituents are well approximated by the form $f(x) \sim (1-x)^{\beta}$. Approximate values for the exponent β are given in Table 1.1; the values are taken from [13,14,15]. Large values of β mean that the contribution of the constituent decreases rapidly at large x. We can see from Table 1.1 that pions are much more likely to have constituents with high x than nucleons and that the valence-quark

component dominates at high x in both types of hadrons. Because pions have a valence antiquark, they have a particularly large fraction of high x antiquarks. Since Drell-Yan events from nuclear targets are produced mainly by the antiquark component of the incident beam, pions beams are preferred in the study of the Drell-Yan process at high x_F.

1.3 Drell-Yan Interactions

We will discuss the Drell-Yan process before discussing J/ψ production because it is a simpler and better-understood process. It also introduces all of the language used to discuss J/ψ production.

Interpretation of the production of $\mu^+\mu^-$ pairs in hadron collisions via a quark-antiquark annihilation model was first considered by Drell and Yan in 1970 [3]. The process is shown in Fig. 1.2. The basic annihilation process is electromagnetic and has a cross section of $4\pi\alpha^2 e_q^2/9M^2$, where e_q is the charge of the quark-antiquark pair and M is the dimuon invariant mass. This cross section includes a factor of $1/3$, which is the probability that a quark-antiquark pair will be colorless as is required to form a $\mu^+\mu^-$ pair. To calculate the Drell-Yan cross section for πN interactions we must sum over the momentum distributions for all of the quarks present; gluons have no contribution because the process is electromagnetic.

$$\frac{d\sigma}{dMdx_F} = \frac{8\pi\alpha^2}{9M^3}\kappa\sum_q \frac{1-\tau}{x_\pi + x_N} e_q^2 \left[f_{q/\pi}(x_\pi) f_{\bar{q}/N}(x_N) + f_{\bar{q}/\pi}(x_\pi) f_{q/N}(x_N)\right] \qquad (1.2)$$

where $f_{q/\pi}(x_\pi)$ and $f_{q/N}(x_N)$ are the quark structure functions for the pion and nucleon evaluated at mass M. For application to a particular nucleus, $f_{q/N}(x)$ is calculated from the proton and neutron momentum distributions weighted by the relative numbers of protons and neutrons in the nucleus. The factor κ is a strong-interaction correction to the initial state of the parton model. First-order calculations of κ have shown it to be nearly independent of the momentum fractions of the quarks and the mass of the dimuon [16]. Measurements of the Drell-Yan cross section show that $\kappa \sim 2$ [17].

Over most of the kinematic range, the dependence of the cross section on x_F and s and the angular distribution of the dimuons are well described by the Drell-Yan model. However, a striking deviation from the model arises in the muon angular distribution at large x_F. In the approximation that the energy of a particle is much greater than its mass, the helicity of the particle will be conserved by the electromagnetic interaction. Summing the cross sections for the helicity-conserving spin states of $q\bar{q} \to \mu^+\mu^-$, we find a muon angular distribution in the final state

of the form $1 + \alpha \cos^2 \theta^*$ with $\alpha = 1$. As noted above, θ^* is the angle between the momentum vectors of the μ^+ and the incident pion in the muon-pair rest frame. We assume the antiquark momentum is parallel to the incident pion momentum.

The angular distribution, as we have presented it above, is a measure of the alignment of the momenta of the outgoing muons with the momenta of the incoming quarks. If we assume that the muons are produced in the decay of an intermediate particle, a virtual photon in the case of Drell-Yan, we can relate the measured angular distribution to the alignment of the spin of the intermediate particle with the quark momenta using the quantum theory of angular momentum. In the discussion below, we assume that the intermediate particle has spin-1 and decays to a pair of spin-$\frac{1}{2}$ particles. The language used to describe the decay angular distribution and the spin alignment of intermediate spin-1 particles comes from electromagnetic theory.

In general, the parameter α can take on values ranging from -1 to $+1$. The case $\alpha = 1$ results when the intermediate particle has equal components of helicity $= +1$ and helicity $= -1$. This corresponds to a photon with electric and magnetic field vectors pointing perpendicular to its direction of motion and is called transverse polarization. The case $\alpha = 0$ arises when all helicity states of the intermediate particle are equally populated. We will refer to intermediate particles which produce an angular distribution with $\alpha = 0$ as unpolarized. The last extreme is when $\alpha = -1$ which occurs when the intermediate particle has helicity $= 0$. This case corresponds to a photon with an electric field vector aligned with its momentum and is called longitudinal polarization.

At low x_F the prediction of $\alpha = 1$ agrees with the experimental results. However, in measurements of the angular distribution at large x_F, Fermilab experiment E444 found a change in the angular distribution [4]. The results were consistent with an angular distribution with $\alpha = -1$ in the limit $x_F \to 1$. The result was confirmed and strengthened by cross section measurements made with two separate data sets taken by E615 [18,15]. The leading explanation of the change in angular distribution comes from a model of 'higher-twist' QCD corrections to the initial state of the reaction. A theory constructed by Berger and Brodsky, which takes into account the interactions of the two valence quarks in the pion, predicts the angular distribution will obey

$$d\sigma \sim (1 - x_\pi)^2(1 + \cos^2 \theta^*) + \frac{4}{9}\frac{k_T^2}{M^2} \sin^2 \theta^* \tag{1.3}$$

where x_π is the longitudinal momentum fraction of the annihilating antiquark and k_T is its transverse momentum [5]. This prediction adequately matches the experimental data.

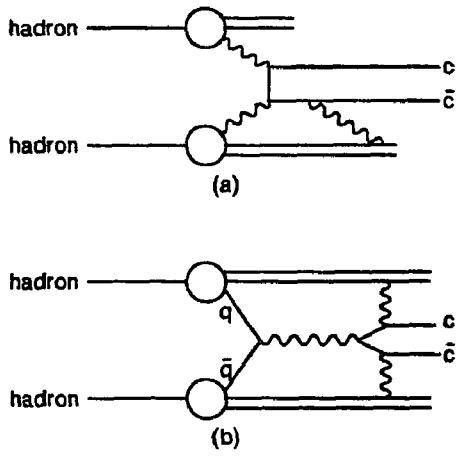

Figure 1.4: J/ψ production processes

1.4 J/ψ Production

The production of J/ψ is more complicated than the Drell-Yan process because the J/ψ participates in the strong interaction. The J/ψ can couple to quarks via the electromagnetic and strong interactions and can couple to gluons via the strong interaction. In addition, the J/ψ can be produced in the decays of other $c\bar{c}$ bound states, predominately the χ. A list of some possible mechanisms is presented below.

1. $gg \rightarrow c\bar{c} \rightarrow \mathrm{J}/\psi + g$
2. $q\bar{q} \rightarrow g \rightarrow c\bar{c} \rightarrow \mathrm{J}/\psi + gg$
3. $gg \rightarrow \chi \rightarrow \mathrm{J}/\psi + \gamma$
4. $q\bar{q} \rightarrow \chi \rightarrow \mathrm{J}/\psi + \gamma$
5. $q\bar{q} \rightarrow \mathrm{J}/\psi + X$

The dominant production mechanisms over most of the range of x_F are thought to be gluon fusion to a $c\bar{c}$ pair with emission of a soft gluon to balance the quantum numbers, and $q\bar{q}$ annihilation via a gluon to a $c\bar{c}$ pair which forms a J/ψ after the emission of two gluons [19]. These two processes (1 and 2 in the list) are shown in Fig. 1.4. Three gluons are required in both cases to form a colorless spin-1 state from massless colored spin-1 gluons. These processes are thought to have a large

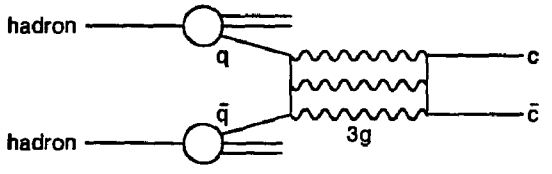

Figure 1.5: Quark-antiquark annihilation to J/ψ

cross section because the radiated gluons can have very low energy. The emission of low-energy gluons is highly probable.

The next most important mechanism is production of intermediate $c\bar{c}$ bound states which then decay to a J/ψ, these are processes 3 and 4 in the list.[1] As we will discuss in more detail below, the cross section for the production of a resonance is proportion to the width of the resonance. The width of the χ states are more than an order of magnitude larger than that of the J/ψ, so their cross sections are correspondingly larger. The fraction of J/ψ produced in this manner can be measured by looking for photons, from χ decay, produced simultaneously with the J/ψ. The WA11 experiment performed such a measurement and found that $18\pm4\%$ of the J/ψ come from χ_1 decays and $13\pm2\%$ come from χ_2 decays [21]. The fraction from χ_0 decays is less than 1%. From measurements of the ψ' cross section from E615 and the known branching ratio to J/ψ, it has been calculated that 7% of the J/ψ are the product of ψ' decays. Combining these data, we see that 40% of J/ψ are produced via decay of intermediate mesons.

Over most of the range of x_F, J/ψ production via quark-antiquark annihilation (process 5 in the list, shown diagrammatically in Fig 1.5) contributes only a small fraction of the total cross section. We will see later that direct production of J/ψ through $q\bar{q}$ annihilation can account for the total observed cross section at high x_F. The rate of the gluon fusion in this kinematic extreme is greatly reduced because the valence-quark component of hadrons dominates over the gluon and sea-quark components at high x_F. Production of J/ψ through intermediate particles becomes unlikely near $x_F = 1$ because the intermediate particle is already at the kinematic limit. The momentum taken by the extra particle involved in the decay tends to decrease the momentum fraction of the J/ψ.

[1] In fact, some authors consider this process more important than the direct J/ψ production processes of the preceeding paragraph [20].

The quark-antiquark annihilation process is of interest to us for two reasons. First, it is important at high x_F, the limit of interest for this experiment. Second, the interaction is closely related to the Drell-Yan model. If quark-antiquark annihilation is the dominant J/ψ production process at high x_F, then we can expect a change in the muon angular distribution, analogous to change observed in the Drell-Yan interactions, at high x_F. The quark-antiquark annihilation model would provide a framework for interpretation of the such an effect. We present a calculation of the process, without any higher-twist corrections, in order to estimate the contribution of $q\bar{q}$ annihilation at large x_F.

Since J/ψ production via the $q\bar{q}$ annihilation process is a resonance phenomena, we can calculate the cross section via the Breit-Wigner resonant production formula. The total width of the J/ψ is $\Gamma_{J/\psi} = 0.062$ MeV. Since the mass resolution of the experiment, 180 MeV/c^2 as discussed later, is much larger than the intrinsic width of the J/ψ, the observed J/ψ resonance will have a shape determined by the experimental resolution function and a total cross section given by the integral over the entire Breit-Wigner resonance cross section. To derive the cross section for hadrons we must weight the point cross section using the structure functions, as for the Drell-Yan case. The resulting integral is identical to the Drell-Yan cross section, Eq. (1.2), with the electromagnetic cross section replaced by the Breit-Wigner cross section [63].

$$\int \frac{4\pi\alpha^2 e_q^2}{3M^2} dM \rightarrow \frac{6\pi^2}{M_{J/\psi}^2} B_{\mu^+\mu^-} B_{q\bar{q}} \Gamma_{J/\psi} \tag{1.4}$$

where $B_{q\bar{q}}$ is the branching ratio of $J/\psi \rightarrow q\bar{q}$, $B_{\mu^+\mu^-} = 0.07$ is the branching ratio of $J/\psi \rightarrow \mu^+\mu^-$, and $M_{J/\psi} = 3.091$ GeV/c^2 is the J/ψ mass. We estimate $B_{q\bar{q}}$ by assuming that the total hadronic width of the J/ψ goes to equally to $u\bar{u}$, $d\bar{d}$, and $s\bar{s}$, and that there are no direct decays to multiquark states. The branching fraction to any given type of quark will then be the hadronic branching ratio divided by three, so $B_{q\bar{q}} = 0.29$. This estimate could be off by as much as a factor of 2. The cross section for $\pi N \rightarrow J/\psi \rightarrow \mu^+\mu^-$ is then

$$\frac{d\sigma_{J/\psi}}{dx_F} = \frac{4\pi^2}{M_{J/\psi}^3} B_{\mu^+\mu^-} B_{q\bar{q}} \Gamma_{J/\psi} F(x_F) \tag{1.5}$$

$$\text{where} \quad F(x_F) = \kappa \sum_q \frac{1-\tau}{x_\pi + x_N} \left[f_{q/\pi}(x_\pi) f_{\bar{q}/N}(x_N) + f_{\bar{q}/\pi}(x_\pi) f_{q/N}(x_N) \right] \tag{1.6}$$

The function $F(x_F)$ is a summation over structure functions. The factor κ is a strong-interaction correction to the parton-model initial state which we assume is the same for the Drell-Yan process and J/ψ production via $q\bar{q}$ annihilation.

The J/ψ production model developed so far is a direct copy of the Drell-Yan model of quark-antiquark annihilation; the only difference is the fundamental cross section used. We can attempt to extend the analogy past the Drell-Yan model and write down a form for the x_F dependence of the angular distribution, based on the higher-twist model of Berger and Brodsky. Experimentally, we know that the J/ψ angular distribution at low x_F is isotropic. A change to longitudinal polarization at high x_F can be parameterized by

$$d\sigma \sim (1 - x_F)^2 + k \sin^2 \theta^* \tag{1.7}$$

We have used x_F in place of the x_π that appears in the Berger-Brodsky formula of Eq. (1.3). The difference between x_F and x_π over the range $x_F > 0.75$ where we will apply this formula is less than 2%.

If $q\bar{q}$ annihilation were the only J/ψ production process then the parameter k would be given by the Berger-Brodsky result, Eq. (1.3), for the Drell-Yan process, $k = 4k_T^2/9M^2$. Since there are other J/ψ production processes, the fraction of J/ψ which are polarized, and therefore the parameter k, are reduced. We will assume that total cross section for unpolarized J/ψ production scales as $(1 - x_F)^2$ relative to the longitudinally polarized part of the cross section. We can then write k as

$$k = A\frac{4}{9}\frac{k_T^2}{M^2} \tag{1.8}$$

where k_T^2 and M are defined in the discussion following Eq. (1.3), and A is the ratio of the cross section for the production of unpolarized J/ψ via $q\bar{q}$ annihilation to the total cross section for the production of unpolarized J/ψ. We can determine the value of A, independent of any angular distribution measurement, by comparing the total observed cross section for J/ψ production to the cross section for J/ψ production from $q\bar{q}$ annihilation calculated from Eqs (1.5)-(1.6). We find that

$$A = \frac{(1 - x_F)^2}{R[(1 - x_F)^2 + \nu] - \nu} \quad \text{where} \quad R = \frac{\sigma_{tot}}{\sigma_{q\bar{q}}} \quad \text{and} \quad \nu = \frac{8}{27}\frac{k_T^2}{M^2} \tag{1.9}$$

The variable R is the ratio of the total cross section for J/ψ production (σ_{tot}) to the cross section for J/ψ production via $q\bar{q}$ annihilation ($\sigma_{q\bar{q}}$), and ν is related to the cross section for the production of longitudinally polarized J/ψ.

1.5 Nuclear Effects

Until now we have considered the interaction of the J/ψ with each nucleon individually. However, we must also consider the possibility that the production

cross section depends on the properties of the nucleus or that the J/ψ will interact in the nucleus after it is created. The lifetime of the J/ψ is 10^{-20} seconds. This is significantly longer than typical strong-interaction time scales. Therefore, it is very possible that the J/ψ will undergo interactions while still within the nucleus.

The dependence of the cross section for J/ψ production, $\sigma_{J/\psi}$, has been studied by the NA3 collaboration at CERN [22] and by Fermilab experiment E444 [23]. NA3 measured $\sigma_{J/\psi}$ for pions (both positive and negative) incident on hydrogen nuclei (atomic weight = $A = 1$) and platinum nuclei ($A = 195.1$). Assuming a power law dependence on the atomic weight, $\sigma_{J/\psi} = A^{\delta}$, the NA3 results give an exponent $\delta = 0.96 \pm 0.01$ for the cross section integrated over all x_F and all P_T.

Also interesting are the results on the x_F and P_T dependence of δ. The NA3 experiments found that δ decreases with increasing x_F. The results quoted are that $\delta = 0.97$ for $0 < x_F < 0.2$ decreasing to $\delta = 0.77$ for $0.8 < x_F < 1.0$. This is in opposition to the E444 data which show no variation in δ versus x_F. There is statistically significant disagreement between the two experiments. For the P_T dependence, the NA3 analysis shows $\delta = 0.91$ for $P_T < 0.5$ GeV/c, increasing to $\delta = 1.09$ for $3.5 < P_T < 4.0$ GeV/c. The E444 results show the same trend. The enhancement in heavier nuclei at large P_T is interpreted as evidence that the J/ψ reacts in the nucleus after it is created. These reactions increase the P_T of the J/ψ after it is produced.

Chapter 2

Apparatus

The E615 apparatus is shown in figure 2.1 [24]. The main parts of the apparatus are:

- the tungsten target.
- the absorber, used to filter out all particles except muons.
- the selection magnet and trigger scintillators, used to select muon pairs with high invariant mass.
- the analysis magnet and wire chambers used to measure the momenta and trajectories of the muons.

A very important consideration in the design of E615 was the reduction of background events. The cross section for the Drell-Yan process is much smaller than the cross sections for the background hadronic processes. In E615 only 1 in 10^9 incident pions produced a $\mu^+\mu^-$ event with invariant mass greater than 4 GeV/c^2. Most of the beam interacted hadronically, producing a large background of particles. In addition, 20% of beam did not interact in the target. Absorbing these background particles without also eliminating the desired muon signal was possible because muons readily penetrate matter. Several meters of material, placed between the target and the detector, stopped electrons and hadrons leaving only muons and neutrinos. Therefore, any charged particle emerging from the absorber was a muon.

The remaining background can be divided into two categories: muons from the decay of beam particles (halo muons), and muons from the decay or interaction of hadrons produced in the target (secondary muons). Holes in the trigger scintillators along the beamline prevented triggering on most halo muons. Halo muons outside the beam pipe were rejected by scintillators upstream of the target. Secondary muons could be produced from secondary hadrons either singly by weak decays, or

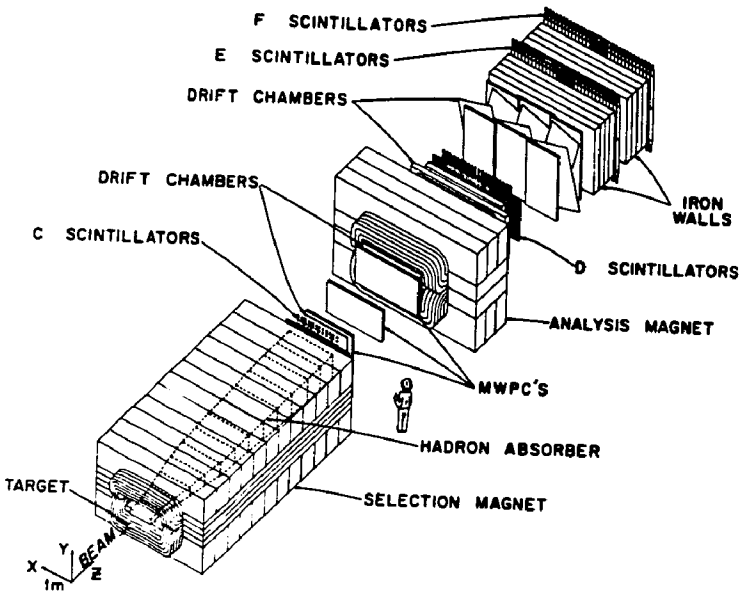

Figure 2.1: The E615 apparatus

in pairs by electromagnetic decays or interactions. Accidental pairs involving singly produced muons and muon pairs generated from secondary hadrons both tended to have low invariant mass and were eliminated by a sophisticated trigger which made an approximate determination of the dimuon mass online. An essential part of the trigger was the 'selection magnet' which was designed to give high-invariant-mass dimuons an easily identifiable signature.

After leaving the absorber, the muons were detected and measured by a series of wire chambers and scintillators. To achieve high resolution in reconstruction of the dimuon trajectories at the interaction point, the absorber was constructed to minimize multiple scattering. The 'analyzing magnet' was needed to determine the muon momenta.

We have just given a quick introduction to the main features of the E615 experiment. Each part of the apparatus will be discussed in greater detail in the remainder of this chapter.

2.1 The Beam and Target

The pion beam used for E615 was a secondary beam produced by steering a 800-GeV/c proton beam, extracted from the Fermilab Tevatron, at a beryllium production target. The secondary beam could be run with either positively or negatively charged particles. The momentum of beam particles reaching the target was restricted to a range of $\sigma_P/P = 3.4\%$. The central momentum of the incident beam for the data analysed here was 252 GeV/c [27].

The central momentum was chosen to maximize the yield of pions while keeping the proton and kaon backgrounds at acceptable levels. Using results of previous studies of the secondaries produced in proton interactions, the hadronic composition of the beam at the experimental target was calculated to be 53.8% pion, 42.1% proton, and 4.1% kaon for the π^+ beam and 92.8% pion, 1.7% proton, and 5.5% kaon for the π^- beam[24]. There were also leptons in the beam. In terms of total number of charged particles, the negative beam was 7.2% e^- and 1% μ^- and the positive beam was 2.4% e^+ and 0.5% μ^+. The lepton component does not contribute to Drell-Yan or J/ψ production, but is important in determining the pion intensity.

The secondary beam had an intensity of 2×10^8 pions per second during the spill of 20 out of 60 seconds per accelerator cycle. This high intensity was necessary to get a large sample of high-mass Drell-Yan events, but precluded tagging or measuring the momentum of the incident particles on an event-by-event basis. Identification of the incident particles was not necessary for the π^- beam, but would have been useful for the π^+ beam. The flux of the beam was monitored by ionization chambers which counted the total number of charged particles. The total number of pions on

target for the 250-GeV/c running was calculated from the measured total number of beam particles, adjusting for the non-pion contamination and the detector dead time, to be 5.5×10^{13} π^- and 3.8×10^{13} π^+.

The experimental target was a 5.1-cm-diameter, 20.3-cm-long cylinder of tungsten. In studies of the Drell-Yan process on a variety of different nuclear targets, the Drell-Yan cross section, σ_{DY}, has been observed to vary linearly with the atomic weight, A, of the nucleus, $\sigma_{DY} \sim A^{1.0}$. However, the cross section, σ_H, for inelastic hadronic interactions scales as $\sigma_H \sim A^{0.76}$ [25]. The fraction of pions producing Drell-Yan events varies as the ratio of the two cross sections, $\sigma_{DY}/\sigma_H \sim A^{0.24}$. Therefore tungsten's large atomic weight produces a favorable yield of Drell-Yan events, an increase of 3.5 relative to a hydrogen target.

There was a gap of 46 cm between the target and the absorber. This gap was made larger than the experimental z resolution of event vertex, $\sim$ 30 cm, in order to separate interactions that occurred in the target from those that occurred in the absorber. Given that we can successfully identify interactions which occur in the target, we can constraint the point of production to lie within the target. Tungsten's short interaction length for pions, 12 cm, allowed use of a small target. The target constraint on the event vertex improved the dimuon mass resolution by roughly 20%.

2.2 The Absorber

The absorber was designed to stop all particles except muons and neutrinos, while minimizing the error in the reconstruction of the muons tracks due to multiple scattering. The length required for the absorber is proportional to the hadronic absorption length, R_0, of the material used. The mean angular deflection of muons in material of given pion absorption power is $\phi_{rms} \sim \sqrt{R_0/X_0}$, where X_0 is the radiation length of the material. To minimize ϕ_{rms} one should use with low Z materials. The absorber was constructed using the two lowest-Z materials readily available. It contained 3.2 m of beryllium ($Z = 4$) and 4.1 m of graphite ($Z = 6$). Graphite was used for part of the absorber because only a limited amount of beryllium was available and was placed downstream because the detrimental effects of multiple scattering decrease nearer to the wire chambers. In total, the absorber was 15 hadronic interaction lengths (35 radiations lengths) long.

2.3 The Selection Magnet

The absorber was contained inside the gap of a 400-ton dipole magnet, the selection magnet. To maximize the magnetic field, the magnet was constructed

in sections. The gap inside each magnet section was made only large enough to contain the tracks of muons produced in the target which would pass through the wire chambers. By minimizing the gap of the magnet, one maximizes the magnetic field produced by a given magnet current and number of turns in the magnet coil, which were fixed by engineering constraints. The integral of the vertical component of the magnetic field was 10.7 Tesla·m along a line passing through the center of the magnet. The field imparted a transverse-momentum kick of 3.2 GeV/c in the horizontal plane to muons passing through the detector. The magnet swept very low-mass (less than 1 GeV/c^2) dimuons out of the detector volume and provided a means for the trigger to identify high-mass (greater than 3 GeV/c^2) dimuons.

2.4 The Trigger

The triggers used information from the scintillators. The A and B scintillators, not shown in Fig. 2.1, were used to veto muons produced by decay of particles in the beam, and the C, D, E, and F scintillators were used to do crude online event reconstruction for the trigger. The level-1 trigger required at least two hits in each bank C-E and no hits in banks A or B. Wire chamber readout began after each level-1 trigger. The level-2 trigger required that the projection of at least two of the tracks in the elevation view pointed back to the target. In the elevation view the magnets cause no bending so the extrapolation of the tracks is straightforward. When the level-2 trigger was satisfied, the scintillator patterns and wire-chamber data were stored in buffers to await the outcome of the level-3 trigger decision.

The level-3 trigger did more sophisticated processing in order to reject dimuons with an invariant mass of less than 2 GeV/c^2. The trigger operated by recognizing patterns indicative of high-mass muon pairs, rather than performing an explicit calculation of the pair's mass. To illustrate the principle, we consider a pair of muons in the bending plane of the selection magnet, each with transverse momentum of 3.2 GeV/c. The invariant mass of the pair is 6.4 GeV/c^2. If the muons are bent inwards, then the transverse momentum of each muon will be exactly balanced by the kick of the magnet and the muons will leave the magnet with parallel tracks. Lower-mass dimuons will have diverging tracks and higher-mass dimuons will have converging tracks. In practice, it was found that the angle between the muon tracks was well correlated with the dimuon mass for a wide range of asymmetric decays.

The level-3 trigger was designed to identify high-mass dimuons based on whether the tracks were diverging or converging. Using the pattern of hits in the C and D scintillators, possible pairs of tracks were identified. For each pair, three distances between the pair of hits within each scintillator bank were calculated, (two in the bending plane and one perpendicular). From the three differences a 16-bit word

was formed. A calculation of the sign of each individual track from scintillator hits provided two additional bits of information. The 18-bits were then fed into a look-up table which generated the decision to accept or reject the event based on prior calculation of the masses of the pairs which could generate that 18-bit description. The average time required for the level-3 trigger decision was 1.5 μs. Events satisfying all three triggers were written to tape. Every thousandth event satisfying the level-1 trigger was also written to tape. It was found that the level-3 trigger reduced the ratio of events with mass less than 2 GeV/c^2 to those with mass greater than 4 GeV/c^2 by a factor of 12 [24].

2.5 The Wire Chambers and Analysis Magnet

The trajectories of muons leaving the absorber were measured using multiwire proportional chambers (MWPC's) and drift chambers to make accurate position measurements. Another magnet, the analysis magnet, was located between the two sets of drift chamber so the momentum of the muons could be determined.

The drift chambers provided position measurements with a resolution of 290 μm, the MWPC resolution was 750 μm. The distance covered by the chambers upstream of the analysis magnet was about four meters. Downstream the chambers extended over five meters. So one would expect the angular resolution of the detector to be about 0.08 mrad upstream and 0.06 mrad downstream. The analysis magnet generated a field integral of 2.9 Tesla·m. For a 125-GeV/c muon this would cause a deflection of 7 mrad. So the expected momentum resolution is roughly 2% for the component of the momentum in the plane perpendicular to the magnetic field. We have somewhat underestimated the resolution because the effective position resolution for a set of chambers is better the resolution of any individual chamber. The momentum resolution found from the event reconstruction is 1% at 125 GeV/c.

Examining the observed invariant mass spectrum near J/ψ mass, we found that the experimental mass-resolution function was roughly gaussian. Fitting to the data for $x_F > 0.80$ we found a width of 178 ± 2 MeV/c^2. A Monte Carlo simulation of the experiment confirmed that the mass resolution was dominated by multiple scattering in the target and absorber.

Chapter 3

Analysis

The goal of this analysis was to obtain the cross section for the production of J/ψ as a function of x_F and P_T and to measure the angular distribution of muons from J/ψ decay. This chapter will attempt to explain how the physical quantities of interest were extracted from the measurements made with the apparatus. First, we will describe the reconstruction of muon trajectories in the detector and the calculation of dimuon kinematics at the interaction point. Second, we will describe the method used to separate the J/ψ signal from events produced by continuum interactions and the limitations on the measurements that could be made with our data sample. We will then discuss how to correct for the acceptance of the detector and determine the true cross section from the measured number of J/ψ.

3.1 Event Reconstruction and Selection

The general event-reconstruction procedure common to all analysis done on E615 data has been described in detail in the other theses from E615 [13,15,17,26,27].

The first step of the event reconstruction used the information from the scintillator banks to do a fast calculation of the mass of the muon-antimuon pair with a mass resolution of about 250 MeV/c^2. The calculation was very quick and was used to eliminate events with mass less than 2 GeV/c^2 early in the analysis. The calculation was essentially a repeat of the online level-3 trigger, but was more accurate because it used a better algorithm and information from more scintillator banks. About 40% of the data passed this cut.

The analysis routines then turned to the data from the wire-chambers. The first operation was to cut events with too many hits in the chambers. These cuts eliminated readout errors, events with showers in the chambers, and most events with more than one pair of muons. For events passing these cuts, track finding in the chambers began starting from the hits in the scintillators. The scintillator hits were used as the starting point for the track-finding algorithm, but were not used

in the track fits.

Upstream tracks were rejected if the fit probability was less than 2% or the number of chambers hits in the track was less than 5 out of the 9 possible. Downstream line segments were required to have a fit probability of better than 5%. If more than 50 msec of Cyber 175 computer time was used for an event, track fitting was stopped. Using the data in the downstream drift chambers, an additional fit was performed to find the time at which the particle passed through the chamber. The beam at Fermilab delivered particles in 1 ns 'RF buckets' separated by 18.6 ns intervals. The chamber timing had a resolution of 2.3 nsec. A cut was placed at 11.5 nsec, relative to the scintillator timing, to eliminate tracks from other RF buckets.

For this analysis, we accepted only events with one μ^+ track and one μ^- track. This removed any possible ambiguity due to mispaired muons. In order to eliminate halo muons, we also required that muons with the same sign of charge as the beam have angles relative to the beam of greater than 5 mrad and momenta less than 200 GeV/c.

A global fit was performed using pairs of upstream and downstream line segments. This fit took into account deflection of tracks in the fringe fields of the magnets. The pairing of an upstream and downstream segment in to a muon track was accepted if the fit probability was better than 3%. The resolutions from this fit were $\sigma(P_{zz})/P_{zz} = 0.0001 \cdot P_{zz}$ for P in GeV/c, $\sigma(\theta_{zz}) = 0.1$ mrad, and $\sigma(\theta_{zy}) = 0.3$ mrad [15]. As a check on the consistency of the fitting procedure the difference between the position of a hit within a chamber and the fitted position at that chamber were plotted for each chamber. The average difference for each chamber is zero to within 20% of the chamber resolution and there are no systematic shifts from chamber to chamber[13].

After the track fitting was done the next task was to simulate the transport of the muons through the absorber to find the muon momenta (magnitude and angle) at the point of production (the 'vertex'). Each muon pair had ten measured quantities. At a given plane perpendicular to the beamline and downstream of the absorber, the five parameters for each muon are: the two slopes of trajectory, the two intercepts of the trajectory, and magnitude of the muon momentum. The quantities of more direct physical interest are the muon trajectories at the point of production. There are nine parameters: the three coordinates of the vertex, and two slopes and a magnitude for each muon's momentum. The z position of the vertex was fixed at the mean z-coordinate of hadronic production in the target, because the vertex z-resolution was larger than the size of the target.

Six of the muon parameters at the vertex, not including the magnitudes of the momenta, were calculated by fitting the ten measured muon parameters using a

transport equation. The transport equation was an approximation to the average of many muon trajectories through our Monte Carlo simulation of the experiment. Some of the parameters of the transport equation (and the simulation) could be calculated from the known properties of the absorber. The additional parameters were found by minimizing the observed width of the J/ψ resonance [17]. If the probability that the muon pair had been produced a the same vertex, as calculated from the transport equation fit, was less than 2% the event was rejected.

The magnitudes of the muon momenta at the interaction point were calculated from the measured momenta taking into account the energy loss in the absorber. The energy loss was found to be nearly constant over the momenta range of interest with a small correction linear in the momentum [17,27].

With the kinematic parameters of the muons at the vertex, we calculated the kinematic variables of the muons in their rest frame at their point of their production. The kinematic variables used in the analysis were calculated from the momentum three vectors of the two muons, as discussed in section 1.1.

3.2 Extraction of the J/ψ Signal

In the previous section we have shown have the kinematic variables of the muons at the interaction point were determined. The data can be considered to be a list of the x_F, P_T, $\cos\theta^*$, ϕ^*, and M of each dimuon event observed in the detector. To find the dependence of the J/ψ cross section on the kinematic variables other than M, we bin the data in the variables and then find the number of J/ψ in each bin by performing a fit to the number of events versus mass within the bin. The dimuons from J/ψ decay form a peak centered near the J/ψ mass with a width equal to the mass resolution of the detector. The number of J/ψ in each bin is given by the number of events contained in the peak. We must have enough events in each bin to extract the J/ψ signal, so we are limited in how finely we can bin the data and how many variables we can examine simultaneously.

The goal of the analysis was to measure the cross section for J/ψ production as a function of x_F and P_T, and the angular distribution in $\cos\theta^*$ and ϕ^* of muons from J/ψ decay as a function of x_F. We split the analysis into separate measurements of the production cross section and the angular distribution because we did not have sufficient data at high x_F to make acceptable mass fits when the data were binned in the four variables simultaneously. The x_F-P_T dependence of the cross section was measured by dividing the data into 90 regions: 15 bins of x_F by 6 bins of P_T. The angular distribution measurement used 375 regions: 15 bins of x_F by 5 bins of $\cos\theta^*$ by 5 bins of ϕ^*. The x_F bins were evenly spaced over the range $.25 < x_F < 1.0$. The $\cos\theta^*$ and ϕ^* bins evenly covered their physical ranges. The P_T bins extended

from $P_T = 0$ to $P_T = 5.0$ GeV/c, increasing in size at large P_T because the number of events dropped off as P_T increased.

Figure 3.1 shows the number of events versus mass along with the fitted function (solid line) and the fit to the background (dotted line) for two x_F-P_T bins. The function used for the fit was the sum of two gaussians for the J/ψ and ψ′ peaks and a quadratic polynomial plus the exponential of a linear function for the continuum. The functional form chosen for the continuum is somewhat arbitrary and was used mainly because it provided good results. The continuum events arise primarily from Drell-Yan pair production which is well characterized by an exponential dependance on mass. At the low end of the mass range there is a sharp drop-off in the number of events due to the effect of the level-3 trigger. The sum of an exponential and a quadratic polynomial provided a good fit to both the low-mass cutoff and the high-mass exponential form. The precise form of the continuum function affected the quality of the fits, but did not affect the fitted number of J/ψ. Similarly, inclusion of the a peak for the ψ′ improved the fits but had little effect on the number of J/ψ. This was expected since the number of ψ′ observed is much smaller than the number of J/ψ. The function used for the fit is given below.

$$\begin{aligned} f(M) &= N_{J/\psi} e^{-(M-M_{J/\psi})^2/2\sigma^2} \\ &\quad + N_{\psi'} e^{-(M-M_{\psi'})^2/2\sigma^2} \\ &\quad + aM^2 + bM + c + e^{\alpha M+\beta} \end{aligned}$$

There were seven free parameters in the function $f(M)$ used to fit the mass distribution: the number of J/ψ ($N_{J/\psi}$), the number of ψ′ ($N_{\psi'}$), and five parameters to fit the continuum (a, b, c, α, β). The function had three fixed parameters used to describe the resonance peaks: the mass of the J/ψ ($M_{J/\psi} = 3.097$ GeV/c^2), the mass of the ψ′ ($M_{\psi'} = 3.685$ GeV/c^2), and the width of the resonances ($\sigma = 180$ MeV/c^2). The width is the experimental resolution rather than the intrinsic width of the particles. We did not allow the central mass and width of the J/ψ and ψ′ peaks to vary because with the addition of these two parameters the fitting function is ill behaved. The J/ψ gaussian and the continuum form are not orthogonal functions so for fits on plots with low statistics, or with a poor choice of initial parameters, both functions attempt to fit the region of highest statistics, the J/ψ peak, causing the fit parameters to diverge. Since we could not allow the width and mass to vary at high x_F, we chose to be consistent and not allow them to vary at any x_F.

To investigate what effect this might have on our determination of the number of J/ψ we performed fits with the masses and the width left as free parameters. These fits were on bins of only x_F, in order to have sufficient statistics over the entire

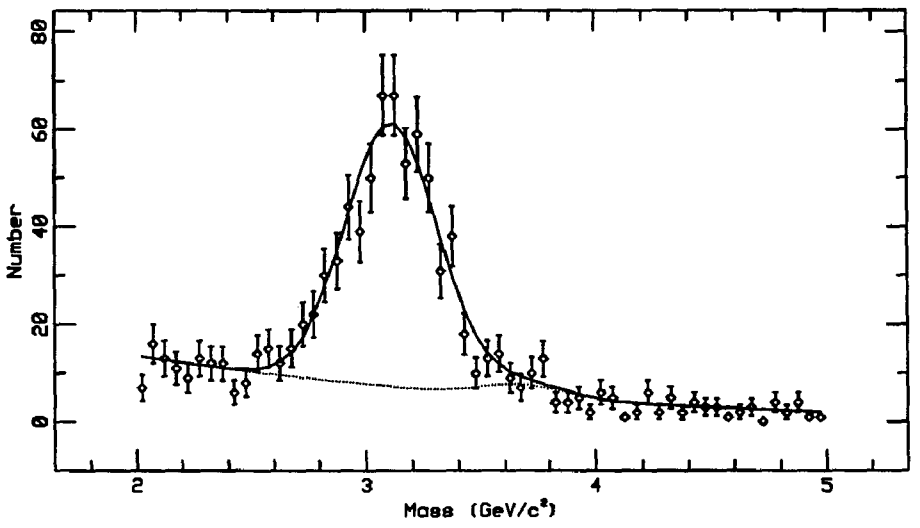

$0.85 < x_F < 0.90$, $2.0 < P_T < 3.0$

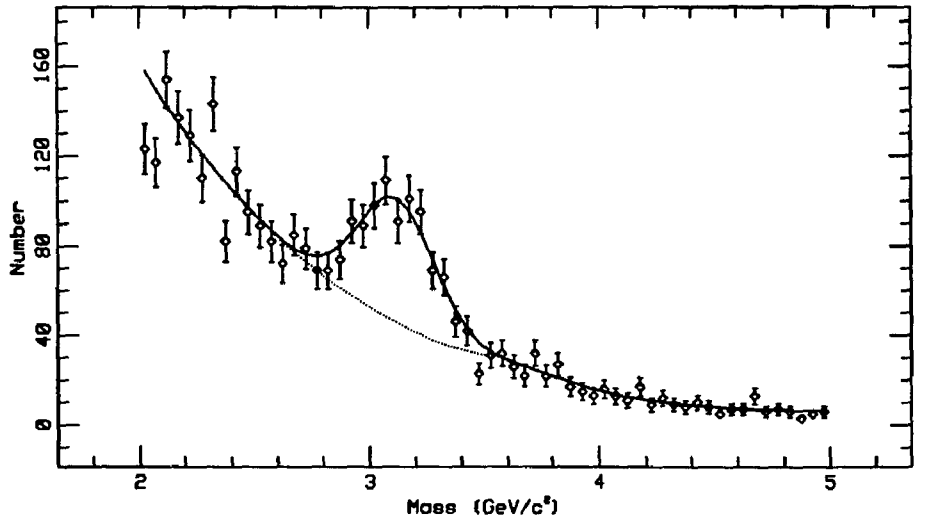

$0.95 < x_F < 1.0$, $0.5 < P_T < 1.0$

Figure 3.1: Plots of the raw mass distribution. Shown in the plots are the data points, the overall fit to the data (solid line), and the fit to the background (dotted line).

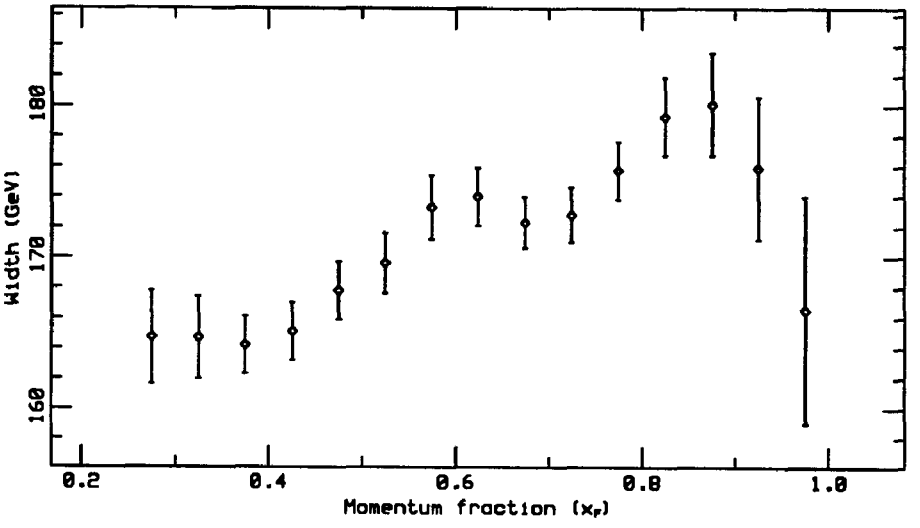

Figure 3.2: Width of the J/ψ peak

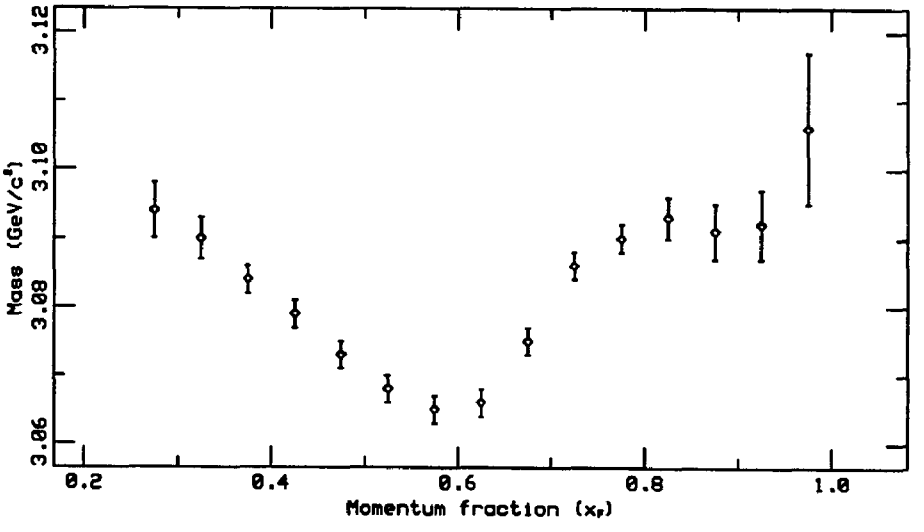

Figure 3.3: Centroid of the J/ψ peak

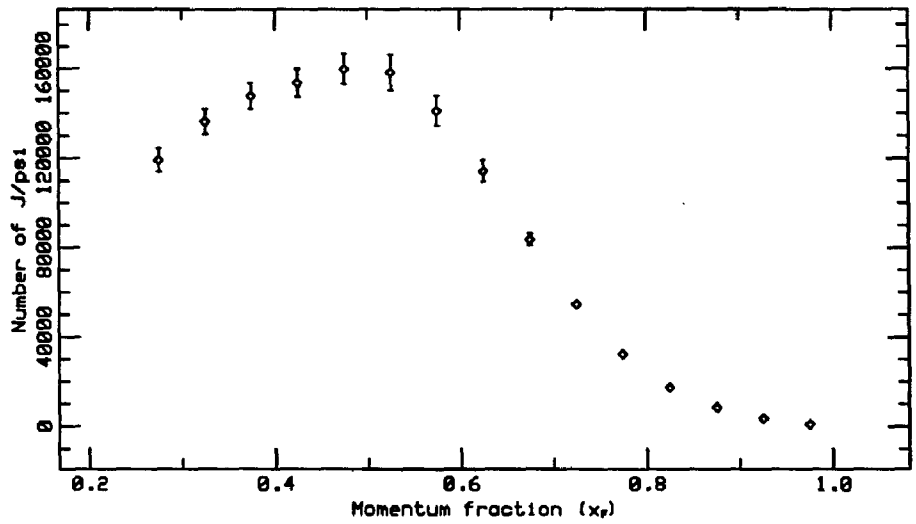

Figure 3.4: Raw number of J/ψ

range of x_F. Figure 3.2 shows the measured width of the J/ψ as a function of x_F. Figure 3.3 is a similar plot for the centroid of the J/ψ peak. Since the uncertainty in the momentum is proportional to the momentum, we would expect the width increase with x_F. This is the case except for the two highest x_F bins. The reason for the shift in the J/ψ centroid is most likely due to inaccuracy in the background subtraction. These variations in the width and centroid of the J/ψ peak had little effect on the number of J/ψ's determined from the fit. The ratio of the number of J/ψ from the fits with constant mass and width to the number of J/ψ from the fits where those two parameters was unity to within 2% except for the highest x_F bin where there was a 7% difference. The difference in the number of J/ψ from the two fitting procedures is significantly smaller than our estimated uncertainty for either fitting procedure across the entire range of x_F.

To account for the uncertainty in the width determination and for errors due to deviations from gaussian in the J/ψ peak, we added a 2.5% error in quadrature to the error in the number of J/ψ's given by the fitting procedure. In cases where the χ^2 was greater than the number of degrees of freedom of the fit (DoF), we multiplied the error in the number of J/ψ's by $\sqrt{\chi^2/\mathrm{DoF}}$, essentially expanding our error bars until the $\chi^2/\mathrm{DoF} = 1$. The observed number of J/ψ versus x_F is plotted

in figure 3.4.

3.3 Correcting for Acceptance

The number of J/ψ events in each bin of the x_F-P_T or the x_F-$\cos\theta^*$-ϕ^* distribution does not directly give the underlying physical cross section. To find the cross section for J/ψ production, or the true decay angular distribution, we must know the probability that a muon pair with a given set of kinematic variables will be detected in the apparatus (the acceptance). A Monte Carlo simulation of the detector was constructed to provide a way to calculate the acceptance. We did not have enough events to bin the data in all the of kinematic variables simultaneously. Therefore, when calculating the acceptance for a particular bin of some of the variables, we must always average the cross section over at least one other variable. However, to do the average we must already know the cross section.

We began with a first estimate of the acceptance and by iteratively improving our values for the cross section, we converged to the correct acceptance and cross section. For our first pass, we used the acceptance calculated from an analysis of Drell-Yan data from E615. We applied this acceptance to the observed J/ψ cross section to calculate a physical cross section. This cross section was used to generate the first set of Monte Carlo events and new values for the acceptance. The new acceptance, together with the data, were used to make an improved calculation of the physical cross section. The new cross section was then used to generate more Monte Carlo events. The process was repeated until the cross section that went into the simulation to generate the acceptances and the distribution calculated from acceptances and the data agreed within the errors of the data.

We actually performed two series of iterations. First we assumed that the angular distribution was flat, *i.e.*, the cross section had no dependence on $\cos\theta^*$ or ϕ^*, and then iterated using the procedure described above, modifying the dependence of the cross section only on x_F and P_T, and using data binned only in x_F and P_T. The parameterization used for the x_F-P_T dependence of the cross section is given in Eq. (4.2). This first stage of iteration was performed because we knew from previous experiments that the angular distribution is flat over most of the range of x_F and P_T. The assumption of a flat distribution greatly reduced the number of Monte Carlo events required to calculate the acceptance. Most of the computer time for Monte Carlo event generation was spent iterating until this approximation to the x_F-P_T distribution converged.

After the x_F-P_T distribution converged we used it to find acceptances for the angular variables, which were used to find the angular distribution. We then changed the cross section in the simulation to reflect the new angular distribution as a func-

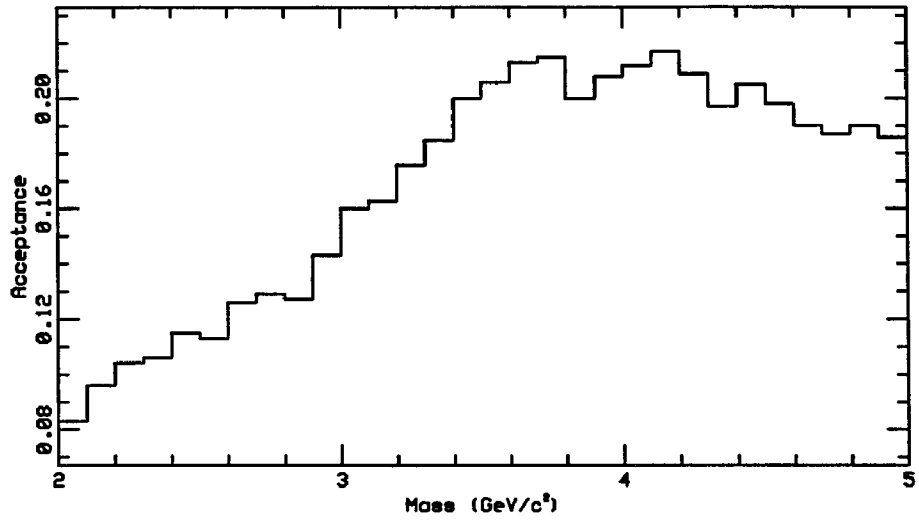

Figure 3.5: Acceptance versus mass

tion of x_F calculated from the acceptances and the data binned in x_F, $\cos\theta^*$, and ϕ^*. With the new angular part of the cross section, we generated events to find a new x_F-P_T acceptance, which we used to recalculate the x_F and P_T dependence of the cross section. We alternated finding the x_F, P_T acceptance and the x_F-$\cos\theta^*$-ϕ^* acceptance until the fits for both the J/ψ production cross section and the angular distribution converged. Fortunately, only a few repetitions were necessary. The parameterization used for the angular part of the cross section is given in Eq. (4.8). The parameterization has no dependence on ϕ^* because we found no significant variation of the cross section with ϕ^*.

Figures 3.5 and 3.6 show the Monte Carlo generated acceptances as a function of each of the five kinematic variables. For each plot the other variables were integrated over according to our best fit of the J/ψ production distribution, which will be presented in the next section. The integration was over the entire physical range of the variable, except when integrating over x_F which we restricted to $0.75 < x_F < 1.0$, since we are interested primarily in the high x_F region.

Another problem in calculating the acceptance would arise if the experimental resolution is comparable to the bin size. In this case, the acceptances are a function of the production distribution and we must know the cross section in order to calculate the contribution to the signal in a given bin from muon pairs over the range

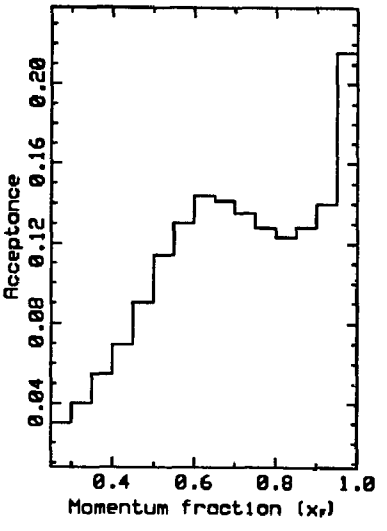
Acceptance
0.04
0.08
0.12
0.16
0.20
0.4
0.6
0.8
1.0
Momentum fraction (x_F)

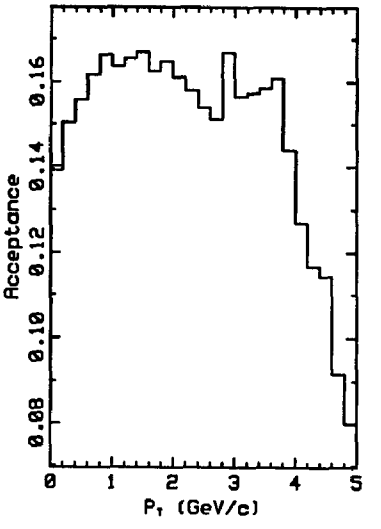
Acceptance
0.08
0.10
0.12
0.14
0.16
0
1
2
3
4
5
P_T (GeV/c)

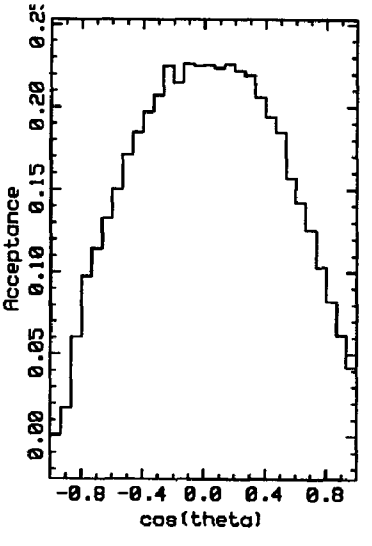
Acceptance
0.00
0.05
0.10
0.15
0.20
-0.8
-0.4
0.0
0.4
0.8
cos(theta)

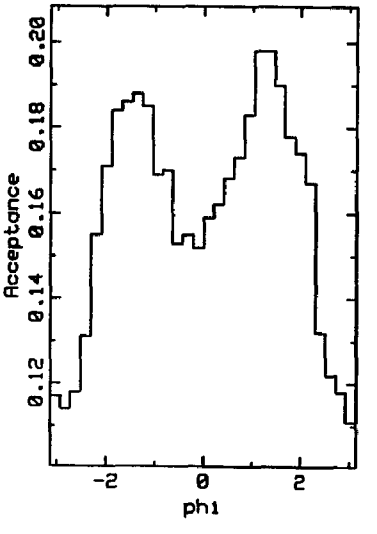
Acceptance
0.12
0.14
0.16
0.18
0.20
-2
0
2
phi

Variable	Resolution
x_F	0.03
P_T	190 MeV/c
M	180 MeV/c^2
$\cos\theta^*$	0.08
ϕ^*	0.28 radians

Table 3.1: Resolutions for the kinematic variables

of the experimental resolution. As long as the magnitude of the cross section varies slowly relative to the bin size, this problem could be handled using the same iterative procedure described above. We avoid this difficulty because the experimental resolution, given below, is smaller than the bin sizes we have selected except in the case of the invariant mass. The mass fitting procedure described above allows for a resolution larger than the bin size.

It is important to note that our simulation of the experiment generated complete simulated events equivalent to the data produced by the apparatus. The simulated events were processed using the same reconstruction procedures used for the data. The simulation takes into account the inefficiencies and resolution of the apparatus. The acceptance was calculated using the kinematic variables of generation in the generated distribution and the reconstructed kinematic variables in the accepted distribution. The acceptance calculation, therefore, takes into account smearing of the cross section near the edges of kinematic-variable bins.

3.4 Resolution

An estimate of the resolution of the detector in the five kinematic variables in shown in Table 3.1. The resolutions were determined using the Monte Carlo simulation. In each case events were produced with the production cross section for $x_F > 0.75$ described below, with one variable held fixed. The root-mean-square deviation from the mean of the reconstructed distribution for the variable held fixed was taken as the resolution for that variable. The invariant-mass resolution agrees with the value found from a fit to the observed width of the J/ψ peak.

3.5 Beam Composition and Cross-section Normalization

To reduce the contribution to J/ψ produced by protons in the π^+ beam, we accepted only events with $x_F > 0.75$ from the positive beam. The number of events produced at high x_F by protons is greatly reduced relative to the yield from pions, because the structure function of the proton goes to zero more rapidly at high x_F than the structure function of the pion. Examination of the relative number of J/ψ produced by the positive and negative beams of E615 as a function of x_F shows that the ratio is constant for $x_F > 0.75$ [27]. We take this as evidence that the proton contribution to J/ψ production above $x_F = 0.75$ is negligible.

The cross section of a process can be calculated if the number of events produced by the process, the number of incident particles, and the properties of the target nucleons are known. The cross section is given by

$$\sigma = \frac{N\sigma_0}{N_\pi f}\left[1 - e^{-L/R_0}\right]^{-1} \tag{3.1}$$

where σ is the cross section of interest, N is the number of events corrected for acceptance, f is the fraction of time during which the experiment was able to accept events (the live time), σ_0 is the pion absorption cross section per nucleon, N_π is the number incident pions, L is the length of the target, and R_0 is the pion absorption length of the target. This formula is derived assuming the the pion beam is exponentially attenuated as it passes through the target. The main part of this chapter has been devoted to describing how N is determined. The live time can determined from information available from the data-acquisition system, the length of the target can be easily measured, and the pion absorption of tungsten is well measured [64].

The main problem in measuring the absolute scale of the cross section is determining the number of pions incident on the target. The intensity of the beam was monitored during each spill by an ionization chamber which recorded the number of charged particles in the beam. The chambers were calibrated by inserting a copper foil in the beamline and later measuring the radioactivity from Na^{24} nuclei produced by collisions of the beam particles with Cu nuclei. The systematic uncertainty in the calibration is on the order of 20%. For this analysis we use a determination of the cross-section normalization made by J. Conway for the π^- beam [15]. The total cross section for the π^+ beam for $x_F > 0.75$ was normalized to match the total cross section for the π^- over the same interval.

Chapter 4

Results

In this chapter we present the cross-section measurements made during our analysis. We present results for the x_F and P_T dependence of the J/ψ production cross section and for the angular distribution of the decay muons. We also present the results of fits of the data to analytic approximations of the cross sections and compare the data to the $q\bar{q}$ annihilation model and the parameterization of the dependence of the angular distribution on x_F presented in section 1.4.

4.1 The x_F-P_T Cross section

The measured cross section for each of the x_F-P_T bins described in the previous chapter is given in Table 4.1. The values $\sigma_{i,j}$ are the cross section (in picobarns) for J/ψ production integrated across the x_F-P_T bin multiplied by the branching ratio $B_{\mu^+\mu^-}$ of J/$\psi \rightarrow \mu^+\mu^-$.

$$\sigma_{i,j} = B_{\mu^+\mu^-} \int dx_F \int dP_T \, \frac{d\sigma_{J/\psi}}{dx_F \, dP_T} \tag{4.1}$$

where the integration extends over a single data bin. The values given are the cross section per nucleon, calculated from the cross section on tungsten assuming that the cross section varies as $\sigma \sim A^\delta$, where A is the atomic mass of the target. Due to the disagreement between the various measurements of the x_F dependence of δ, (see section 1.5), we chose to ignore the issue completely and take $\delta = 1$ for all x_F and all P_T. The error quoted for the cross sections is due to statistical error in the J/ψ mass fit and the acceptance calculation and disregards the overall systematic uncertainty of 20% in the beam normalization. We do not present results for bins with fewer than 20 J/ψ events.

Our cross-section measurement depends on the acceptance calculated from the Monte Carlo simulation of the experiment, and on the parameterization of the cross section used to generate events in the simulation. As discussed above, we generated

	P_T in GeV/c					
x_F	0.25	0.7	1.25	1.75	2.5	4.0
0.275	125±10	259±19	204±15	126±10	88±9	10±3
0.325	119±8	219±17	163±11	104±7	61±5	13±3
0.375	93±6	178±13	136± 9	75±5	53±4	7.8±1.7
0.425	84±6	148±11	109±8	60±4	41±3	4.3±0.8
0.475	64±5	122±9	92±7	44±3	29±2	6.8±1.7
0.525	54±4	95±8	72±5	38±2	21.3±1.3	4.4±0.6
0.575	44±3	73±6	54±4	29.5±1.8	18.0±1.2	2.3±0.3
0.625	32±2	56±4	38±2	19.8±1.2	11.0±0.7	1.9±0.3
0.675	24.6±1.4	41±2	26.6±1.6	14.0±0.8	7.5±0.5	1.3±0.2
0.725	18.9±1.1	28.5±1.3	17.7±0.8	9.2±0.5	4.9±0.4	0.63±0.13
0.775	12.7±0.6	18.7±0.7	11.5±0.4	5.1±0.2	3.19±0.17	0.28±0.04
0.825	7.6±0.4	10.6±0.4	6.2±0.3	2.73±0.17	1.32±0.11	0.20±0.04
0.875	3.72±0.24	5.4±0.3	2.78±0.16	1.19±0.09	0.58±0.07	0.05±0.02
0.925	1.65±0.16	2.10±0.14	1.06±0.09	0.32±0.04	0.15±0.03	-
0.975	0.29±0.05	0.39±0.05	0.13±0.02	0.058±0.015	0.014±0.005	-

Table 4.1: The J/ψ cross section versus x_F-P_T. The values are the cross section times branching ratio per nucleon in picobarns integrated across each data bin. The x_F and P_T values are the bin centers. All x_F bins have a width of 0.05. The four lower P_T bins have a width of 0.5 GeV/c. The upper two P_T bins are 2.0 GeV/c < P_T < 3.0 GeV/c and 3.0 GeV/c < P_T < 5.0 GeV/c.

several sets of simulated events, varying the cross section used in the simulation until it agreed with the cross section calculated from the experimental data and the acceptances generated using the simulation. To parameterize the J/ψ production cross section, we used the form

$$\frac{d\sigma}{dx_F\, dP_T} = A\frac{(1-x_F)^C}{[1+(P_T/\,\mathrm{GeV}/c)^B/S]^T} \tag{4.2}$$

where

$$S = D + E(1-x_F)^2 \tag{4.3}$$
$$T = F + G(1-x_F)^2 \tag{4.4}$$

This form was taken from a previous analysis of J/ψ data from E615 [13]. It is not physically motivated, but has been found to provide a good fit to the data. The form has seven free parameters, including the normalization. The normalization, A, has units of nanobarns/GeV/c. The normalization values are for the differential cross section multiplied by $B_{\mu^+\mu^-}$. With our choice of GeV/c as a scale for P_T, all of the other parameters are dimensionless. To provide sufficient statistics across the entire range of x_F while limiting the total number of events required, we calculated the acceptance and fit the cross section in three regions of x_F. The final parameter values for the cross-section fits in each region are shown in Table 4.2. The rightmost column of the table is the result of a fit for the cross section over the entire range $0.25 < x_F < 1.00$. The error quoted for the normalization (parameter A) includes the uncertainty due to the beam normalization.

Figure 4.1 shows our measurement of the differential J/ψ production cross section, $d\sigma_{J/\psi}/dx_F$, versus x_F multiplied by $B_{\mu^+\mu^-}$. The data from this analysis are shown as diamonds; the filled boxes are data from the NA3 collaboration [22]; the solid line is an empirical fit calculated by numerically integrating the cross-section parameterization of Eq. (4.2) using the parameters from the overall fit of Table 4.2; and the dashed line is the cross section calculated from a $q\bar{q}$ annihilation model for J/ψ production. The position in x_F of each data point is the average x_F of the events used for the data point. Weighting the x_F of the data produces a noticeable shift from the center of the bin only at very high x_F where the cross section changes very rapidly. The weighted average values of x_F for the three highest-x_F bins are 0.872, 0.920, 0.964.

Our data are in good agreement with the measurements of the NA3 collaboration. The NA3 data were taken with a 280-GeV/c pion beam on platinum nuclei ($A = 195.1$), and so are directly comparable to the data presented here. If we assume the cross section scales as A^δ with the worst case value of $\delta = 0.77$ (see

	$0.25 < x_F < 0.50$	$0.50 < x_F < 0.75$	$0.75 < x_F < 1.0$	$0.25 < x_F < 1.0$
A(nb/ GeV/c)	94 ± 20	89 ± 18	64 ± 13	87 ± 17
B	1.61 ± 0.03	1.53 ± 0.03	1.45 ± 0.03	1.53 ± 0.02
C	2.13 ± 0.19	1.88 ± 0.04	1.70 ± 0.05	1.88 ± 0.02
D	3.4 ± 0.4	4.8 ± 0.2	5.5 ± 0.1	5.2 ± 0.3
E	40 ± 2	55 ± 1	162 ± 3	67 ± 9
F	7.0 ± 0.6	8.8 ± 0.1	12.5 ± 0.2	10.3 ± 0.5
G	30 ± 1	50 ± 1	149 ± 10	56 ± 8

Table 4.2: Results of the x_F-P_T cross-section fits

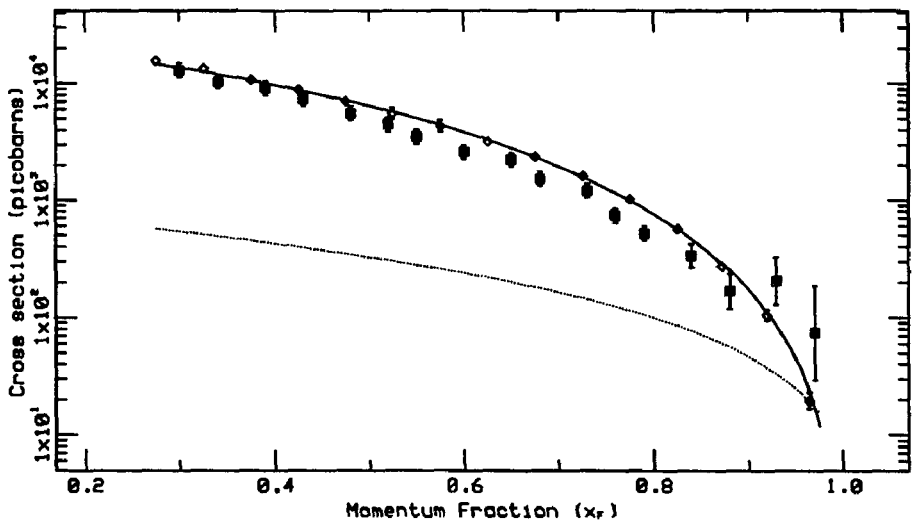

Figure 4.1: $B_{\mu^+\mu^-} \cdot d\sigma_{J/\psi}/dx_F$ versus x_F. Shown in the plot are the data from this analysis (diamonds), data from the NA3 collaboration (filled boxes), an empirical fit to the data from this experiment (solid line), and the cross section calculated from a $q\bar{q}$-annihilation model (dotted line).

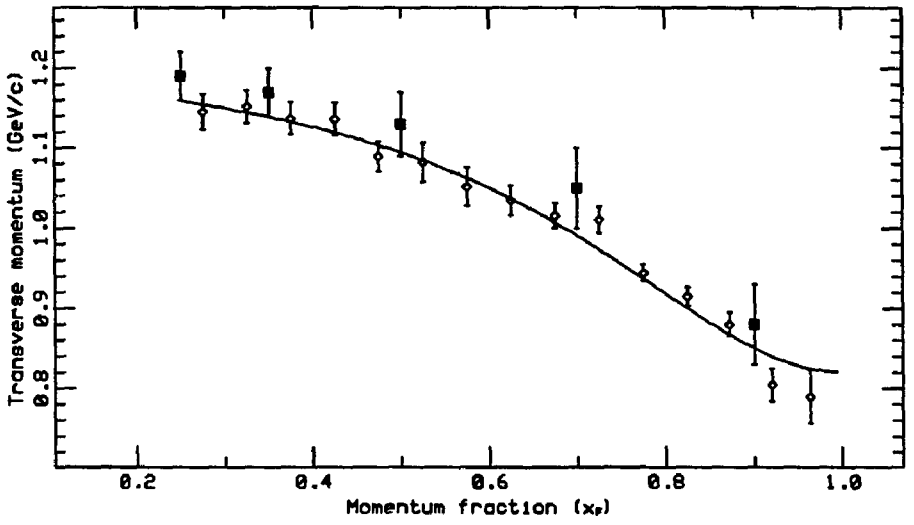

Figure 4.2: Average P_T versus x_F. Shown in the plot are the data from this experiment (diamonds), data from the NA3 collaboration (filled boxes), and an empirical fit to the data from this experiment (solid line).

section 1.5) then the maximum possible relative error in the normalization of the nuclear cross sections to nucleon cross section is 4%. This is well below the uncertainty in our beam normalization. The only significant discrepancy between the two sets of measurements is at very high x_F, where the better acceptance and data sample of E615 should make our measurement more reliable than the NA3 result.

The dotted line in Fig. 4.1 is the result of a calculation of the $q\bar{q}$-annihilation cross section using the resonance-production model and hadron structure functions, Eqs (1.4-1.6) described in section 1.4. We have matched the $q\bar{q}$ cross section to our measured cross section at the highest x_F data point. This leads to a strong-interaction initial-state correction factor $\kappa = 2.1$, (see section 1.3 for discussion of κ), which is in agreement with the values of κ obtained from Drell-Yan data [17,15]. We take this as an indication that a substantial fraction of J/ψ production at high x_F is due to $q\bar{q}$ annihilation.

Figure 4.2 shows our measurement of the average P_T of the J/ψ as a function of x_F. The diamonds are data from this analysis, the filled boxes are results from NA3 [22], and the curve is an emperical fit calculated by numerically integrating the cross-section parameterization of Eq. (4.2) using the parameters from the overall

cross-section fit of Table 4.2. The average P_T was calculated using a weighted average of P_T for the six P_T bins at a each x_F.

$$\langle P_T \rangle_i = \frac{\sum_j P_{T_{i,j}} \sigma_{i,j}}{\sum_j \sigma_{i,j}} \tag{4.5}$$

where i is the index over x_F, j is an index over P_T, $\sigma_{i,j}$ was defined in Eq. (4.1), and $P_{T_{i,j}}$ is the mean P_T within bin i,j. We calculated $P_{T_{i,j}}$ for each x_F-P_T bin by numerically integrating the cross-section parameterization of Eq. (4.2), using the parameters from the fit in the appropriate region of x_F in Table 4.2, over the bin. There is good agreement between our data and the NA3 results over the entire range of x_F. The average P_T decreases with x_F, but is nonzero in the limit $x_F \rightarrow 1$.

4.2 The Angular Dependence

The primary goal of this analysis was to determine the angular distribution of muons from J/ψ decays. We measured the cross section for J/ψ production and decay to muons as a function of x_F, $\cos\theta^*$, and ϕ^*. The angular variables are measured relative to the incident pion momentum in the t-channel, as described in section 1.1.

The cross-section measurements for $x_F > 0.75$ are shown in Table 4.3. The values $\sigma_{i,j,k}$ are the cross section in femtobarns integrated across the data bin multiplied by the branching ratio, (this is analogous to Eq. (4.1) defined for the x_F-P_T cross section).

$$\sigma_{i,j,k} = B_{\mu^+\mu^-} \int dx_F \int d\cos\theta^* \int d\phi^* \, \frac{d\sigma_{J/\psi}}{dx_F \, d\cos\theta^* \, d\phi^*} \tag{4.6}$$

where the integration extends over a single data bin. Results are not presented for bins with fewer than 15 J/ψ events. The Monte Carlo events for the acceptance calculation were generated according to an angular distribution which was uniform in ϕ^* but not uniform in $\cos\theta^*$; this will be discussed further. As in the case of the x_F-P_T cross sections, the estimated error disregards the overall systematic uncertainty of 20% in the beam normalization.

The rightmost column in Table 4.3 shows the cross section calculated using sums over ϕ^* for each x_F-$\cos\theta^*$ bin of the number of J/ψ events and of the number of accepted and generated events used in the acceptance calculation. All x_F-$\cos\theta^*$-ϕ^* bins, including those suppressed in Table 4.3, were used in the summation over ϕ^*.

To parameterize the cross sections, we used the most general form possible for the parity-conserving decay of a spin-1 particle, shown in Eq. (4.7). The form has three free parameters (neglecting the normalization). The parameter α describes

			$0.75 < x_F < 0.80$			
$\cos\theta^*$	$\phi^* = -0.8\pi$	-0.2π	0.0	0.4π	0.8π	Total
−0.8	4150 ± 1080	2290 ± 170	2080 ± 130	2210 ± 170	2080 ± 480	11090 ± 340
−0.4	2760 ± 240	2160 ± 120	2010 ± 120	2090 ± 110	2880 ± 230	11000 ± 250
0.0	2180 ± 140	2030 ± 120	2040 ± 120	1970 ± 120	2180 ± 130	10340 ± 230
0.4	1970 ± 110	1940 ± 120	980 ± 180	2020 ± 110	2190 ± 120	10140 ± 230
0.8	1860 ± 120	1740 ± 200	3420 ± 1260	2200 ± 220	1910 ± 120	9600 ± 290
			$0.80 < x_F < 0.85$			
$\cos\theta^*$	$\phi^* = -0.8\pi$	-0.2π	0.0	0.4π	0.8π	Total
−0.8	2070 ± 620	1260 ± 130	1250 ± 110	1240 ± 120	1610 ± 480	6420 ± 250
−0.4	1590 ± 190	1120 ± 80	1100 ± 80	1160 ± 80	1690 ± 170	6050 ± 180
0.0	1320 ± 100	1020 ± 80	1130 ± 90	1090 ± 80	1190 ± 90	5680 ± 160
0.4	1150 ± 90	1000 ± 70	1140 ± 130	1180 ± 90	1120 ± 90	5560 ± 160
0.8	1000 ± 80	820 ± 110	–	950 ± 120	1210 ± 100	5180 ± 200
			$0.85 < x_F < 0.90$			
$\cos\theta^*$	$\phi^* = -0.8\pi$	-0.2π	0.0	0.4π	0.8π	Total
−0.8	2400 ± 1200	770 ± 120	590 ± 70	540 ± 90	2120 ± 1010	3310 ± 190
−0.4	770 ± 110	510 ± 60	480 ± 50	710 ± 70	910 ± 130	3040 ± 120
0.0	560 ± 70	560 ± 60	590 ± 60	580 ± 60	630 ± 70	2900 ± 110
0.4	510 ± 60	490 ± 50	510 ± 70	530 ± 60	510 ± 50	2540 ± 100
0.8	480 ± 60	360 ± 70	–	540 ± 90	540 ± 70	2490 ± 140
			$0.90 < x_F < 0.95$			
$\cos\theta^*$	$\phi^* = -0.8\pi$	-0.2π	0.0	0.4π	0.8π	Total
−0.8	–	200 ± 40	210 ± 30	280 ± 50	–	1190 ± 100
−0.4	320 ± 60	230 ± 30	260 ± 30	220 ± 30	270 ± 50	1240 ± 80
0.0	230 ± 30	190 ± 30	210 ± 30	210 ± 30	250 ± 40	1060 ± 60
0.4	240 ± 30	170 ± 30	180 ± 40	220 ± 30	220 ± 30	1040 ± 60
0.8	160 ± 30	220 ± 50	–	240 ± 50	140 ± 30	890 ± 80
			$0.95 < x_F < 1.00$			
$\cos\theta^*$	$\phi^* = -0.8\pi$	-0.2π	0.0	0.4π	0.8π	Total
−0.8	–	31 ± 14	30 ± 11	51 ± 14	–	195 ± 34
−0.4	36 ± 16	39 ± 8	43 ± 10	41 ± 11	49 ± 19	208 ± 24
0.0	38 ± 12	43 ± 8	39 ± 10	40 ± 8	53 ± 16	214 ± 22
0.4	27 ± 9	30 ± 12	35 ± 18	29 ± 9	42 ± 10	164 ± 23
0.8	10 ± 9	19 ± 14	–	–	25 ± 9	87 ± 23

Table 4.3: The J/ψ cross section versus x_F-$\cos\theta^*$-ϕ^*. The values are the cross section times branching ratio per nucleon in femtobarns integrated across the data bin. The x_F, $\cos\theta^*$, and ϕ^* values shown are the bin centers. The bins are uniformly spaced in each of the three variables.

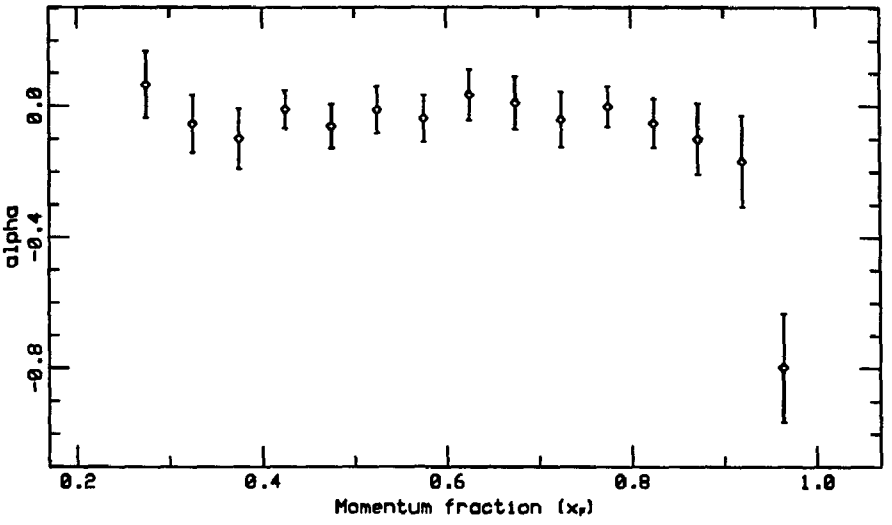

Figure 4.3: The angular distribution parameter α versus x_F

the alignment of the μ^+ momentum with the momentum of the incident pion. From the angular distribution, we can infer the alignment of the spin of the J/ψ with the momentum of the incident pion. We will describe the alignment of the J/ψ spin in terms of polarizations as discussed in section 1.3.

$$\frac{d\sigma}{d\cos\theta^* \, d\phi^*} \sim 1 + \alpha \cos^2\theta^* + \beta \sin 2\theta^* \cos\phi^* + \gamma \sin^2\theta^* \cos 2\phi^* \qquad (4.7)$$

The form in Eq. (4.7) was fit to the 25 cross-section measurements, versus $\cos\theta^*$ and ϕ^*, available for each of the 15 regions of x_F. The sum of the χ^2's for the complete set of fits was 355 for 315 degrees of freedom. The coefficients obtained from the fits are plotted in Figures 4.3-4.5. The striking feature of the data is that $\alpha \to -1$ as $x_F \to 1$. We measure $\alpha = -0.80 \pm 0.17$ for $0.95 < x_F < 1.0$. This indicates that J/ψ become longitudinally polarized at high x_F, though there is no net polarization over most of the range of x_F. The other two angular distribution parameters, β and γ, are consistent with zero across the entire range of x_F.

To study the dependence of α on x_F, and because we needed a form for the Monte Carlo event generator, we parameterized the dependence of α on x_F using a form suggested by a model of higher-twist effects in QCD, see Eq. (1.7).

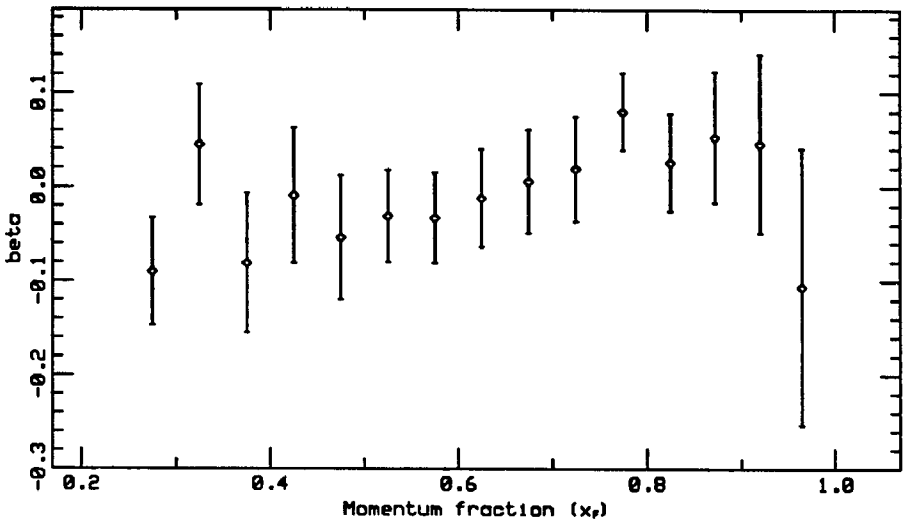

Figure 4.4: The angular distribution parameter β versus x_F

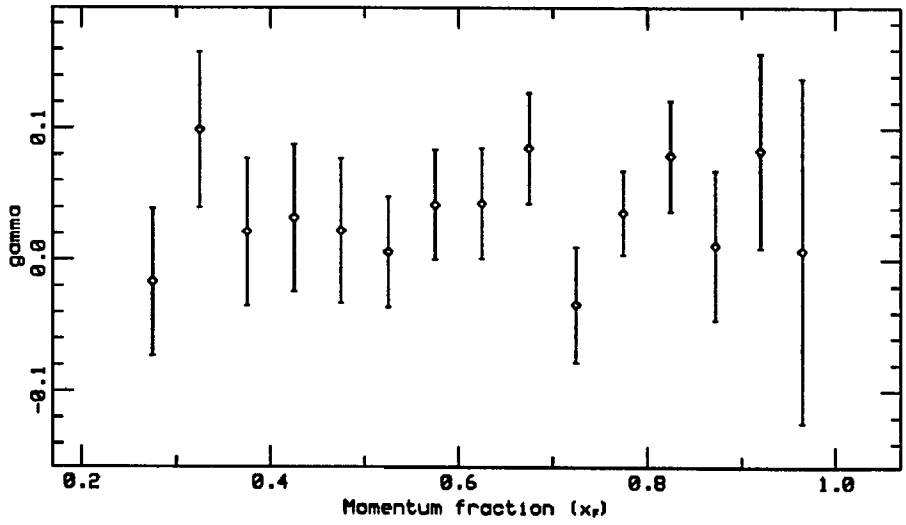

Figure 4.5: The angular distribution parameter γ versus x_F

$k = 0.0024$	$0.85 < x_F < 0.90$	$0.90 < x_F < 0.95$	$0.95 < x_F < 1.00$
α	-0.10 ± 0.11	-0.17 ± 0.14	-0.80 ± 0.17
β	0.05 ± 0.07	0.05 ± 0.10	-0.11 ± 0.15
γ	0.01 ± 0.06	0.08 ± 0.07	0.01 ± 0.13
χ^2	34.8	19.3	12.6
$k = 0.0$	$0.85 < x_F < 0.90$	$0.90 < x_F < 0.95$	$0.95 < x_F < 1.00$
α	-0.07 ± 0.11	-0.06 ± 0.15	-0.57 ± 0.23
β	-0.05 ± 0.07	0.08 ± 0.10	-0.14 ± 0.17
γ	0.07 ± 0.06	0.06 ± 0.07	-0.09 ± 0.13
χ^2	37.1	24.5	20.5

Table 4.4: Angular distribution parameters

$$\frac{d\sigma}{dx_F\, d\cos\theta^*} \sim k \sin^2\theta^* + (1 - x_F)^2 \tag{4.8}$$

This form has no ϕ^* dependence because we did not find any significant variation of the cross section with ϕ^*. To be self-consistent in our calculation of k, we repeated generating acceptances and calculating k until the value of k used in the simulation event generator and the value calculated from the experimental data agreed. Our final, self-consistent, value was $k = 0.0024 \pm 0.0006$.

To be sure that our observation of longitudinal polarization at high x_F was not an artifact of having $k \neq 0$ in the simulation, we also generated acceptances with $k = 0$. Table 4.4 shows the angular distribution results for $x_F > 0.85$ using acceptances calculated with $k = 0.0024$ (self-consistent) and $k = 0$. The value of α in the highest x_F bin for self-consistent acceptance calculation is $\alpha = -0.80 \pm 0.17$. This is four standard deviations from zero. For the acceptance calculation with a flat angular distribution, the size of the effect decreases, $\alpha = -0.57 \pm 0.23$. However, the χ^2/DoF for the high x_F bin has increased from 0.6 to 1.0 and the total χ^2/DoF for $x_F > 0.85$ has increased from 67/63 to 82/63. This gives us confidence that the self-consistent acceptance calculation is correct.

Because we have a significant result in only one x_F bin, one might be wary that the effect is due to the binning of the data. In Fig. 4.6 we plot the previous set of data together with the data with the bins shifted in x_F by half the bin width. The original data points are shown as diamonds and the shifted data points are shown are open squares. The solid curve is α calculated from Eq. (4.8) with $k = 0.0024$. The shifted data agree with the original data and the calculated curve.

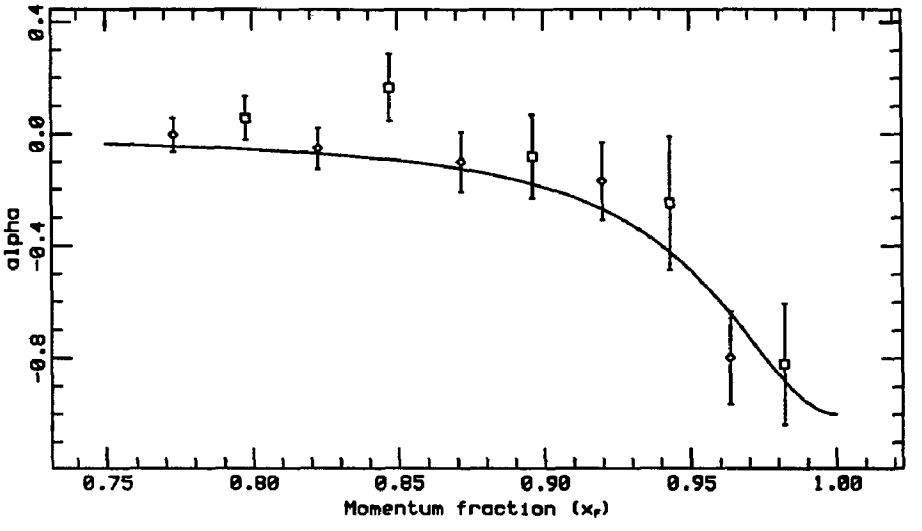

Figure 4.6: The angular distribution parameter α at high x_F. Shown in the plot are the original data points (diamonds), the data in bins shifted by half the bin width in x_F (open boxes), and α calculated from Eq. (4.8) with $k = 0.0024$ (solid curve).

It is of interest to compare the angular distribution we observe with the angular distribution predicted by our extension of the Berger-Brodsky model of higher-twist initial-state corrections presented in section 1.4. To make the comparison, we need a numerical prediction for the parameter k of Eq. (4.8).

We can find a value for k and the three variables which determine k, (the parameters A, M, and k_T^2 defined in Eqs (1.8)-(1.9) in section 1.4) using our measurements of the J/ψ cross section presented earlier. Applying Eq. (1.9) to the results presented in Fig. 4.1 in the region $x_F > 0.75$, we find that $A = 0.10 \pm 0.02$. The appropriate mass in Eq. (1.8) is the J/ψ mass. Since we cannot measure k_T^2, the square of the transverse momentum of the annihilating antiquark, we chose to use as an approximation our measured value of the average $P_T^2 = 0.62 \pm 0.08$ GeV$^2/c^2$ for J/ψ with $x_F > 0.95$.

Inserting these three values into Eq. (1.8), we find that $k = 0.0029 \pm 0.0007$. This value is in agreement with our measured value of $k = 0.0024 \pm 0.0006$. Therefore, our extension of the Berger-Brodsky higher-twist model adequately explains the observed change in the angular distribution.

The data presented here are the first evidence for longitudinal polarization of J/ψ produced in hadronic collisions. The previous best measurement of the parameter α in Eq. (4.7) found $\alpha = 0.06 \pm 0.25$ for $0.8 < x_F < 1.0$ [4], which is consistent with our value of $\alpha = -0.07 \pm 0.09$ over the same interval. The observed x_F dependence of the J/ψ polarization is in agreement with the prediction of our extension of the Berger-Brodsky model of QCD corrections to the quark-antiquark annihilation process. We conclude that the observed spin alignment is an indication of QCD higher-twist effects in the hadronic production of J/ψ.

Part II

Calibration of a Large BGO Electromagnetic Calorimeter

Chapter 5

Introduction

A massive vector boson, the Z_0, was first predicted as a by product of the unification of the electromagnetic and weak interactions in the electroweak theory of Weinberg, Salam, and Glashow [28]. The existence of the particle was later confirmed experimentally by the UA1 collaboration, headed by C. Rubbia [29]. This thesis is the culmination of the author's work on the construction and calibration of a detector designed to study the physics of electron-positron collisions at center-of-mass energies near the mass of the Z_0. In this chapter we will discuss the electroweak theory and examine some predictions for physics processes in e^+e^- collisions near the Z_0 resonance. We emphasize the what is required experimentally to make accurate measurements of the processes. In the following chapter, we describe an electron-positron collider (LEP) and a detector (L3) designed to study high energy e^+e^- collisions. We describe the the design of the electromagnetic calorimeter and its associated electronics in detail. In the final chapter (of this part) we present the results of the calibration of the electromagnetic calorimeter in an electron test beam.

5.1 The Standard Model

Particle physics currently has a theoretical structure, referred to as the 'Standard Model', into which (essentially) all the results of all experiments done to date can be accommodated. The elements of the standard model are three types of particles, quarks, leptons, gauge bosons, and two gauge theories which describe the interactions of the particles. Quarks and leptons are the building blocks of matter. Both varieties of particles have spin $\frac{1}{2}$ and therefore are fermions which obey the Pauli exclusion principle. They are differentiated by their interactions, which are listed in Table 5.1. Only quarks participate in the 'strong' interactions which are described by the gauge theory called quantum chromodynamics. The quarks have an additional quantum number, not shared by the leptons, which can take on one

	Interactions			
Particle	Strong	Electromagnetic	Weak	Charge
Quarks				
u (up)	Yes	Yes	Yes	$+\frac{2}{3}$
d (down)	Yes	Yes	Yes	$-\frac{1}{3}$
Leptons				
e^- (electron)	No	Yes	Yes	-1
ν (neutrino)	No	No	Yes	0

Table 5.1: Interactions of the fundamental fermions

Generation	Quarks		Leptons	
1^{st}	u (up)	d (down)	e^- (electron)	ν_e (e-neutrino)
2^{nd}	c (charm)	s (strange)	μ^- (muon)	ν_μ (μ-neutrino)
3^{rd}	t (top)	b (bottom)	τ^- (tau)	ν_τ (τ-neutrino)

Table 5.2: Generations of fermions

of three possible values and is referred to as 'color'. Color, in quantum chromodynamics, is analogous to electromagnetic charge. All of the fermions participate in the other type of interaction which is described by a gauge theory known as the electroweak theory. At low energies the phenomenology of the electroweak theory splits into the electromagnetic interaction and the weak interaction. The leptons are divided into those which participate in the electromagnetic interaction, electrons, and those that do not, neutrinos. All known fermions participate in the weak interaction.

The fermions listed in Table 5.1 suffice to construct all of the matter we observe at low energies. However, in high energy interactions, additional particles are produced. To date a total of eleven fundamental fermions are believed to have been observed. The Standard Model does not provide any motivation for the existence of more than four fundamental fermions but it does provide a means to classify them. The fermions are organized into 'generations', as shown in Table 5.2. Each generation consists of a doublet of leptons and a doublet of quarks. On basis of this classification, the existence of a twelfth fermion has been predicted, the t or top quark, but has not yet been observed.

The Standard Model predicts very few of the parameters of the fermions. The only prediction is that each doublet of particles is complete and that the charge of

the members of each doublet differs by one unit [30]. The theory has no explanation of why the charges are $+\frac{2}{3}, -\frac{1}{3}, -1$, and 0 or any predictions for the masses or the other properties of the fermions. The standard model is more successful in predicting the interactions of the particles. Precise predictions of the strong interactions are still out of reach because quantum chromodynamics seems mathematically intractable, but significant results have been achieved in the electromagnetic and weak interactions.

5.2 Electroweak Interactions

In 1967, Weinberg suggested a gauge theory which unified the weak and electromagnetic interactions[28]. A similar theory was independently suggested by Salam[31] and both theories used results derived by Glashow in 1961[32]. In the theory, the fundamental carriers of the forces are a triplet of vector bosons, $\vec{W} = (W^{(1)}, W^{(2)}, W^{(3)})$, and a singlet vector boson, B. The four bosons are massless and are assumed to have couplings of comparable strength. These vectors bosons, in their fundamental form, do not describe the real world because there is only one massless vector boson observed, the photon. Furthermore, the theory would imply that the weak interactions were long range and had a strength comparable to that of the electromagnetic interaction, which is in contradiction with the fact that the observed strengths differ by many orders of magnitude.

To extract a theory closer to the real world from the four massless bosons, Weinberg invoked a process known as 'spontaneous symmetry breaking'. With the introduction of a doublet of scalar mesons, spontaneous symmetry breaking rearranges the four massless vector bosons and the doublet of scalars, into three massive vector bosons, one massless vector boson, and a single scalar. The doublet of scalar particles is necessary to provide the additional degrees of freedom needed to make three of the vector bosons massive, since the number of fundamental fields (four from the doublet of scalars and two from each of the massless bosons) must match the number of synthesized fields (three for each massive vector boson, two for the massless vector boson, and one scalar field). The new massless vector boson is identified with the photon and is the carrier of the electromagnetic force. The three massive vector bosons, W^+, W^-, Z_0, are carriers of the weak force. The scalar particle, referred to as the Higgs, has not been observed experimentally. It plays an essential role in determining the mass of the vector bosons and will be discussed in greater detail below.

The usefulness of spontaneous symmetry breaking lies in giving mass to the carriers of the weak force. As originally pointed out by Yukawa, in 1935, the range of a force mediated by a massive particle is exponentially damped with a

characteristic length of $\hbar/mc$, where $\hbar$ is Plank's constant divided by 2π, c is the speed of light, and m is the mass of the particle. For convenience we will follow the standard practice in high energy physics and take $\hbar$ and c equal to 1. The effective weakness of weak interactions is due to the mass of the carrier rather than any intrinsic weakness of the coupling constant. Therefore it is possible to describe the observed world with electromagnetic and weak coupling of similar magnitudes.

Starting with the coupling constant of the electromagnetic interaction and the strength of weak interactions at low energies, the electroweak theory predicts the masses of the vector bosons and the couplings of all known fermions to the bosons. Before presenting the predictions, we first summarize the relevant experimental data as three parameters[1]. The first parameter is the electromagnetic coupling constant $e = \sqrt{4\pi/137.036} = 0.30282$. This parameter is extremely well measured by a wide range of experiments. The best values come from measurements of atomic spectra[34].

The other two parameters involve weak interactions. Weak interactions can be either charge carrying or neutral. The two cases correspond to the interactions mediated by the $W^{\pm}$, a charged particle, versus those mediated by the Z_0, a neutral particle. An example of a charge carrying interaction is $\nu_e + n \rightarrow e^- + p$. A similar neutral interaction would be $\nu_e + n \rightarrow \nu_e + n$. Our second parameter, G, characterizes the strength of charged weak interactions. To define the parameter we use the measured rate of the charge carrying weak decay $\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$. In addition to providing a very direct correspondence between the experimental result and the value of G obtained, this experiment also gives the most accurate determination of G. We define

$$G = 1.0022\sqrt{\frac{192\pi^3\Gamma}{m_\mu^5}} = 1.1664 \times 10^{-5}\ \text{GeV}^{-2} \tag{5.1}$$

where m_μ is the mass of the muon and Γ is the rate of the decay. The factor of 1.0022 is a purely electromagnetic radiative correction. Since this is the only observed channel for muon decay (the limit on all other channels combined is better than 10^{-5}), the rate of the decay is simply the inverse of the, extremely well measured[35], lifetime of the muon.

Our third parameter characterizes the relative strength of charged and neutral weak interactions. To conform with tradition, we prefer to parameterize the relative strength as a angle, θ_W. The parameter can calculated from measurements of the ratio, R, of neutrino-quark and antineutrino-quark cross-sections

[1] We present numerical values for the parameters in the rationalized Heaviside-Lorentz system of units, see for example [33].

$$R = \frac{\sigma(\nu_\mu q \to \nu_\mu q)}{\sigma(\bar{\nu}_\mu q \to \bar{\nu}_\mu q)} \tag{5.2}$$

at neutrino energies much greater than the proton mass. Due to strong interaction radiative corrections and quark mixing, calculation of $\sin^2\theta_W$ from neutrino-quark interactions is quite complex, so we will not present it here. The current best value for $\sin^2\theta_W$ is 0.234 ± 0.013 [36]. We quote errors only for this parameter because the errors of the other parameters are completely negligible in comparison.

A crucial, and the most directly testable, prediction of the electroweak theory is the mass of the vector bosons. The predicted masses of the $W^\pm$ and the Z_0 are

$$M_W = \frac{e}{2\sin\theta_W}(G\sqrt{2})^{-1/2} = 77 \pm 2 \text{ GeV} \tag{5.3}$$

$$M_Z = \frac{e}{2\cos\theta_W \sin\theta_W}(G\sqrt{2})^{-1/2} = 88 \pm 2 \text{ GeV} \tag{5.4}$$

The particles were first observed in high energy proton-antiproton collisions in the Super Proton Synchrotron (SPS) at the European Organization for Nuclear Research (CERN). The discovery was first made by the UA1 collaboration, headed by C. Rubbia, and then quickly confirmed by the UA2 collaboration. The measured masses, from the combined data of UA1 and UA2, are $M_W = 81 \pm 2$ GeV and $M_Z = 93 \pm 2$ GeV. The agreement with the theory is quite good. The slight discrepancy is probably due to errors in the measurement of $\sin\theta_W$ rather than flaws in the theory.

Once the vector boson masses are fixed, their coupling constants are also fixed in order to recreate the observed low energy interactions. The coupling constant of the $W^\pm$ is $g = e/\sin\theta_W$ and the coupling constant of the Z_0 is $g' = e/\cos\theta_W \sin\theta_W$. A significant prediction of the electroweak theory is that the vector bosons couple to all fundamental fermions with the same coupling constant. An interesting consequence of the 'universality' of the electroweak interaction will be discussed below. The theory also fixes the exact form of the interactions. So the polarization and angular dependence of physical processes can be calculated.

It is clear that high energy collisions with sufficient energy to produce the Z_0 and the $W^\pm$ are an extremely useful way to test the electroweak theory. The first such collisions, $p\bar{p}$ at the CERN SPS, confirmed the existence of the Z_0 and $W^\pm$. Future experiments, particularly with e^+e^- collisions, should provide precise tests of the electroweak theory, and may provide a look at physics beyond the theory. In the next section, we will discuss some specific examples of the physics of e^+e^- collisions with center of mass energies on the order of 100 GeV. As a specific example, we will use the machine parameters of the Large Electron Positron collider (LEP),

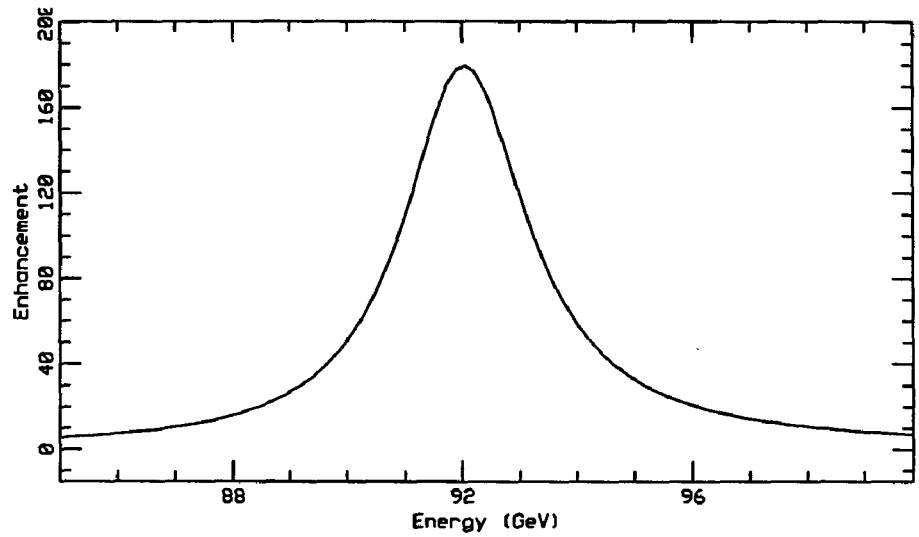

Figure 5.1: $\sigma(e^+e^- \rightarrow \mu^+\mu^-)/\sigma_0$ near the Z_0 mass

being constructed at CERN (see section 6.1), when discussing the practicality and accuracy of measurements to be made. The LEP collider is planned to provide center of mass energies of up to 120 GeV at a luminosity of 10^{31} cm^{-2} sec^{-1} in its first stage, after which the energy should be increased to 200 GeV. We will concentrate on physics of interest for the first stage of LEP.

5.3 Physics at the Z_0

The physics of e^+e^- collisions at energies near 100 GeV is dominated by the Z_0. Figure 5.1 shows the enhancement in the cross-section for $e^+e^- \rightarrow \mu^+\mu^-$ relative to the pure electromagnetic cross-section,

$$\sigma_0 = \frac{64\pi^3\alpha^4}{3E_{cm}^2} \tag{5.5}$$

as a function of center of mass energy, E_{cm}, for energies near the Z_0 resonance. The presence of the Z_0 increases the $\mu^+\mu^-$ cross-section by more than two orders of magnitude . There is an even larger cross-section for quark production, leading to an enhancement of 10^3 in the total cross-section $\sigma(e^+e^- \rightarrow \text{anything})/\sigma_0$ at the

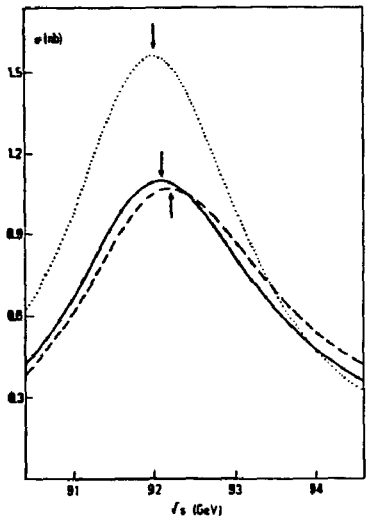

Figure 5.2: Shape of the Z_0 resonance

peak of the resonance. This vast increase in cross-section makes direct production of Z_0's in e^+e^- collisions an ideal laboratory for the study of electroweak physics and an excellent means to search for new particles and new phenomena. [2]

The first experiment done at the Z_0 resonance will almost certainly be to scan, in energy, across the width of the resonance to precisely determine its shape. A detail graph of the predicted shape of the resonance is shown in Fig. 5.2 taken from [37]. The shape is greatly affected by radiative corrections, *i.e.* radiation of photons by the incident e^+e^- before the collision. Figure 5.2 shows the uncorrected shape (dotted line), the first order correction (dashed line), and the second order correction (solid line). An efficient way to perform the scan would be to vary the energy of the beams in a sequence of fine steps, (*e.g.* 1 GeV), across the resonance, collecting sufficient statistics at each step to make a precise determination of the cross-section.

At LEP, it is estimated that the beam energy at each step can be determined from machine parameters to an accuracy of 3×10^{-4}. This energy resolution is significantly better than the resolution of any of the proposed detectors, and is the

[2]We have assumed that the width of the resonance is 2.6 GeV and that the branching ratio to $\mu^+\mu^-$ is 3%. These assumptions will be discussed later.

limiting error in determining the mass of the Z_0. Other sources of error are uncertainty in the calculation of the radiative corrections, which can be made accurate to a few times 10^{-6}, and the uncertainty is measuring the beam luminosity at each step. A precise knowledge of the luminosity is necessary to convert the number of events observed to a cross-section for each energy. However, only the relative cross-sections are required to determine the mass of the Z_0. It should be possible to measure the relative luminosity to an accuracy of 2%, which induces an error on the Z_0 mass of 1×10^{-4}.

Unfortunately a measurement of the Z_0 mass to this accuracy is not, when taken alone, a test of the electroweak theory. The prediction of the Z_0 mass is only good to 5% because of uncertainty in the experimental value of $\sin^2\theta_W$. It is hoped that in the near future the value of $\sin^2\theta_W$ will be measured, via the ratio of the W and Z_0 masses in $p\bar{p}$ collisions at the SPS and the Fermilab Tevatron, to an accuracy of 2×10^{-3}. However, this is still almost an order of magnitude less accurate than the LEP result is expected to be. To make a precision test of the electroweak theory we need to make another measurement at LEP.

The best tests of electroweak theory come from predicted asymmetries in the final state particles. There should be asymmetries in the distribution of charge and in the polarization of the particles. The more straightforward measurement is the charge asymmetry. We present a short discussion of the effect below. The discussion is based on treatments of charge asymmetry in [37] and [38].

We consider the process $e^+e^- \rightarrow (Z_0 \text{ or } \gamma) \rightarrow e^+e^-$. At the collider, the relative orientation of the incoming e^+ and e^- are fixed. To measure the charge asymmetry of the process, we divide our (hypothetical) detector into two parts separated by a plane orthogonal to the beam line and passing through the interaction point, and then count the relative number of final state positrons and electrons emitted into one of the halves. We define the 'forward' half of the detector to be on the side of the incoming positron. If we define n_p and n_e as the number of positrons and electrons emitted into the forward half then the charge asymmetry is given by

$$A = \frac{n_e - n_p}{n_e + n_p} \tag{5.6}$$

The theoretical prediction in the standard electroweak theory for the charge asymmetry for e^+e^- produced at the peak of the Z_0 resonance with unpolarized beams is

$$A = 3(1 - 4\sin^2\theta_W)^2 \tag{5.7}$$

Since $\sin^2\theta_W$ is very close to $\frac{1}{4}$, the predicted asymmetry is quite small. For $\sin^2\theta_W = 0.234 \pm 0.012$, the asymmetry can lie anywhere in the range $0.001 < A <$

0.038. To analyze the experiment, we are interested in the accuracy with which we can determine $\sin^2\theta_W$ given a measured value of the asymmetry.

$$\Delta \sin^2\theta_W = \frac{1}{24}\left(\frac{A}{3}\right)^{-1/2}\Delta A \tag{5.8}$$

The expected asymmetry is very close to zero, so the error on $\sin^2\theta_W$ depends very sensitively on the actual value of the asymmetry. From the current best value for $\sin^2\theta_W$ we would predict an asymmetry of 0.012. If the asymmetry is 0.012, then to determine $\sin^2\theta_W$ to an accuracy of 0.002, we would need an error in the measurement of the asymmetry of < 0.003. Sufficient statistics for this level of sensitive could be obtained in 100-200 days of running at the Z_0 peak. However, the response of the detector must be uniform to within a few parts per thousand. Achieving and maintaining such uniformity over a period of a few hundred days would be very difficult.

The situation improves if the e^+e^- beams are longitudinally polarized. A polarization of 50% is considered to be achievable at LEP[39]. With a polarization of 50%, the charge asymmetry is

$$A = 3\eta\frac{\eta + 0.25}{1+\eta} \tag{5.9}$$

where $\eta = 1 - 4\sin^2\theta_W$. The expected asymmetry is 0.057, significantly larger than in the unpolarized case. In addition to the asymmetry being larger, the dependance of the weak mixing angle on the asymmetry is more benign. To achieve a precision of 0.002 on $\sin^2\theta_W$, we need to measure A to an accuracy of 0.008. Polarization also offers the possibility of measuring, and correcting for, systematic experimental biases by performing measurements with the polarization reversed.

Detailed study of the Z_0 is extremely important to test the electroweak theory. However, the Z_0 is also a useful tool for the study of non-electroweak physics because of the universality of the its couplings and its large resonance cross-section. The rate for the decay of the Z_0 into a massless fermion pair of charge Q is

$$\Gamma(Z_0 \to f\bar{f}) = \frac{Gm_Z^3}{24\pi\sqrt{2}}\left[1 + (1 - 4Q\sin^2\theta_W)^2\right] \tag{5.10}$$

where m_Z is the mass of the Z_0, and G was defined above. For quarks, the rate must be multiplied by three in order to account for the three possible colors of quarks. The assumption that the fermions are massless is quite good for all known fermions except there is a slight error for the b quark. If the t quark mass is $< m_Z/2$ then the Z_0 will decay to $t\bar{t}$ pairs and the phase space adjustment for the mass of the t quark will be significant.

Fermion	Partial Width
u (up)	305 MeV
d (down)	393 MeV
e^- (electron)	86 MeV
ν (neutrino)	170 MeV

Table 5.3: Contribution of fermion pairs to the Z_0 width

The rate for Z_0 decay into a pair of fermions is nonzero for all known fermions. If the Z_0 coupling is universal for all fermions, as the electroweak theory suggested, then all fermions with mass less than $m_Z/2$ should contribute to the total width of the Z_0. This allows us to measure the total number of fermions with mass less than $m_Z/2$ by simply measuring the total width of the Z_0. The predicted contribution to the width for each type of known fermion is listed in Table 5.3 [37]. The quark widths include the factor of three for color. There is an uncertainty in the width for decay into quark pairs of $\pm 1\%$ due to uncertainties in the strong interaction radiative corrections. Because the mass of the top quark must be a significant fraction of the mass of the Z_0, its contribution to the width of the Z_0 depends strongly on its mass. There will be an additional ambiguity until the top quark mass is known, or shown to be greater than $m_Z/2$. In the best case the measurement will be useful only to an accuracy of ± 20 MeV.

A better way to count the number of neutrinos is too use the process

$$e^+e^- \rightarrow \gamma(Z_0 \rightarrow \nu\bar{\nu}) \tag{5.11}$$

One makes collisions at an energy a few GeV above the Z_0 mass, and looks for events which have a single photon with no other energy deposited in the detector. The number of neutrino-like particles is proportional to the cross-section observed. The main advantage of this technique is that it eliminates the uncertainty in the Z_0 width due to the quark decay modes. The main background to the process is radiative Bhabba scattering, $e^+e^- \rightarrow e^+e^-\gamma$, where the final state e^+ and e^- were not observed. This happens most often when the momenta of the e^+ and e^- are not significantly changed by the emission of the photon. This background can be reduced to a negligible level by requiring the photons to have a large angle ($> 25°$) relative to momenta of the incident e^+ and e^-. The error in the measurement is likely to be dominated by either uncertainty in the efficiency of the trigger for low energy photons or by the accuracy of the cross-section normalization from the luminosity monitoring.

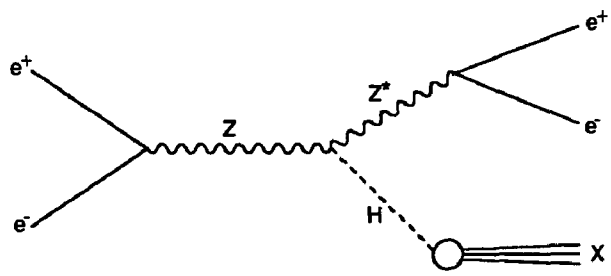

Figure 5.3: Radiative production of Higgs by Z_0

If one is willing to assume that all fermions are grouped into 'generations', as discussed previously, and that the neutrinos of all generations have masses of less than 45 GeV, then the width of Z_0 will determine the total number of generations. However, it is very likely that the world is more complicated than that. There are many possible types of particles waiting to be discovered. One only has to glance through the myriad of exotic particles concocted by theorists, or simply trust that Nature still has a few more surprises. Many of the particles couple to the Z_0 and several have masses of less than 45 GeV. All particles that meet these two criteria will contribute to the total width of the Z_0. Many, (probably most since they have not yet been detected), of these particles have very weak interactions, and so will effectively mimic neutrinos. Simply measuring the width of the Z_0 will not be sufficient to distinguish the various possibilities. However, those particles that have sufficiently short lifetimes that they decay in the detectors will produce a bounty of new physics.

One particle that will be searched for in Z_0 decays is the Higgs scalar, mentioned above. The Higgs is the current 'Holy Grail' of particle physics. Its fundamental role of generating particle masses makes it one of the most interesting particles ever predicted. The couplings of the Higgs are specified by the electroweak theory, but the theory makes no prediction of its mass or charge. In the discuss below, we will assume the minimal Higgs allowed by the theory: a single neutral scalar particle. If the Higgs has a mass is in the range from 10 GeV to 60 GeV and behaves as predicted, it should be discovered at LEP.

One way to search for the Higgs is to look for the process in Fig. 5.3. The Higgs is radiated by the Z_0, turning the real Z_0 into a virtual Z_0 which then decays with an invariant mass less than the physical Z_0 mass. The relative rates of the decays

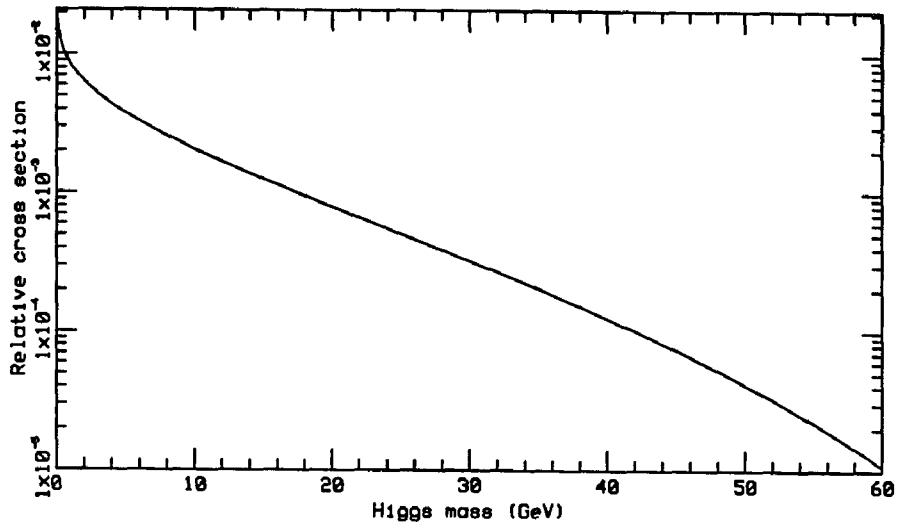

Figure 5.4: Relative rates $\Gamma(Z_0 \to He^+e^-)/\Gamma(Z_0 \to e^+e^-)$

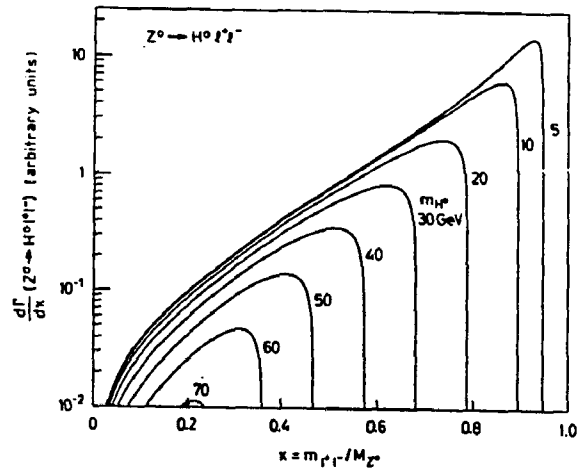

Figure 5.5: Invariant mass distribution for final state leptons in $Z_0 \to He^+e^-$

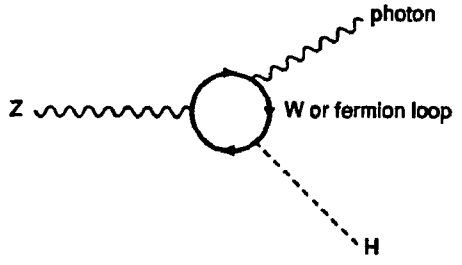

Figure 5.6: Decay of Z_0 to $H\gamma$

$Z_0 \to He^+e^-$ and $Z_0 \to e^+e^-$ is shown in Fig. 5.4 [40]. At the nominal luminosity for LEP of 10^{31} cm^{-2} sec^{-1}, one would expect to observe on the order of 1 Higgs per day if the Higgs mass is 20 GeV. The expected invariant mass distribution for the final state leptons is shown in Fig. 5.5 taken from [40]. To find the signal, one would examine the invariant mass spectrum of dielectrons in events with $e^+e^- + X$, where X = anything in the final state, and then fit the distribution to determine the mass of the Higgs. The primary background comes from Z_0 decays involving the top, $Z_0 \to (t \to e^+X)(\bar{t} \to e^-X)$, if the top mass is $< m_Z/2$. One can identify most background events because the top decays will have two widely separated pairs of $e^\pm + X$, while the Higgs signal should have two isolated electrons and the remaining particles should have a total momentum antiparallel to the total momentum of the e^+e^- pair.

Another decay of the Z_0 which involves the Higgs is $Z_0 \to H\gamma$. This decay proceeds via the one-loop process shown in Fig. 5.6. The particle in the loop can be either a fermion or a W. The loop involving the W has a much larger amplitude than any of the fermion loops. Therefore we consider only the W loop and calculate the total rate for the process to be

$$\frac{\Gamma(Z_0 \to H\gamma)}{\Gamma(Z_0 \to \mu^+\mu^-)} = 5.9 \times 10^{-5} \left(1 - \frac{m_H^2}{m_Z^2}\right)^3 \left(1 + 0.11\frac{m_H^2}{m_Z^2}\right) \tag{5.12}$$

The rate is significantly lower than the rate for $Z_0 \to He^+e^-$. It is unlikely that the Higgs will be discovered via this process. However, for center of mass energies below $2M_W$, this process is the best way to measure the coupling of the Higgs to the W bosons.

The source of interesting physics near 100 GeV might come from bound states

of the top quark. As discussed previously, the top quark is the six quark. It has been predicted, in order to fill out the doublet occupied by the bottom quark, but has not yet been observed. The best current limit on the mass of the bottom quark comes from experiments done by the UA1 collaboration. They set a lower limit on the top quark mass of 45 GeV[41]. This is just within the range accessible to the first stage of LEP. Toponium is more interesting than 'just another' $q\bar{q}$ bound state because the large mass of the quark makes the system very tightly bound, providing a probe of the strong interactions at very short distances. A top quark mass greater than 30 GeV will allow us to test the principle of asymptotic freedom and measure the inter-quark potential in regions where QCD predicts a coulombic potential.

The spectroscopy of toponium is likely to be very different from the spectroscopy of the lighter quark systems. Figure 5.7 shows the predicted branching ratios for the various decay modes of the $2S$ $t\bar{t}$ state. The modes are

SQD Single Quark Decays. In these decays one of the top quarks decays, via emission of a W, into a bottom quark. The W decays into a doublet of quarks or leptons. The other top quark is a spectator.

l^+l^- Purely leptonic decays. The rate shown is for decays to $\mu^+\mu^-$. These decays proceed by $t\bar{t}$ annihilation in a Z_0 for neutrinos, or annihilation into a Z_0 or a photon for charged leptons.

$\Sigma q\bar{q}$ Decays to lower mass $q\bar{q}$ pairs. These decays proceed by $t\bar{t}$ annihilation into a Z_0 or a photon, the same process as for charged leptons. Decays to $b\bar{b}$ can also occur via W exchange between the t and $\bar{t}$ quarks.

ggg Decays via the strong interaction. The t and $\bar{t}$ annihilate producing either three gluons or two gluons and one photon. The gluons give rise to hadronic jets.

1P,γ Electromagnetic decay to the 1P state (referred to as χ_t) via emission of a photon.

Due to the presence of the Z_0 resonance, the rates of the decay modes vary greatly over a relatively narrow mass range. Study of electromagnetic transitions between the $t\bar{t}$ resonances will be possible only below the Z_0 mass. Above the Z_0 mass, the decays will be dominated by weak processes leading, primarily, to single quark decay modes. The strength of the electroweak decays will probably limit detailed study of the $t\bar{t}$ system to the lowest few states. The inter-quark potential can be determined from the energies of the 1S and 2S states. However,

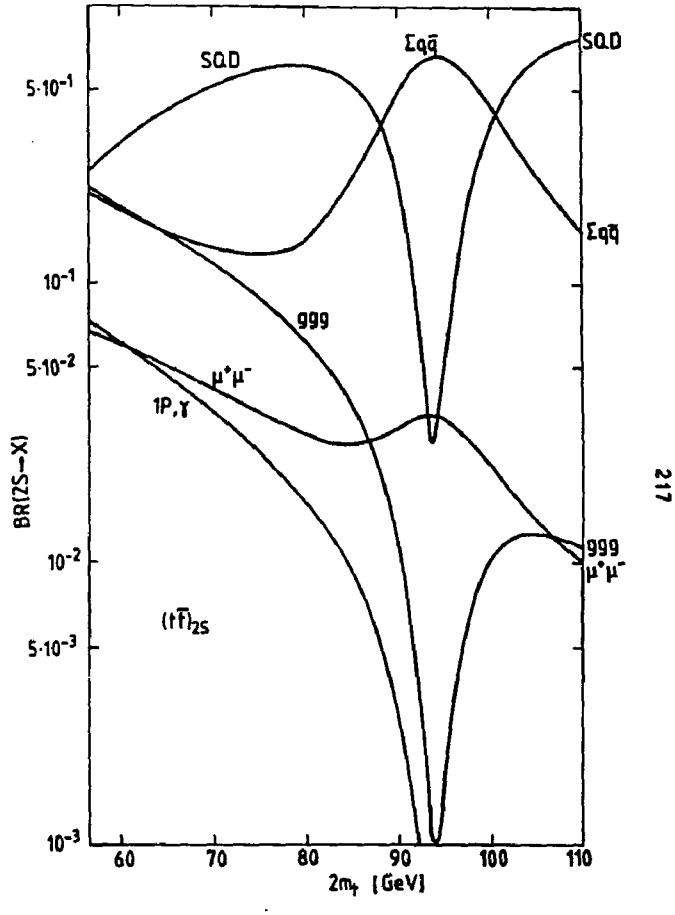

Figure 5.7: Branching ratios for the decay modes of 2S $t\bar{t}$ as a function of $2m_t$

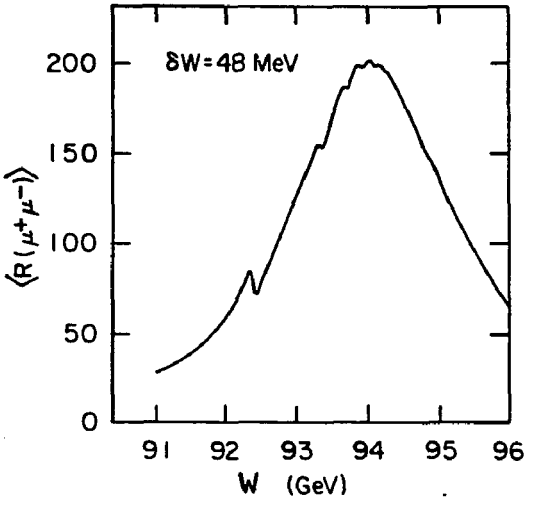

Figure 5.8: Cross-section for $e^+e^- \to \mu^+\mu^-$ when $2m_t = M_Z = 94$ GeV

more information can be derived from the leptonic decays of the states, which provide a direct measurement of the overlap of the quark wave functions.

If the $t\bar{t}$ mass lies within the Z_0 resonance, then it will be possible to measure the top spectrum via the interference of the Z_0, γ, and $t\bar{t}$ propagators. Figure 5.8 shows one prediction for total cross-section for $e^+e^- \to \mu^+\mu^-$ for $2m_t = M_Z = 94$ GeV. The figure includes the effect of the nonzero momentum spread of the beams (the assumed spread in the center of mass energy is 48 MeV). Close proximity of $2m_t$ and m_Z would complicate study of the $t\bar{t}$ system and all of the Z_0 decay measurements discussed above, but it would also lead to some very interesting physics.

Toponium also offers another way to search for the Higgs. If the top quark is light enough to be produced and the mass of the higgs is less than the toponium mass, then the Higgs will be produced in radiative decays of toponium. The process is shown in Fig. 5.9. Since the Higgs coupling is proportional to the mass of the top quark, the large mass of the top produces an appreciable rate for this process. The first order rate for the process is

$$\Gamma(\theta \to H\gamma) = \frac{1}{8\sin^2\theta_W}\frac{m_\theta^2}{m_W^2}\left(1 - \frac{m_H^2}{m_\theta^2}\right)\Gamma_0(\theta) \qquad (5.13)$$

where m_θ is the mass of the $t\bar{t}$ bound state and $\Gamma_0(\theta)$ is the rate of the decay

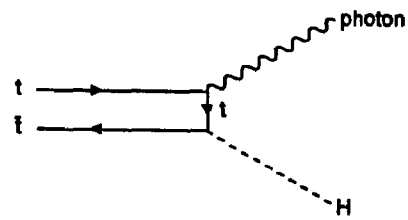

Figure 5.9: Higgs production in toponium decays

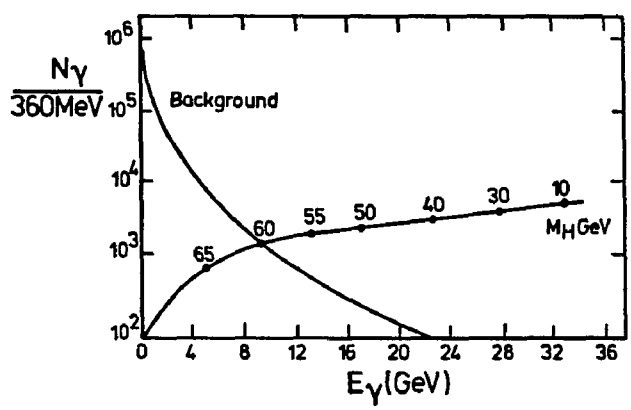

Figure 5.10: Rate of Higgs production in toponium decays

to a l^+l^- lepton pair. There are radiative corrections due to gluons exchanged between the three t quark legs which reduce the rate by approximately a factor of 2. The number of $t\bar{t}$ events, including radiative corrections, and the number of background events for 1000 hours of running at $m_\theta = 70$ GeV has been plotted for various Higgs masses in Fig. 5.10 taken from [43]. The main backgrounds come from $e^+e^- \rightarrow (Z_0, \gamma) \rightarrow \gamma X$ and other single photon $t\bar{t}$ decays. The decay has an easily identifiable signature: monochromatic photons. The experimental requirements to observe this decay are good photon energy resolution and an effective trigger for single photon events.

It is clear from the, by no means exhaustive, list of physics processes presented here, that e^+e^- collisions at center-of-mass energies near the Z_0 mass will produce a substantial amount of interesting results. However, in order to study the physics, certain experimental requirements must be met. The most important parameters of the collider are its luminosity and energy. A very luminous collider will allow detection of particles which account for only a small fraction of Z_0 decays, most notably the Higgs. An energy range which extends somewhat above the Z_0 peak, combined with accurate energy and luminosity measurements, will allow precise measurements of the total number of neutrinos and possibly indicate the existence of exotic weakly interacting particles. The important requirements for detectors are: uniform coverage over a large solid angle, good energy resolution (particularly for leptons), and accurate particle identification and separation. Uniform coverage over a large solid angle and accurate luminosity monitoring are necessary for the tests of the electroweak theory described above. Good energy resolution and the ability to distinguish individual particles are necessary for new particles searches and to study the top quark.

Chapter 6

Hardware

6.1 LEP

The Large Electron Positron collider (LEP) is an electron-positron accelerator designed to exploit the Z_0. The first stage of construction of LEP will provide energies up to 49 GeV per beam and a luminosity which will produce 10,000 Z_0's per day when running at the Z_0 peak. The second phase of LEP will increase the energy to 100 GeV per beam, in order to allow production of W^+W^- pairs.

The collider is currently under construction at the European Organization for Nuclear Research (CERN) under the French countryside near Geneva, Switzerland. Preliminary proposals for the collider were made as early as 1976. The first detailed design study of the collider in its current form was presented in August, 1978 [44]. The initial design of the collider was done in light of the electroweak theory of Weinberg, Salam, and Glashow, but predated the discovery of the Z_0 in 1981. It is a testament to the influence of the electroweak theory, that a decade long accelerator construction program was initiated on one of its predictions, long before the prediction was experimentally verified.

The main ring of the collider has a circumference of 26.7 km and ranges from 30 to 150 meters below the ground surface. The ring has 3328 dipole bending magnets and more than 1200 quadrupole and sextupole focussing magnets. The dipole magnets have a bending radius of 3.1 kilometers, and a maximum field strength of 0.13 Tesla. The same magnets will be used in the second stage of LEP and are strong enough to keep 110 GeV electrons in orbit. The collider is designed to provide a luminosity of 10^{31} $cm^{-2}sec^{-1}$ with collisions of electron and positron bunches (beam crossing) at a given interaction point every 11 μsec.

The RF system which will provide high intensity electric fields for the acceleration of particles in the first stage of LEP will consist of high efficiency klystrons powering acceleration cavities coupled to low-loss RF cavities. All three components are made of conventional (*i.e.* non-superconducting) materials. Stored energy os-

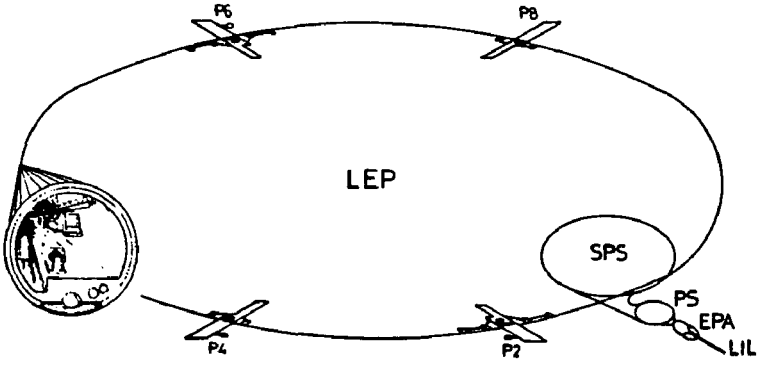

Figure 6.1: LEP and its injection system

cillates between the acceleration cavities and the storage cavities with the phase adjusted to have the energy in the acceleration cavity when a bunch enters the cavity. The storage cavity has a significantly higher Q value than the acceleration cavity which reduces the power dissipation. For the second stage of LEP, the cavities will be replaced by superconducting cavities now under development. The number of RF cavities will increase with time. At start-up there will be two acceleration section with a total length of 271.5 meters, dissipating a total of 16 Megawatts.

The minimum energy at which particles can be maintained in the main ring of LEP is approximately 14 GeV. To accelerate the electrons and positrons to this energy before they are injected into the main ring, a series of lower energy accelerators is used. The injection system for LEP, shown in Fig. 6.1, consists of two specially constructed linacs and an accumulator ring, and two accelerators previously built at CERN. The two linear accelerators are called the Linear Injector for LEP (LIL). Electrons beams are produced in an electron gun and then accelerated to 200 MeV in the first linac. Positrons are produced from the 200 MeV electron beam via a converter. Both beams are accelerated via the second linac to 600 MeV. The 600 MeV beams are stored in an accumulator ring, referred to as the EPA. The accumulator collects 1250 linac pulses to form four bunches. The four bunches are then transferred to the CERN Proton Synchrotron (PS). The PS has been modified

to accelerate electron and positron bunches to 3.5 GeV. The 3.5 GeVbeams are fed into the CERN Super Proton Synchrotron (SPS) and accelerated to 20 GeV for injection into the main LEP ring. A total of 30 bunches from the SPS of each sign are required to fill LEP. The total time for one filling cycle is 16 minutes. It is estimated that LEP will need to be refilled every few hours.

At the startup of LEP, the electron and positron beams in the main ring will be made to collide at four points (interaction points). Each interaction point will be the site of a detector designed to measure the particles produced in the collisions. In the next section, we describe one of the four experiments currently being installed at LEP.

6.2 The L3 Detector

The L3 detector is designed to make precision measurements to test the Standard Model and to search for new particles and new phenomena. The detector is primarily designed to study muons, electrons, and photons. Hadrons are measured crudely, with an order of magnitude less precision than electromagnetically interacting particles. The choice to concentrate on leptons and photons is historically sound. A large fraction of the significant new results in particle physics over the past two decades, including many results on the nature of hadrons, have come from precise measurements of leptons in the final state. Examples include the discovery of the J/ψ and the Υ, discovery of point-like constituents of nucleons through deep inelastic scattering, and study of the Drell-Yan process.

Leptons and photons are preferred probes because it is much more difficult to accurately measure hadrons than to measure leptons and photons. There are several reasons for this, all of which can ultimately be traced to the strength of the strong interaction. The primary difficulty is that the hadrons observed in nature are not fundamental particles. Amplitudes for high energy processes are calculated in terms the interactions of quarks, rather than of mesons or hadrons. However, each primary quark produces a 'jet' of secondary hadrons. Deducing the properties of the primary quarks from the observed hadrons can be done only imprecisely, since there is no calculable theory that accurately describes how quarks cloth themselves to become hadrons.

More difficulties arise even if we limit ourselves to only measuring the total energy and direction of jets. Hadronic interactions produce a variety of types of particles which have widely differing detection efficiencies. Fluctuations in the spectrum of particles produced, particularly in the first few interactions, limit the achievable resolution. A related problem is that since hadrons interact with nuclei, a large fraction of the initial energy is spent exciting or causing the breakup of nuclei. Only a

small part of the energy is converted to a detectable signal. Additionally, jets tends to have significant transverse momentum, since the transverse momentum of a jet is proportional to the mass of its constituents and hadrons have masses upward of a few tenths of a GeV. The transverse momentum increases the lateral extend of the shower, making it difficult to distinguish overlapping jets from a single jet.

In contrast, muons, electrons, and photons are, as far as we know, fundamental particles. The particle in the physical process of interest retains its identity until it interacts in the detector. There is no additional layer of calculation necessary to determine the nature of the primary particle, and only a single particle appears in the detector. Since the interactions are only electromagnetic, the vast majority of secondary particles are electrons and photons. Essentially all methods used to coax measurable signals from physical interactions rely on electromagnetism, so electrons and photons are the most readily measurable secondaries. Additionally, the long range character of electromagnetism greatly reduces fluctuations in the production of secondary particles relative to the case of hadrons.

Proof of the relative difficulty of measuring hadrons versus electromagnetically interacting particles lies in the results that have been achieved. The best energy resolution achieved in a hadron calorimeter is $\Delta E/E = 5\%$ at 50 GeV. For electrons and photons a calorimeter has been constructed (described below) with a resolution of $\Delta E/E < 1\%$ at 50 GeV, and for muons a magnetic spectrometer has been constructed (also described below) with a momentum resolution of $\Delta p/p < 1.9\%$ at 50 GeV.

The choice in the design of L3 was to concentrate on the subset of events which can be measured very accurately, (*e.g.* electron and muon pairs have a combined branching ratio of only 6% in decays of the Z_0). The increase in precision and the ability to unambiguously resolve each component of the selected events more than compensate for the loss in statistics. Confucius had a similar idea when he said, 'A single gem is more pleasing to the eye than a multitude of common stones'.

A sketch of the L3 detector inside the experimental hall is shown in Fig. 6.2 [45]. The hall is 21.4 m in diameter and 26.5 m long. The floor of the hall is 50 m below ground surface. The dominant feature of the detector is the magnet. The entire detector rests inside 0.5 T magnetic field, a so called 'magnetic cave'. The magnet yoke has a outer radius of 7.8 m and a length of 11.6 m. The inner radius of the magnet coil, and therefore the maximum size of the detectors housed inside, is 5.93 m. There are four distinct particle detectors. Moving outward from the beam, the detectors are the Vertex Chamber, the Electromagnetic Calorimeter, the Hadron Calorimeter, and the Muon Chambers. Each subdetector provides a specialized measurement, as indicated by the name of the subdetector. Each subdetector will be discussed, briefly, below. Before preceding, we should introduce

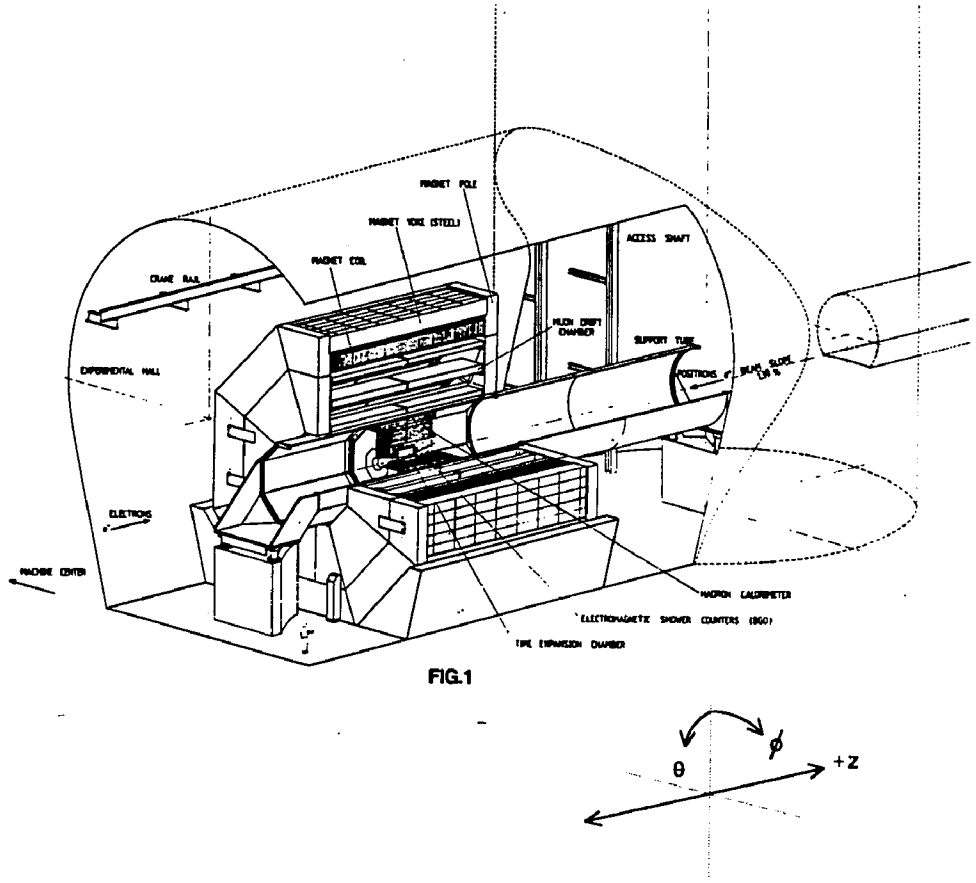

Figure 6.2: The L3 detector

the coordinate system which will be used through out the this thesis. As shown in Fig. 6.2, we define z as the direction along the beam, ϕ as the angle in the plane perpendicular to the beam, r as the radial distance from the interaction point, and θ as the angle in a plane containing the beamline. Note that we define $\theta = 0$ to be perpendicular to the beamline.

The muon chambers are designed to precisely measure the momenta of high energy muons. The magnet and the muon chambers form a magnetic spectrometer. The chambers are assembled in 16 units, referred to as octants. Each octant consists of five chambers in three layers which cover one eighth of a circle in ϕ and an angular range in θ from 0° to ±46°. The main chambers measure only the ϕ coordinate of the particle tracks. Since the magnetic field is parallel to the beamline, only the ϕ coordinate of a muon track is bent by the magnetic field. The resolving power of a magnetic spectrometer is proportional to $B(L/d)^2$, where B is the magnetic field strength, L is the dimension of the spectrometer transverse to the field, and d is the position resolution on measurements of the muon track. The decision for L3 was to maximize L/d. This reduces the requirements on the magnetic field strength, which makes use of a magnet constructed out of conventional materials feasible. The outer most chamber is 5.44 m from the beam line, while the inner chamber is 2.53 m from the beam line. So, the lever arm for momentum measurements is 2.91 m. The sagitta of a 45 GeV muon across 2.91 m in a 0.5 Tesla magnetic field is only 3.7 mm.

To have momentum resolution on the order of 1%, the muon tracks must be measured to an accuracy of several tens of microns. This accuracy is achieved by measuring each track multiple times in each of the three layers of chambers. The accuracy of the position determination should decrease with the square root of the number of measurements. There are 16 planes of measuring wires in the inner and outer chambers and 24 planes in the middle chambers. The position resolution achieved is 50 μm. The alignment of the chambers within each octant is monitored and corrected by an opto-mechanical system. The alignment can be checked with a UV laser system or by using beam or cosmic ray muons if the magnet is off. Tests of the chambers have demonstrated a systematic error in position measurements of less than 30 μm. It is important to note that the alignment between different octants need not be known to such great accuracy because muons with momenta greater than 3 GeV/c will traverse only one octant. The overall momentum resolution of the chambers is $\Delta p/p = 1.9\%(50\ \mathrm{GeV}/p)$, which gives a mass resolution for dimuons of 1.4% at 100 GeV.

In addition to the chambers which measure ϕ, there are four layers of chambers, arranged in pairs that surround the innermost and outermost layers of ϕ chambers, that measure the z-coordinate of the tracks. Since the magnetic field induces no

bending in the z-coordinate, the track of muons through the z-chambers is a straight line. These 'z-chambers' are used to resolve ambiguities in track reconstruction. Also, straight-line tracks which pass through the interaction point are used as a fast trigger signal to identify muon events. There is only a single plane of wires in the z-chambers which has position resolution on the order of 250 μm.

The next detector, as we move towards the interaction point, is the hadron calorimeter, (HC). The HC is designed to measure the total energy of events (in conjunction with the electromagnetic calorimeter), to measure the direction of energy flow in hadronic jets, and to filter and identify muons passing into the muon chambers. It is a sampling detector, constructed as a barrel, two end caps, and a muon filter. The barrel and endcaps are sandwiches of alternating layers of depleted Uranium absorber and Argon-CO_2 proportional chambers. The muon filter, mounted between the muon chambers and the HC barrel, has brass absorber, rather than Uranium, in order to block passage of radioactive decay products of the Uranium into the muon chambers. There is also a layer of stainless steel on the inner surface of the HC to shield the electromagnetic calorimeter from the Uranium. The HC covers more than 98% of the total solid angle with more than four hadronic interaction lengths of material in the radial direction.

The proportional chambers have a 5 mm $\times$ 10 mm cell size, which matches the width of a single hadronic shower. There are more than 10,000 planes of proportional chambers, with a total of more than 400,000 cells. The cells are read out in groups in order to reduce the number of readout channels to about 26,000. The measured energy resolution is $55\%/\sqrt{E/\text{ GeV}} + 5\%$ for pions and $34\%/\sqrt{E/\text{ GeV}}$ for electrons.

Inside the hadron calorimeter is the electromagnetic calorimeter, (EC). The EC is designed to provide good energy and angular resolution for electrons and photons. The calorimeter also acts as the first layer of the hadron calorimeter, providing an additional hadronic interaction length of material. The calorimeter is a total absorption scintillating device. Bismuth Germanate (BGO) was chosen as the scintillating material because of its very short radiation length. Absorption of more than 98% of the energy of a 100 GeV electron is possible with only 24 cm of BGO. The calorimeter consists of a two barrel pieces and two endcaps. Construction of the endcaps has been delayed due to financial constraints. The angular coverage of the barrel extends down to $\theta = \pm 48°$. The endcaps will extend the coverage to $\theta = \pm 78°$.

The calorimeter is made of 10752 BGO crystals, 7680 in the barrel and 3072 in the endcaps. The crystal size, 2 cm $\times$ 2 cm at the inner surface of the calorimeter, gives an angular resolution of a fraction of 1° on the centroid of energetic electromagnetic showers. Collection of scintillation light is performed by silicon photodiodes.

The use of more sensitive phototubes is ruled out because of the 0.5 Tesla magnetic field which permeates the L3 detector. Photodiodes have no internal amplification, so a low-noise charge-sensitive preamplifier is used to boost the signal. The energy resolution of the calorimeter is dominated by electronic noise at energies below 1 GeV and is determined by shower fluctuations in BGO above 10 GeV. The energy resolution of the calorimeter has been demonstrated to be 5% at 100 MeV improving to 1% at energies greater than a few GeV.

The inner most detector is the vertex chamber. The vertex chamber is a high resolution gaseous wire chamber, which gives a high resolution picture of the projection of charged particle tracks on to the $r - \phi$ plane, surrounded by proportional chambers, which give the z-coordinates of tracks leaving the wire chamber. The design goals are to measure particle life times to better than 10^{-13} seconds, to determine the charge of electrons with transverse momenta less than 50 GeV/c, to provide momentum resolution adequate to distinguish between electron and pion showers in the calorimeters via the energy (measured in the calorimeters) versus momentum correlation, and to measure the charge particle multiplicity and aid in jet analysis. The lifetime requirement leads to a position resolution requirement of 30 μm. A 4σ determination of the sign of 50 GeV electrons in the 0.5 Tesla magnetic field within the 36 cm active radius of the chamber requires position resolution of 50 μm.

To achieve the position resolution required, the decision was made to employ a 'Time Expansion Chamber', (TEC), design. The TEC has a low electric field drift region, covering most of the detector volume, and a high field detection gap, containing charge multiplication anode wires and sense wires. Charged particles traversing the drift region ionize the gas, producing tracks of clusters of charge. The charge clusters drift into the detection gap where the charge is multiplied and absorbed onto the sense wires. The signals from the sense wires are digitized by flash ADC's with a sampling rate of 100 MHz. The high sampling rate allows one to reconstruct the charge profile of the track in great detail. The accuracy of track reconstruction is limited by fluctuations in the charge clusters during diffusion across the drift region. The chamber is segmented so that the drift distance never exceeds 6 cm. The expected position resolution is 30 μm.

6.3 The Electromagnetic Calorimeter

The L3 BGO electromagnetic calorimeter has the following design criteria:

- Electron and photon energy resolution of 1% for energies greater than a few GeV, and 5% at 100 MeV.

Material	BGO	NaI(Tl)	CsI(Tl)	BaF_2
Radiation Length (cm)	1.12	2.59	1.86	2.1
Density (g/cm^3)	7.13	3.7	4.51	4.9
Photons/GeV	10^6	10^7	5×10^6	10^6

Table 6.1: Properties of various scintillators

- Angular resolution of 1° for electrons and photons above 1 GeV.
- Good separation of photon and electron showers in jets.
- Pion/electron identification at the level of 10^{-3} for showers with energy greater than 1 GeV.

There are significant constraints on the design of the calorimeter due to the overall design of the L3 experiment. The concept of the 'magnetic cave' with muon chambers completely outside the calorimeters requires compact calorimeters. This limits the choice of scintillating materials, and translates the angular resolution requirement into a rather strict position resolution requirement. Additionally, the cylindrical vertex chamber forces the EC to also have a cylindrical shape, whereas a spherical shape would be preferred.

The vertex chamber is a cylinder 1 meter long and 1 meter in diameter. To make the design as compact as possible, the EC is designed as a cylinder, with 1 meter inner diameter, and two end caps, covering the end of the vertex chamber. Bismuth Germanate ($Bi_4Ge_3O_{12}$), commonly referred to as BGO, was chosen because it has the shortest radiation length of the available scintillators, (see Table 6.1), which leads to the most compact total absorption calorimeter. The short radiation length is also important because it reduces the lateral shower spread, making good angular resolution possible even though the calorimeter is only 50 cm from the interaction point.

To achieve the required resolution the crystals must be long enough to completely contain the shower. Monte Carlo calculations have shown that 22 radiation lengths of BGO are sufficient to keep the fluctuations in deposited energy well below the 1% level for energies up to 100 GeV. The overall shape of the BGO must therefore be a cylinder of 1 m inner diameter with at least 24 cm of BGO along rays originating at the interaction point. The only remaining question is how to segment the BGO, or equivalently how large to make the BGO crystals.

The extent of lateral segmentation of the calorimeter is a trade off between the granularity required for accurate position measurements and the effective increase

in electronic noise. To find the energy of a shower, we sum over all electronic channels covered by the shower. For energies, below a few hundred MeV, where the electronic noise makes a significant contribution to the resolution, we need to sum over channels within about 5 cm of the center of the shower in order to reduce fluctuations in the deposited energy to the level of a few percent. Increasing the lateral granularity increases the number of channels in the sum, thereby increasing the noise in the signal. The electronic readout, discussed below, has an equivalent noise of about one MeV per channel. In the best case, completely uncorrelated noise, the total noise is the quadrature sum of the noise of each channel. If we sum over all channels within 5 cm of the shower center, then the electronic noise is

$$\Delta E^2 \sim (1.0 \text{ MeV}^2/\text{channel})(\frac{\pi r^2}{d^2} \text{ channels}) \quad \Rightarrow \quad \Delta E \sim 1.8\frac{r}{d} \text{ MeV}, \qquad (6.1)$$

where $r = 5$ cm is the extent of the shower and d is the lateral dimension of the BGO crystals. To achieve 5% resolution at 100 MeV, we must have $d >$ 1.8 cm. It has been found, from Monte Carlo studies and test beam data, that crystals which are 2.0 cm $\times$ 2.0 cm at the inner surface of the detector give a position resolution of 1 mm at 10 GeV and 3 mm at 1 GeV. This is sufficient for the required angular resolution. Roughly 11000 crystals of this size are required to completely cover the inner edge of the calorimeter. This is an economically and technically feasible number of electronics channels, and the required crystals are small enough that they can be reliably grown in the quantities required.

6.4 The BGO Crystals and the Mechanical structure

The calorimeter has four mechanically seperate pieces: two half barrels and two end caps. The end caps have not yet been constructed and will not be discussed here. Two orthogonal cross sectional views of the barrel portion of the calorimeter are shown in Fig. 6.3. Each half barrel consists of 160 rows of crystals, symmteric in ϕ. Each ϕ row consists of 24 BGO crystals. The crystals are truncated pyramids. All crystals have a 2 cm $\times$ 2 cm front face, are 24 cm long, and point directly at the interaction point. The dimensions of the rear faces are adjusted to minimize the gaps between the crystals, so the shape and size of the crystals vary with the position of the crystal in θ. The calorimeter is symmetric about the beam axis, so there is no variation in the crystal dimensions as a function of ϕ. The crystals are arranged to be as close to the interaction point as possible, within the space allotted for the electromagnetic calorimeter. As can be see from Fig. 6.3, the innermost edge of each crystal lies on a cylinder with radius 52 cm.

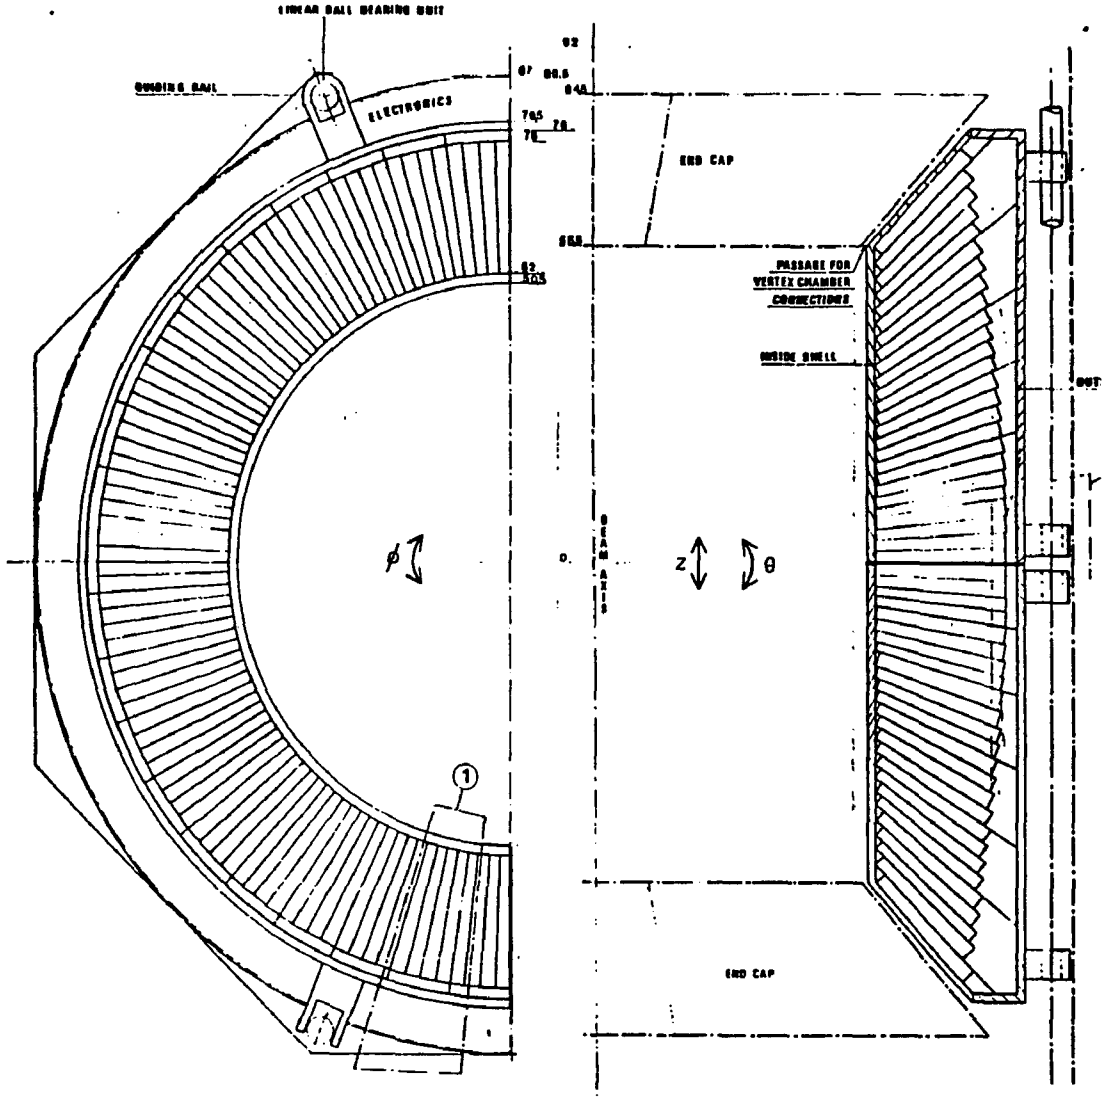

Figure 6.3: The Electromagnetic Calorimeter

The mechanical structure of the calorimeter consists of individual cells, referred to as alveoli, which containing the crystals and a supporting structure. Both parts of the structure are made of epoxy-resin carbon-fiber composite material. The alveoli hold each crystal separately such that the weight of any crystal is transmitted only to the carbon fiber material and not to any of the surrounding crystals. Each crystal is pressed into its alveola by a spring loaded device at its back (the end further from the interaction point). The alveolar structure is attached to the supporting structure made of two roughly cylindrical shells with conical supporting flanges at the ends of the barrel. The part of the supporting structure between the crystals and the beam is 10 mm thick. The material transversed by particles originating at the interaction point varies between 0.04 and 0.09 radiation lengths.

The BGO crystals were grown, cut to size, and polished at the Shanghai Institute of Ceramics in China. The crystals were manufactured by first heating a mixture of pure Bi_2O_3 and GeO_2 powders to 1000°C, and then, after the powders have melted, allowing the material to cool slowly as the crystal grows from a monocrystal seed. The crystals were cut to size and then polished to tolerance of 300 μm in the transverse dimensions, 400 μm in length, and 50 μm in planarity of the faces. The close mechanical tolerances allow the calorimeter to be constructed with a minimum of dead space between the crystals.

The crystals are required to have a minimum optical transparency of 55% at 400 nm and 65% at 630 nm. The average optical transparency of the crystals is 10% above the minimum values. The photodiodes were attached to the crystals with an epoxy glue with index of refraction of 1.55. A number of different surface treatments for the crystals were investigated to maximize the total light output and the uniformity of the light output across the length of the crystal, *e.g.* polishing a number of the faces, depolishing certain faces, painting the crystals, wrapping the crystals in mylar foil, etc. The method chosen was to optically polish the crystals on all six faces and, after the photodiode was attached, paint all faces with a 30 to 80 μm thick coat of high reflectivity NE560 paint. The paint, manufactured by Nuclear Enterprise of Scotland, is water-based with Titanium oxide as the reflective compound. It provides a diffuse reflection of the photons which are not internally reflected at the faces of the crystal. Crystals prepared in this way have a total light output comparable to any of the surface preparation techniques investigated and variations in the light output along the length of the crystal of less than 10%, on average.

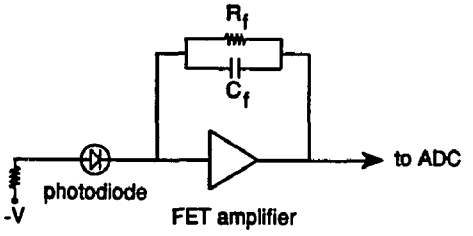

Figure 6.4: The preamplifier

6.5 Analog electronics

The front end electronics for the calorimeter measures the light output of each of 11000 BGO crystals. We have chosen to use a system with photodiodes for light collection, a low noise charge sensitive preamplifier mounted a few centimeters away from the photodiodes, and a wide dynamic range ADC with dual resettable integrators mounted a few meters away from the preamplifiers. The motivation and design of the system is discussed below. The system must meet three basic criteria:

- The system must operate in a 0.5 Tesla magnetic field.
- The electronic noise must be less than 1 MeV per channel.
- The electronics must be able to handle signals over a range from a few MeV to one hundred GeV.

The light output at the outer end of the BGO crystals is a few thousand scintillation photons per incident MeV. The most common choice of light collection device for low light levels is the photomultiplier tube. However, the requirement that the system operate in a magnetic field excludes the use of conventional phototubes. The next alternative is a solid state light collection device, *e.g.* a photodiode or phototransistor. We have chosen to use a silicon photodiode, the Hamamatsu S-2262. This photodiode has a spectral range relatively well matched to the scintillation light of BGO, and has low capacitance, which is important in the optimization of the low noise preamplifier. There are two photodiodes, each with 1.5 cm^2 active area, glued to the rear of each crystal.

The disadvantage of solid state light collection devices is that they have no internal amplification. Because the photodiodes cover only a third of the back face

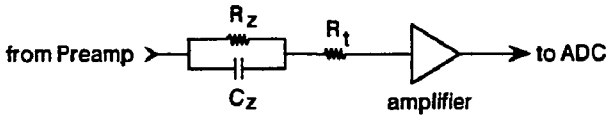

Figure 6.5: A pole-zero

of the crystal and have a quantum efficiency of 60% for BGO scintillation light, approximately one quarter of the photons leaving the backface of the crystal are converted to photoelectrons. The sum of the signal from the two photodiodes in the final system is about one thousand electrons per MeV. This necessitates the use of a low noise preamplifier. The preamplifier used was designed at Lyon[46]. The preamps are mounted a few centimeters away from the photodiodes in order to minimize the cabling and degradation of signal due to transmission in the cables. There are 24 preamps mounted on one preamp card for each ϕ row of a calorimeter half barrel.

A simplified schematic of the preamp is shown in Fig. 6.4. The active component of the preamp is a low noise FET (Field Effect Transistor) current amplifier. The preamp is charge sensitive and produces a output voltage pulse with an amplitude equal to the charge collected from the two photodiodes divided by the feedback capacitance, $C_f = 4pF$. The feedback capacitance was chosen to produce a 10 volt peak, the largest signal the ADC system can handle, for the charge produced by the photodiodes when 100 GeV of energy is deposited in a crystal. The pulse at the output of the preamp has a rise time of a few hundred nanoseconds, corresponding to the decay time of BGO scintillation light, and a decay time constant, determined by the product of the feedback capacitance, C_f, and resistance, R_f, which is roughly 800 μsec. The long time constant is due to the large resistance of the feedback resistor, $R_f = 200\text{M}\Omega$. Thermal fluctuations in the feedback resistor contribute noise which is inversely proportional to the resistance. Therefore, it is advantageous to maximize the feedback resistance.

The output pulse from the preamp is fed into the ADC circuit, designed at Princeton[47]. There are two ADC boards, each with 12 ADC channels, for each 24 channel preamplifier board. The ADC boards are mounted just outside the hadron calorimeter, a few meters away from the preamplifiers. The first stage of the ADC electronics is the 'pole-zero'. The essential components of the pole-zero circuit are shown in Fig. 6.5. As discussed above, the pulse produced by the preamp has an

800 μsec tail. However, beam crossings at LEP occur every 11 μsec. The pole-zero converts the pulse from the preamp, with its 800 μsec tail, into a new pulse with a decay time constant of 1.0 μsec. The pole-zero is able to eliminate the 800 μsec tail because the tail is known to be a simple exponential decay. The circuit works as follows.

The leading edge of the preamp pulse deposits a charge on the pole-zero capacitor, $C_z = 680\ pF$, matching the voltage on C_z to the peak voltage of the pulse from the preamp. There are two ways that charge can be removed from the capacitor. Either the charge can move across the pole-zero resistor, $R_z \sim 1\text{M}\Omega$, or it can move through the resistor $R_t \sim 1\text{k}\Omega$ into the amplifier. The charge will, therefore, decay with two time constants. One time constant, corresponding to charge moving into the amplifier, is given by R_tC_z. The other time constant, corresponding to charge moving across R_z, is given by the R_zC_z. The pole-zero resistor, R_z, of each ADC channel is adjusted so that time constant R_zC_z matches the time constant of the attached preamp. When the two time constants are matched, the exponential decay of charge on the C_z will cancel the exponential tail of the preamp output pulse. The amplitude of the 800 μsec tail at the output of the pole-zero is equal to the amplitude of the pulse from the preamp multiplied by

$$\frac{R_t}{R_z}\left(\frac{\tau - R_zC_z}{\tau}\right) \tag{6.2}$$

where τ is the time constant of the pulse from the preamp. In practice the time constants are matched to an accuracy of 1%, which cancels the 800 μsec tail to an accuracy of 10^{-5}. The charge leaving C_z through R_t provides an output pulse, with a time constant ($R_tC_z = 1$ μsec), which is used by the ADC.

After the pole-zero the signal is split into three parts: the low energy channel, the high energy channel, and the analog sum channel. A simplified schematic of the complete analog part of the ADC is shown in Fig. 6.6. There is a simple RC filter at the input of each of the three channels. The three filters have matched time constants of 0.5 μsec and are used to smooth the leading edge of the pole-zero output pulse, to allow the relatively slow front end amplifiers to accurately follow the shape of the pulse. This eliminates nonlinear feedback from the front end amplifier into the pole-zero.

The analog sum channel converts the signal, with an adjustable attenuation, into a current pulse. The current pulses from groups of crystals are summed to give a fast approximate determination of the energy in the calorimeter used for triggering. The attenuation is adjustable for each crystal to compensate for variations in light output and preamplifier gain of the different crystals.

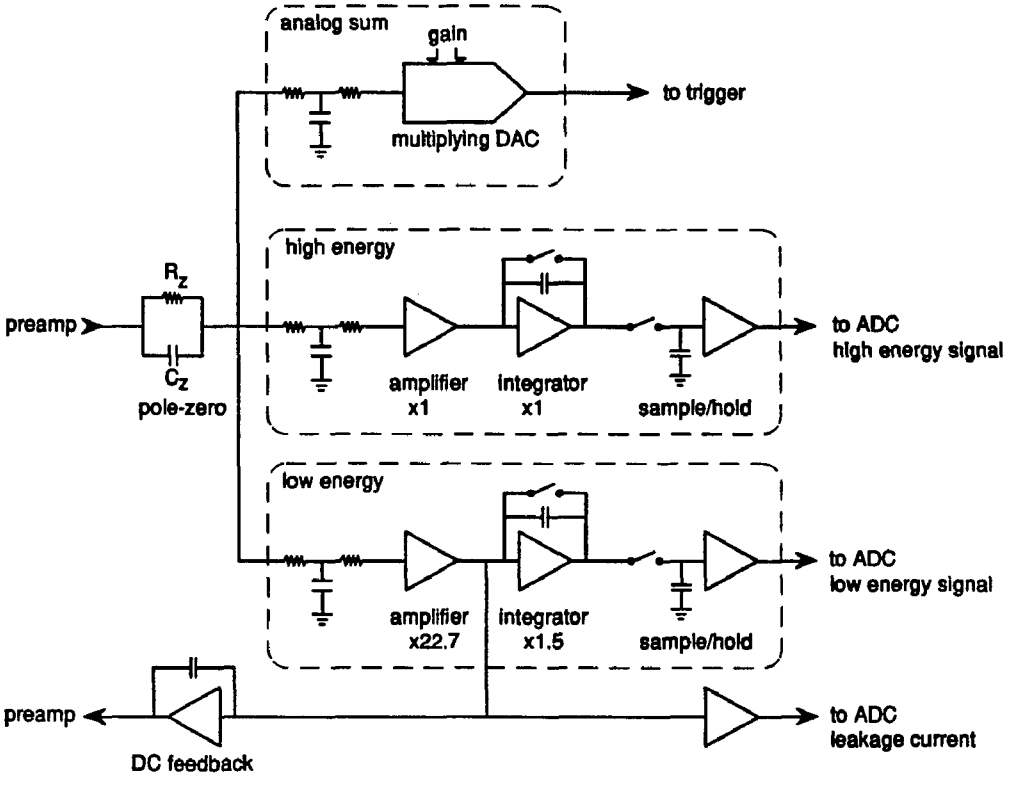

Figure 6.6: Analog part of the ADC

The low energy channel and the high energy channel are two separate circuits used to achieve the wide dynamic range required. The low energy channel is optimized for small signals. It has a front end amplifier with a gain of 22.7 and an additional gain, relative to the high energy channel, of 1.5 due to the integrator. This gives a total relative gain of 34. Low noise often implies slow speed. The low energy amplifier is a very low noise op-amp which is too slow to follow large amplitude signals. There is circuitry, not shown in Fig. 6.6, which limits the gain of the op-amp for large signals and additional circuitry which limits the amplitude of signals in the low energy integrator. These circuits are necessary to prevent the low energy channel from saturating on large signals, but preclude the use of the channel at high energies, so a separate high energy channel is required.

The op-amp labelled 'DC feedback' averages the signal from the preamp. The op-amp generates a slowly varying voltage, which is an average of the signal from the preamp over an interval of about one second. The voltage measured is the DC offset at the output of the photodiode, which is determined, primarily, by the photodiode leakage current. The voltage is inverted and fed back into the preamp to compensate for variations in the leakage current and stabilize the preamp. It is also sent to the ADC to provide a monitor of the photodiode leakage current. This is useful when diagnosing problems in faulty channels.

Before digitization, the voltage pulses from the front end amplifiers are converted to a DC voltages. A rough sketch of the path of the signal from the photodiode though either ADC channel is shown in Fig. 6.7. The charge from the photodiode on the BGO is collected by the preamp, sent through the front-end of the ADC, and fed into the integrator. The integrator and sample/hold were designed for the synchronous beam timing of LEP. Figure 6.8 shows a timing diagram of one beam crossing interval at LEP. The integrator is 'resettable'. The charge on the integrator capacitor is zeroed, before each beam crossing, by closing the (FET) switch across the capacitor. After the charge is zeroed, the switch is opened, and the integrator begins to build up charge. At this point, the sample/hold switch is closed, and the voltage on the sample/hold capacitor tracks the output voltage of the integrator. The switch remains in sample for 5 μsec after the beam crossing. The pulse at the input of the integrator has a time constant of 1 μsec. After 5 μsec, more than 99% of the pulse has been integrated. More importantly, after 5 μsec the pulse has decayed sufficiently that variations in the integration time or time constant of the pulse result in changes in the output voltage of less than 0.1%. The signal is held until just before the next beam crossing. The decision on whether or not to accept the event must be made by that time. If the event is rejected, the cycle repeats, if the event is accepted the signal is held until digitization in complete.

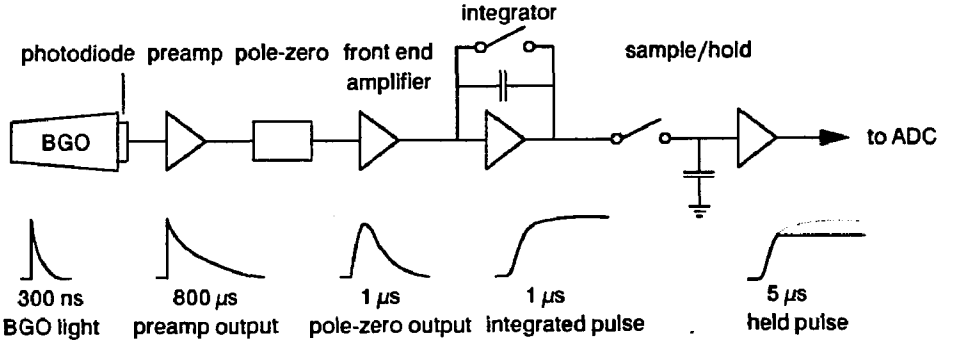

Figure 6.7: Path of the signal before digitization

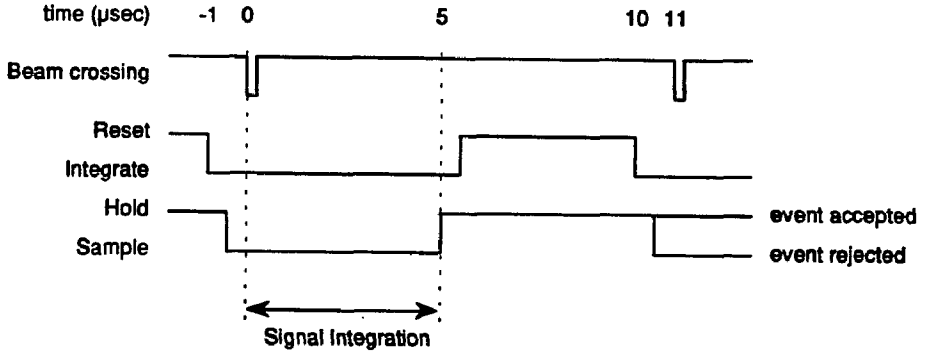

Figure 6.8: ADC timing at LEP

The energy deposited in each crystal can range from a few MeV to two hundred GeV. To prevent digitization errors, the least count of the ADC is roughly one tenth of an MeV, so the required ADC range is 21 bits. However, 0.1 MeV precision is not required for most signals. Digitization need never be more accurate than 1/1000th of the input signal. The wide dynamic range is achieved by performing a 12-bit digitization of the signal after it is amplified with a selectable gain, varying from 1 to 544. The 12-bit digitization and the 9-bit range of amplification combine to give a 21-bit dynamic range.

The outputs of the two sample/holds feed into ADC comparator/amplifier chain as shown in Fig. 6.9. The two amplifier/comparator chains present a total of six levels of amplification of the input signal. The input voltage scale for each comparator, and an approximate equivalent energy scale are presented in Table 6.2. There is a factor of four between successive comparators, except at the crossover between the high and low energy chains. The first step in digitization is to choose which comparator to use. The best comparator is the one with the highest input voltage which is below saturation. To ensure that there are the voltage scale is linear for all readings, we define 'saturation' to mean that the input voltage is greater than 7/8 of full scale of the DAC. After the most sensitive comparator with a reading of less than 7/8 of full scale is found, digitization proceeds with a conventional successive approximation algorithm on that comparator. Since the DAC used for the conversion has 12 significant bits, the gain of four between successive comparators ensures that the digitization is accurate to at least 10 bits on all comparators except on comparator one. On this comparator, the least count is 0.1 MeV, so we are assured that the conversion is accurate to either 0.1 MeVor 0.1% of the input signal, which ever is greater.

There are two additional comparators in the ADC circuit. One is used to digitize the photodiode leakage current measured by the DC feedback circuit. The other comparator allows us to measure current produced by temperature sensors mounted on the BGO, to measure the power supply voltages, and to identify each individual ADC card. Each of the 12 channels on an ADC card measures a different quantity, as shown in Table 6.3. Channels 1, 2, 11, and 12 are used to monitor the power supply voltages on the ADC cards. Channels 3-4 and 5-7 measure the voltage across identification resistors mounted on the preamplifier card or the ADC card, respectively. The identification resistors are soldered in by hand after the boards are assembled. The three resistors for each ADC board and two resistors for each preamp board uniquely identify each circuit board.[1]. Channels 8-10 read a voltage, produced by an AD-590 thermosensor, which is a linear function of the

[1]For the preamps there are two indentical pairs of resistors (one pair for each of the two attached ADC cards) and the identification is unique only within each half barrel of the detector

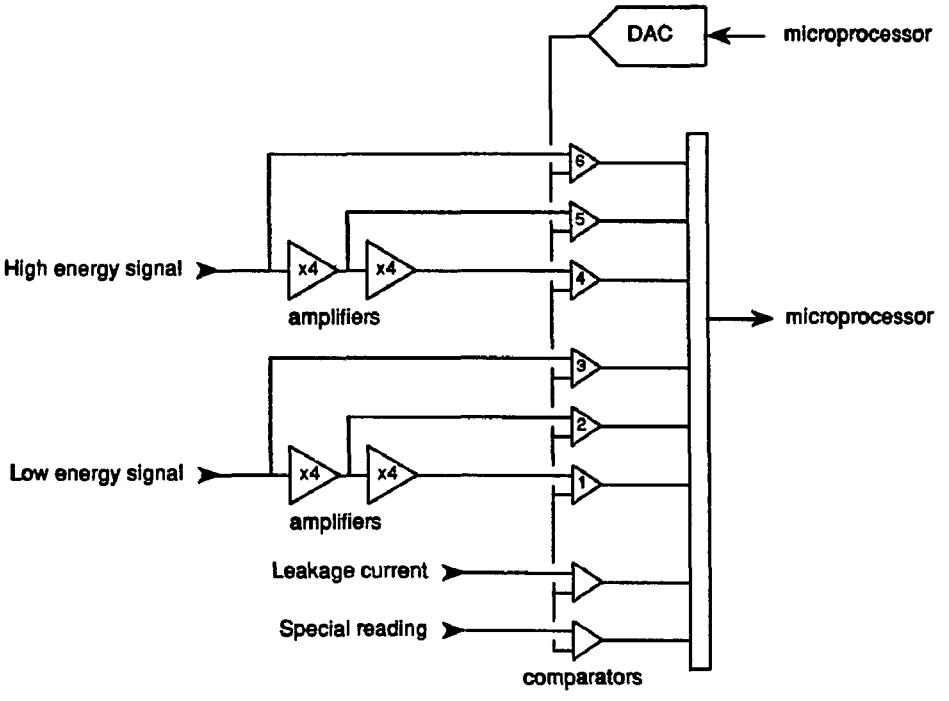

Figure 6.9: The ADC amplifier/comparator chain

	Voltage		Energy	
Comparator	Least count	Full Scale	Least count	Full Scale
6	2.4 mV	10 V	50 MeV	200 GeV
5	0.6 mV	2.5 V	12 MeV	50 GeV
4	0.15 mV	0.6 V	3 MeV	12 GeV
3	71 μV	290 mV	1.5 MeV	6 GeV
2	18 μV	74 mV	0.4 MeV	1.5 GeV
1	4 μV	18 mv	0.1 MeV	0.4 GeV

Table 6.2: The ADC comparators

Channel	Voltage digitized
1	+15 Volts monitor
2	-12 Volts monitor
3	Preamp identification
4	Preamp identification
5	ADC board identification
6	ADC board identification
7	ADC board identification
8	Temperature sensor
9	Temperature sensor
10	Temperature sensor
11	-15 Volts monitor
12	+5 Volts monitor

Table 6.3: Signals on the 'special' comparator

temperature of the sensor. The sensors are mounted on the BGO crystals and on the supporting carbon fiber structure. The temperature sensors will be discussed in detail in section 7.6.

The digitization is controlled by a single chip microcomputer, a HD63B05VOP manufactured by Hitachi. There is one microcomputer for each crystal. In addition to controlling the ADC, the microprocessor also controls the other programmable parts of the analog electronics, *e.g.* the analog sum attenuation, and interfaces with the data collection system. Additionally, the 6305 is capable of performing tests of a significant part of the ADC electronics and, in conjunction with the upper levels of the readout, most of the data collection system.

Except for the analog sum, which is sent to the trigger system, no analog signals pass beyond the ADC cards. The all digital character of the readout minimizes the necessary cabling and eliminates any degradation of the signals during transmission.

6.6 Data Collection

After the light generated in each of the 11000 BGO crystals is measured, the digitized signals must be combined into a single data stream. The large number of crystals and the corresponding fan-in of data naturally leads to a hierarchical scheme, with data collection performed at several tiers. Levels 1 through 3 of the

system were designed at Princeton and level-4 was designed at Nymegen. A decision was made early on in the design of the system to incorporate processing capabilities at each stage of the readout, in order to maximize the flexibility of the system and to reduce the overall data transfer rate via suppression of signals below threshold.

The system has four levels of readout, as shown in Fig. 6.10. The lowest level, the ADC cards, reside within the magnetic cave of the L3 detector and are located just outside the hadron calorimeter endcaps. The ADC for each crystal is controlled by a single chip microcomputer, which handles the A to D conversion, and interfaces with the next level of the readout system. The upper levels of the readout reside in the counting room, a hundred meters away from the calorimeter. Each of the upper levels consists of a number of (relatively) high performance microcomputers. The level-2 microcomputers monitor each trigger in order to verify that the level-1 processors are functioning properly. Levels 3 and 4 operate asynchronously, and are driven by the data coming up from the lower levels. In this section we present a brief overview of the readout system. A more detailed discussion is reserved for Appendix B.

The lowest level, level-1, consists of the 6305 microcomputers controlling the ADC for each crystal in the detector. The 6305's are organized into 'token passing rings' of 60 crystals each. All the processors in a ring share common data bus and control lines. 'Token passing' refers to a system by which control of the common lines is transferred from processor to processor. The details of the system are described in Appendix B. The token passing system uses relatively slow signals (no transition need be faster than 1 μsec) and a very small number of signal lines (four plus the data lines). This makes it extremely reliable and tolerant of noise on the signal lines. It provides an adequate data transfer rate with a bare minimum of hardware. Since each level-1 must access the common lines sequentially the time required to read out the entire ring is proportional to the size of the ring. The chosen size, 60, limits the total number of rings, and therefore the number of level-2 processors, to a reasonable number without degrading the overall performance of the system.

After an event is selected by the trigger, the signal for each crystal is digitized. The digitization requires 220 μsec. The digitized signal is then stored in a small buffer (41 readings) in the 6305. Moving the data to the upper levels of the readout is performed as a background task. If the reading exceeds a preset threshold then the data is passed up to a level-2 microcomputer, otherwise the data is discarded.

Each ring of 60 crystals is controlled by a level-2 processor. The level-2 processor downloads adjustable parameters, such as the analog sum attenuation and the zero suppression threshold, controls when special events, such as pedestals or temperatures, are taken, and handles all the data produced by 6305's in the ring. The level-2 microcomputers are commercially built Mizar 8115 single board micro-

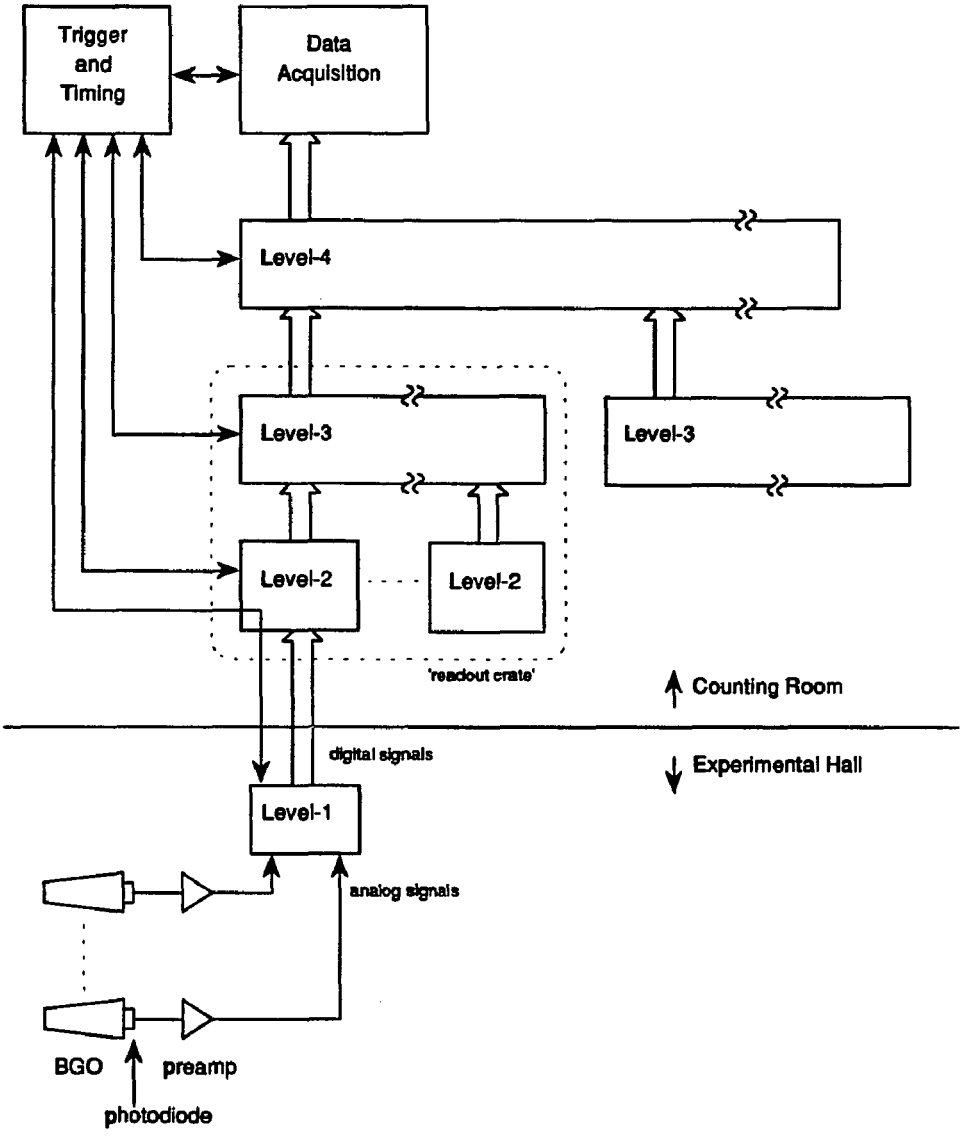

Figure 6.10: The Readout System

computers based on the Motorola MC68010 microprocessor. Each Mizar board is attached to a Princeton built interface board to allow it to communicate with its associated level-1 microcomputers. A pair of interface boards, one attached to the level-2 Mizar computer, the other attached to the ring of five ADC cards, translate the 32 digital signals required for communication to differential TTL levels which are transmitted over more than 100 meters of multiconductor cable.

Each Mizar computer has large (512k byte) memory capable of storing at least several hundred events. Incoming data from the 6305's are reformatted and then stored in this memory. The level-2 processors are organized in crates of 16. Each crate is controlled by a level-3 processor, also a Mizar 8115. The sixteen level-2's and the level-3 communicate via a digital signal bus (VME bus) carried on the backplane of the crate. The memory on the Mizar boards can be accessed either by the board's microprocessor or by the level-3 processor via the VME bus.

When all the level-2's in a crate have assembled a complete event in their memory, the level-3 organizes the data and sends it via a First-In-First-Out buffer to level-4. Since the level-3 can access the level-2 memory directly, no recopying of the data is done. The level-3 does not buffer events, but instead transmits then directly to level-4. We have found that eliminating buffering by the level-3 processors greatly increases the speed of the system. Level-4 collects data from all the level-3's and passes the complete event to the data acquisition system.

The major part of the dead time comes from the time required for the digitization. Each 6305 requires 220 μsec after an event is selected to digitize the signal and store the reading in its internal buffer. Since all events are buffered, no further processing need be done immediately. Therefore, the maximum instantaneous rate of the system can be as high as 4 kHz for bursts of fewer than 41 events (the size of the 6305 buffer). The maximum average event rate depends on the fraction of channels with nonzero signals. We will calculate the event rate for two cases: 10% occupancy, the level expected during normal running at LEP, and 100% occupancy, as used in pedestal and certain types of calibration runs.

Suppression of zero signals occurs in the lowest level of the system. After the reading is saved in the buffer, but before it can be passed to the level-2 processor, 300 μsec of processing are required to format the reading (add error checking bits) and decide if the reading should be zero suppressed. If there are multiple events in the buffer, these operations are done while a different 6305 in the ring is transferring data, so no additional time is required.[2] The operation which must be done sequentially is communication with the level-2 processor. Each 6305 requires 86 μsec to pass up a single reading, or 10 μsec to indicate that it has no data, and

[2]There will be a single event in the buffer only if the event rate is low enough to allow the extra processing time.

Transfer	Rate μsec/byte	Fan-in	10% occupancy Data bytes	10% occupancy Time μsec	100% occupancy Data bytes	100% occupancy Time μsec
$1 \rightarrow 2$	see text	60	18	1750	180	6000
$2 \rightarrow 3$	-	16	40	-	256	-
$3 \rightarrow 4$	1.65	11	652	1100	4100	6800
$4 \rightarrow$ DAQ	0.08	1	7200	600	45000	3800

Table 6.4: Data transfer rate

8 μsec, after all the transfers are completed, to 'clear the token' and return to the quiescent state. The total data handling time is 220 μsec for digitization and an additional 94 μsec/processor for each 6305 with data and 18 μsec/processor for each 6305 without data.

The processing performed by the level-2's is done while waiting for the 6305's. The level-2 reformats the data as it is read in, and does an integrity check of the complete event during the 480 μsec taken to 'clear the token'. Since the level-3 can read the event directly from the level-2 memory, no time is used to transfer the event to the level-3. When all the level-2's in a crate have assembled complete events, the level-3 adds a few more error checking and status words and transfers the data to level-4.

The data transfer rates, and quantity of data handled by individual processors at each level are shown in Table 6.4. It is not useful to measure the data transfer rate between levels 1 and 2 in bytes/sec because time is required to process crystals which have signals below the suppression threshold, and therefore transfer no data. The total time for the transfer, as explained previously, is $220 + 94n + 18(60 - n)$ μsec, where n is the number of crystals in the ring with data, $60 - n$ is the number of crystals in the ring without data. The average transfer rate when all crystals have data is 100 μsec/crystal. However, when only some of the crystals have data, the average transfer rate per byte of data is significantly lower due to the fixed time required for digitization and to process crystals without data.

The bottleneck for normal data is the level-1 processing time, and for full calorimeter data is the level-3 to level-4 transfer rate. The maximum average event rate of the readout system is 570 Hz for 10% occupancy and 170 Hz for 100% occupancy. The data acquisition rate, for the conservative estimate of 10% occupancy, is within the target data taking rate for L3.

We have just described the data collection hierarchy. The arrangement for timing and control is somewhat different. The data from the ADC cards is sent up through

levels 2,3, and 4 to the data acquisition system. The signals that control the timing and control signals are connected directly between each level and the main trigger system, for all levels except level-1. The signals for level-1 are passed through the corresponding level-2 driver/receiver board. This allows a level-2 to disable its level-1 cards in case of failure and also simplifies the cabling.

Chapter 7

Calibration

7.1 Introduction

The data produced by the BGO electromagnetic calorimeter for each event consists of an ADC reading from each crystal that has a signal which exceeds its zero suppression threshold. From this set of data, we wish to reconstruct the energy and impact position of the particles traversing the detector in that event. To accomplish this task, we must understand the response of the system. We must know the dependence of the energy deposited on the impact position and type of the particle, the signal output for a given energy input, and the variation of the signal with environmental factors, primarily the temperature of the BGO.

We have studied the response of the calorimeter with both test beam data and computer simulations. We prefer to rely on the test beam data, using the simulations only to supply information that cannot be derived from the beam data because of the limited granularity and coverage of the detector. Since the calorimeter is primarily intended to measure electrons and photons, we use an electron beam for the calibration. The beam is sent into each crystal in turn, in order to calibrate each crystal and its associated electronics independently. To successfully calibrate a crystal, we must know the position and momentum of the incident electrons, the amount of energy deposited in the crystals, the efficiency for the generation and collection of scintillation light, and the response of the electronics for signals from the photodiode. In the following chapters we will discuss each of these topics, two algorithms for finding the calibration constants, and the performance of the detector.

Figure 7.1 is a sketch of main features of the measurement of a particle's energy using the calorimeter. A particle enters a BGO crystal, producing an electromagnetic shower. The shower spreads through the crystal and into the adjacent crystals. Some of the particles leak out the rear of the crystal, or back scatter out the front of the crystal. The fraction of shower remaining in the crystal depends on the position

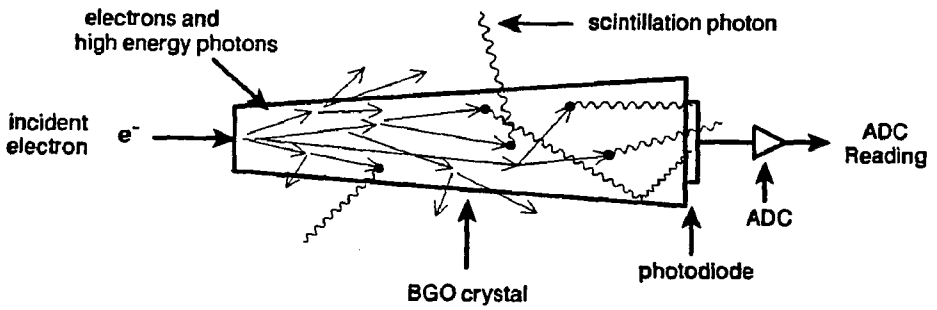

Figure 7.1: From electron to ADC

and direction and energy of the incident electron. A small fraction of the energy in the shower excites the atoms in the BGO which then produce scintillation light. The amount of light produced for a given amount of energy deposited depends on the temperature of the BGO and the quality of the crystal. This light is generated isotropically. To contribute to the signal in the ADC, the light must reach the photodiode. The probability that a given photon will reach the photodiode varies with the position within the crystal according to the geometry of the crystal, the optical properties of the crystal, and the paint on the outside of the crystal. Each photon arriving at the photodiode has some probability of producing a photoelectron. The photoelectrons are collected by the preamp and the resulting signal is digitized. Each of these processes must be understood in order to successfully use the calorimeter.

Equation (7.1) is an expression for the signal from a given crystal. The electromagnetic shower is described by $\varepsilon(\vec{x})$ which is the energy deposited as a function of position in the crystal. The function $C(T)$ is the amount of light generated per unit energy deposited. We include only the dependence of C on temperature. The light generation also depends on the quality of the BGO used to make the crystal and the mechanical stress of the crystal, but these quantities are sufficiently constant that we can neglect any variations. The light generation is a property of bulk BGO. It depends on the process used to manufacture the BGO but not on the geometry of individual crystals. Therefore, one function $C(T)$ is sufficient for all crystals. The temperature at each point is given by $T(\vec{x})$. The efficiency for collecting light produced at a given point in the crystal is described by $F(\vec{x})$. This is highly dependent on the geometry and surface treatment of each crystal. The ADC reading from the

crystal is V, and k is the calibration constant.

$$kV = \int_{vol} d\vec{x}\, \varepsilon(\vec{x}) C(T(\vec{x})) F(\vec{x}) \tag{7.1}$$

For our analysis we make some simplifications. The electron is always incident along the long axis of the crystal and the crystal is much longer than it is wide. Therefore, we average the shower development in the transverse plane and take the shower profile and light collection efficiency to be functions only of the position along the long axis of the crystal, l.

For temperature variation on the order of a few °C, the change in scintillation light generated for a given energy deposited is a proportional to the temperature change, with a coefficient of -1.55%/°C. The average temperature of the detector varied by a few degrees over the course of the calibration. It is clear that we must take into account the average temperature of each crystal. However, the effect of temperature gradients on the light output is less significant. The average gradient from the front to the back of the crystals was less than 2°C during the 1988 beam calibration. The gradient affects the light output only if the shower profile moves within the crystal. Over a range of energies from 2 to 50 GeV, the shower profile moves only a few centimeters, a small fraction of the 24 cm length of the crystal. The change in light output, due to temperature gradients, is on the order of one part per thousand. Therefore, we neglect temperature gradients and take the temperature as constant across the crystal. We define the temperature of the crystal to be the temperature at the shower maximum, $\bar{T}$. So, the temperature correction is

$$C(T(\vec{x})) \simeq C(T_0)\left[1 + c(\bar{T} - T_0)\right], \tag{7.2}$$

where $T_0 = 18$°C is the nominal BGO operating temperature. Furthermore, we absorb the coefficient $C(T_0)$ into the overall calibration constant.

The intrinsic light yield variations over the crystals is on the order of 10%, leading to a few percent change in total light output. We must make a correction for this effect. We choose to recast the light collection efficiency $F(\vec{x})$ as a overall coefficient, F_0, which we absorb into the calibration constant, and a normalized function, $f(l)$, describing the change in light yield as a function of longitudinal position along the crystal. We normalize $f(l)$ to be zero at the center of the crystal. Then we have $F(l) = F_0(1 + f(l))$ where $F_0 = F(L/2)$. Now equation 7.1 becomes

$$kV = E\left[1 + c(\bar{T} - T_0)\right]\left[1 + \frac{1}{E}\int_0^L dl\, \varepsilon(l) f(l)\right], \tag{7.3}$$

where $E = \int_0^L dl\, \varepsilon(l)$ is the total energy deposited. Note that F_0 and $C(T_0)$ have been absorbed into the calibration constant.

This is the equation that defines the calibration constant. In order to calibrate we measure the incident electron before it enters the calorimeter and calculate the shower profile from the electron's energy and impact point via a Monte Carlo simulation. The light yield as a function of position along the crystal is known from cosmic ray measurements made before the crystals were installed in the detector. The integrated charge from the photodiode and the temperature inside the detector are measured by the ADC system. We can then solve for the only unknown in the equation, the calibration constant.

To find the energy deposited in the crystal given the ADC reading we invert the previous equation,

$$E = \frac{kV}{\left[1 + c(\bar{T} - T_0)\right]\left[1 + \frac{1}{E}\int_0^L dl\, \varepsilon(l) f(l)\right]} \tag{7.4}$$

This is the equation we must use in the reconstruction of events at LEP. It is circular in that the corrections depend on shower profile which depends on the energy. Since the corrections are small and the shower profile is a slowly varying function of the of energy and impact position, we can find the shower profile using the uncorrected energy.

To demonstrate that the detector is adequately calibrated, we take data at several energies, over a variety of temperatures, and over as long a period a time as possible. We must show that the calibration is stable over time, and that the temperature and energy nonlinearity corrections are well understood.

7.2 Calibration Algorithms

The energy deposited in the crystal depends on the incident energy and impact of the position of the incident electron. The total energy of an electron with energy greater than a few tens of MeV is never deposited in a single crystal, but rather spread over a few crystals. The fraction of energy of an electron incident at the center of a given crystal deposited in squares of 1 and 9 crystals as a function of energy is shown in Fig. 7.2. The error bars on the plot represent the root mean square fluctuation in the energy deposition. The plot highlights the basic choice to be made in calibrating the crystals. We would prefer to calibrate each crystal independently. However, the fluctuations in the quantity of interest are greatly reduced if we choose to use a group of crystals to find each calibration constant. We have chosen to use both single crystal and multiple crystal methods. The major sources of error for each method are different, therefore comparison of the results will indicate the accuracy of the methods.

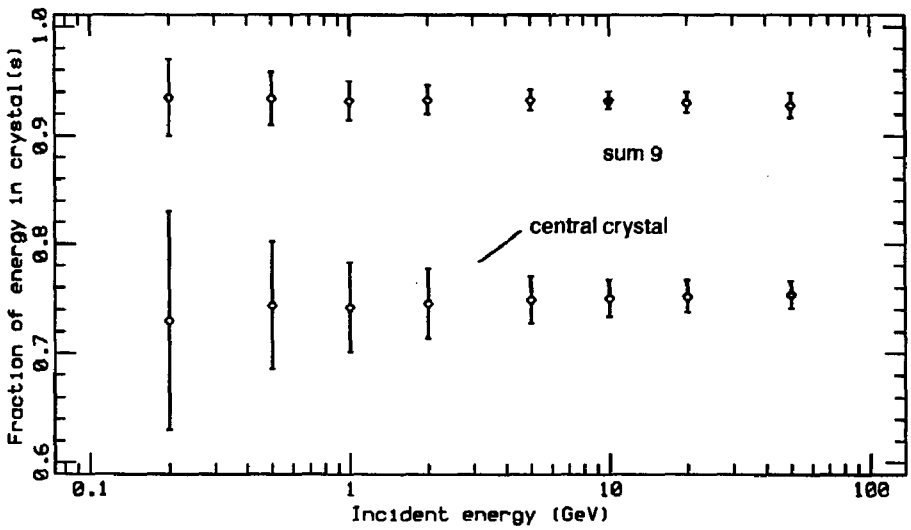

Figure 7.2: Fraction of incident energy deposited in sum of 1 and 9 crystals

The multiple crystal method is quite simple. One begins with some first guess calibration constants. Then on an event by event basis, the energy deposited in an N×N matrix is calculated from the electron's momentum and impact position. This knowledge can come from either Monte Carlo simulation of the process or from previous test beam results. The energy deposited in all crystals except the central is summed using the calibration constants from the previous iteration. This sum is subtracted from the expected total energy and divided by the reading on the central crystal. The result is a new value for the calibration constant for the central crystal. The value from each event is histogrammed. After all the events for the crystal are collected then either the mean or the fitted peak position of the histogram contents are taken as the new calibration constant. This is done for every crystal in the detector; then one begins again using the new, more accurate, calibration constants for the neighbor crystals. The entire procedure is repeated for the entire detector until we reach the desired precision of a few parts per thousand.

Because each calibration constant depends on those of the crystals around it, we must repeat this procedure several times before we converge on the correct constants. Fortunately, since a large fraction of the energy is deposited in the central crystal, the algorithm converges quickly. We expect the error on the calibration constants to be reduced by a factor of the energy in the central crystal divided by

the sum of energy in the neighbors, which is roughly a factor of four at 10 GeV, on each iteration. In practice, three or four iterations are required.

The primary advantages of the method is that, for a complete 5×5 matrix, the fluctuations in the total energy deposited are small and mean value of the total energy depends only weakly on the impact position of the electron. The energy distribution has a relatively narrow width and has a predominately gaussian shape. The gaussian shape makes locating the position of the peak computationally straight forward. The narrow width makes the determination of the position of the peak more accurate, since the statistical uncertainty in the mean of a gaussian varies as the width of the gaussian divided by the square root of the number of points in the gaussian. The weak dependence on the impact position makes the method less sensitive to errors due to a poorly centered beam. The impact position correction can actually be made by finding the impact position of the electron from a center of gravity calculation with the signals from the calorimeter themselves. Therefore, we do not need to accurately know the relative positions of the beam chambers and the BGO crystals.

There are, however, some difficulties with the multiple-crystal method. If a crystal is near a dead channel or near either end of the half barrel, then we cannot sum over the complete set of neighbors and must substitute values calculated from a Monte Carlo simulation of electrons striking the detector. While the simulation can be made relatively accurate, including the dependence on the impact point of the electron, we can only use the average value from the simulation and therefore cannot correct for shower fluctuations. Additional problems arise when we take into account the digitization of the signal. The signal in the neighboring crystals is of much lower energy and is measured by the low energy chain. As we will see later, the pedestal on the low energy chain of the ADC has large fluctuations and was not well measured during part of the data taking. Also the relative gain between the high and low energy chains, a quantity difficult to determine accurately, must be well known. Both of these difficulties must be overcome in order to use the detector and test its resolution, and they add complications to the calibration procedure.

The single-crystal method uses only signals from the high energy chain when calibrating at 10 and 50 GeV. At 2 GeV the signals are from the low energy chain, but are significantly larger than the pedestal. Thus the effect of the inaccurate pedestal subtraction is greatly reduced. However, the energy deposited in a single crystal varies on the order of a few percent for a beam spot of a few millimeters. Ideally, one would know the exact impact position of the electron and correct accordingly. Due to inaccuracies in the beam chambers and the turning table, we do not trust the accuracy of the extrapolated absolute position at the BGO. The chambers do, however, provide accurate information on the relative positions of electron impacts

for a given crystal. Our algorithm is designed to require only the relative positions of impacts. We bin each event according to its impact position. After the events for the crystal are collected, we find the mean and RMS energy for each bin and then fit a function describing the energy as a function of impact point.

$$V(x,y) = V_0(1 - \beta[(x - x_0)^2 + (y - y_0)^2]) \tag{7.5}$$

This function has four free parameters: the ADC reading for a perfectly centered electron, V_0, the x and y of the center of the crystal, x_0, y_0, and a parameter that describes how quickly the energy falls off with distance from the center of the crystal, β. The function has been found to fit Monte Carlo data and beam data very well. To find the calibration constant, we fit the same form, with $V(x,y)$ and V_0 replaced by $E(x,y)$ and E_0, to Monte Carlo data and extract the energy deposited by a centered electron, E_0. The calibration constant is the energy for a centered electron divided by the fitted ADC reading for a centered electron, $k = E_0/V_0$.

The method has no problems with dead channels, the edges of the detector, or the difficulties with the low energy chain. It can also be done in a single pass. However, difficulties arise if the beam is not adequately centered on the crystal: if there are an insufficient number of events near the center of the crystal the procedure will not work. Also, the effective width of the peak (though hard to compare directly, given our technique) is larger than the corresponding width in the multiple crystal case. This, together with the fact that we ignore events not near the center, leads to greater statistical uncertainty. Fortunately, all crystals seem to have an adequately centered beam. The effective RMS at 10 GeV is about 0.2 GeV. To have enough statistics for an accuracy of two parts per thousand we need 100 events. For each crystal we collect about 1200 events, around 500 to 700 of which are used in the fit. Therefore, we have an adequate data sample.

7.3 Energy Deposition

Of primary importance to the calibration is knowledge of the total energy deposited in each crystal in the shower and the dependence of the energy on the impact position of the electron. Most of the information can, in principle, be found using the test beam data. The only quantity not measured during the calibration is the energy that leaks out of the BGO. If the energy leakage is known (from Monte Carlo, or other experiments) and the crystals are calibrated to a sufficient accuracy, it is a simple procedure to determine the fraction of energy deposited in each crystal and how the energy deposition varies with impact position. However, if we use beam data, the accuracy of the result is dependent on our knowledge of the conditions under which the data were taken. In particular, we are sensitive to temperature

gradients and uncertainties in the light collection efficiency of the reference crystals. We therefore chose to rely on a Monte Carlo simulation. The Monte Carlo results described below were produced using the GEANT program written at CERN[57]. The production program and the simulation of the geometry of the BGO detector were written by M. Maire, of LAPP Annecy[58].

The BGO calorimeter is cylindrical on both the inside and the outside. This fact, together with the requirements that each crystal point at the interaction point and there should be the minimum possible dead space between the crystals, leads to the conclusion that the crystals can not all have the same shape. The shape of the crystals must vary with θ. A cross-sectional view of the crystals in shown in Fig. 7.3. The variation in the cross-sectional area of the crystals at a given distance from the front face is on the order of a few percent.

Figure 7.4 shows the energy deposited in a single crystal for 10 GeV electrons incident at the center of the crystal over the range of θ. The energy deposited shows variations on the order of one percent. This variation is well explained by the changing size of the crystals. Since the crystals have nearly the same shape (with one exception discussed below) the change in energy deposited should scale with the cross sectional area of the crystal. The superposed curve in Fig. 7.4 is the cross-sectional area of the crystals at the position of the shower maximum at 10 GeV scaled and shifted to match the data at $\theta = 10$ and 22. The fit is quite good, so we can generate events at only two positions in θ and calculate the energies for the other crystals. However, a difficult arises with $\theta = 1$ because crystals at θ have a significantly different shape from the others. We can not determine the energy deposited by scaling according the cross-sectional area of the crystal and must, therefore, also generate events for $\theta = 1$. The energy deposited in a 3×3 crystal matrix is far less sensitive to variations in crystals size. The energy in the sum of nine crystals shows less than 0.1% variation over the complete range of θ.

In addition to varying with the crystal shape, the energy deposited in a crystal also depends on the position of the electron's impact relative to the crystal. Figure 7.5 shows the energy in a single crystal and in a 3×3 matrix of crystals, from a Monte Carlo simulation, as a function of the square of the distance from the center of the crystal to the impact point of a 10 GeV electron. The curves are normalized so the energy for a centered electron is unity. The error bars are the statistical uncertainty, on a sample of 2000 events, for each point. The data shown are for a crystal at $\theta = 13$. The fitted curves are of the form $E/E_0 = (1 - \beta d^2)$, where d is the distance from the center of the central crystal to the impact point. The fit is extremely good. The χ^2/DoF is 0.6 for the central crystal and 0.5 for the sum of nine crystals.

Up to now we have only considered the total energy deposited in the crystals.

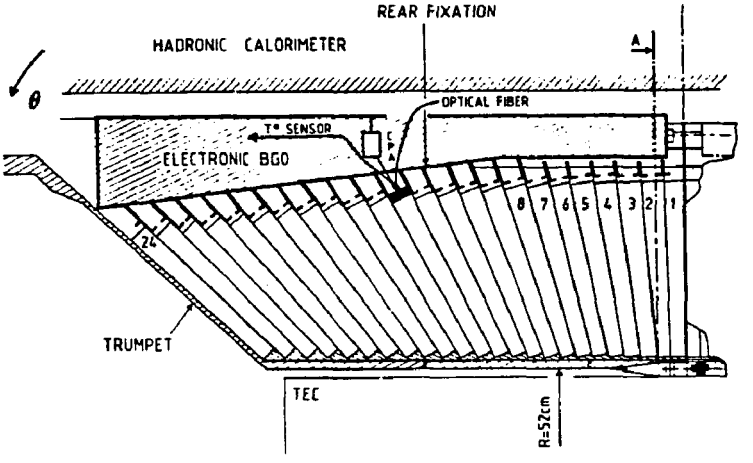

Figure 7.3: Cross sectional view of the crystals

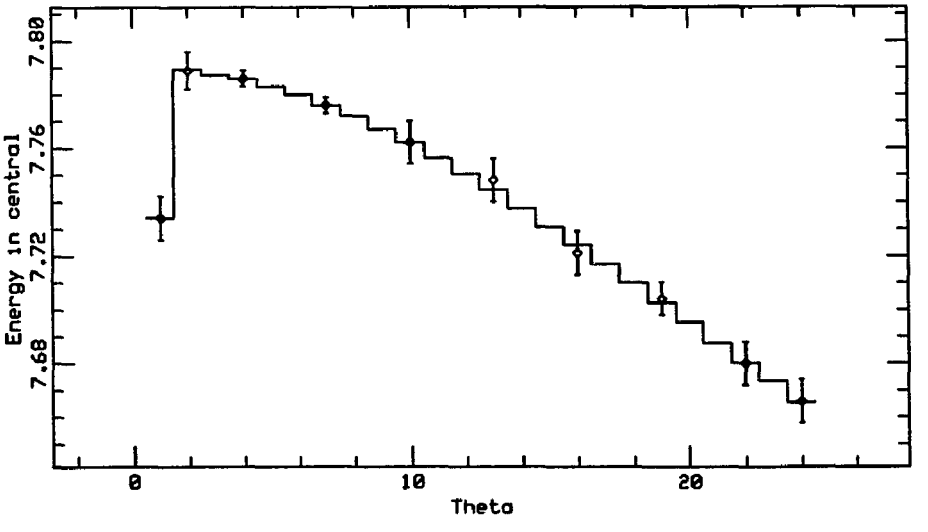

Figure 7.4: Energy in the central crystal versus θ

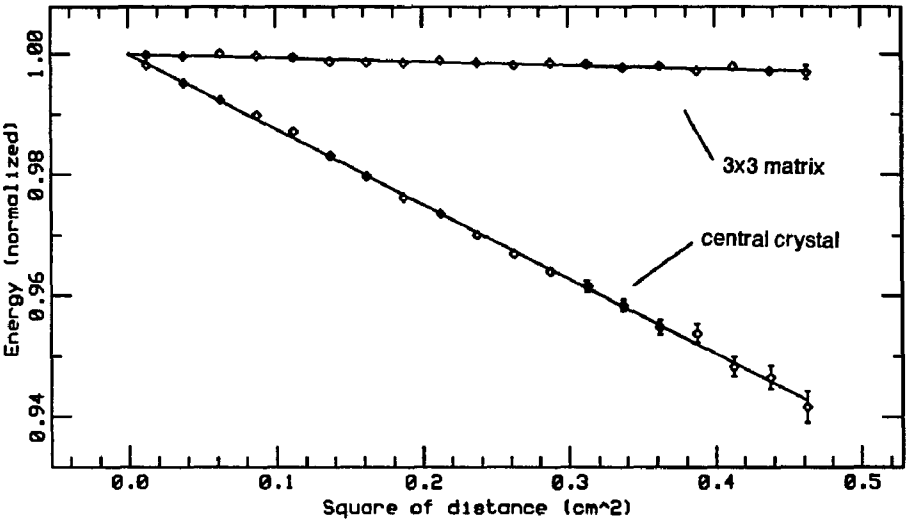

Figure 7.5: Energy deposited versus impact position

However, the amount of light collected in the photodiode per unit energy deposited varies as a function of where the energy is deposited within the crystal. The longitudinal profile of the electromagnetic shower is a function of energy of the incident electron. The position of maximum energy deposition varies logarithmically with the incident energy. Since the crystals do not have uniform light collection, we must take into account the changing shape of the shower profile when we apply calibration constants taken at one energy to different energies. In addition, the shape of the shower within a crystal varies with the position of the crystal relative to the center of the shower. The shower has an overall conical shape, so the shower maximum is closer to the rear of crystals farther from the center of the shower. The shower profile for a crystal is a function of the electron impact position and the position of the crystal within the matrix, *i.e.* whether it is the central, side, or corner crystal in a 3×3 matrix.

Due to differences in the treatment of the surface of the crystals, the light collection curve varies from crystal to crystal. The total light output from a crystal is the convolution of the electromagnetic shower profile with the light collection curve for that crystal. Since the calorimeter is not segmented longitudinally, neither quantity can be determined directly from the calibration data, so we must turn to other sources. Since there are no adequate experimental data on the longitudinal devel-

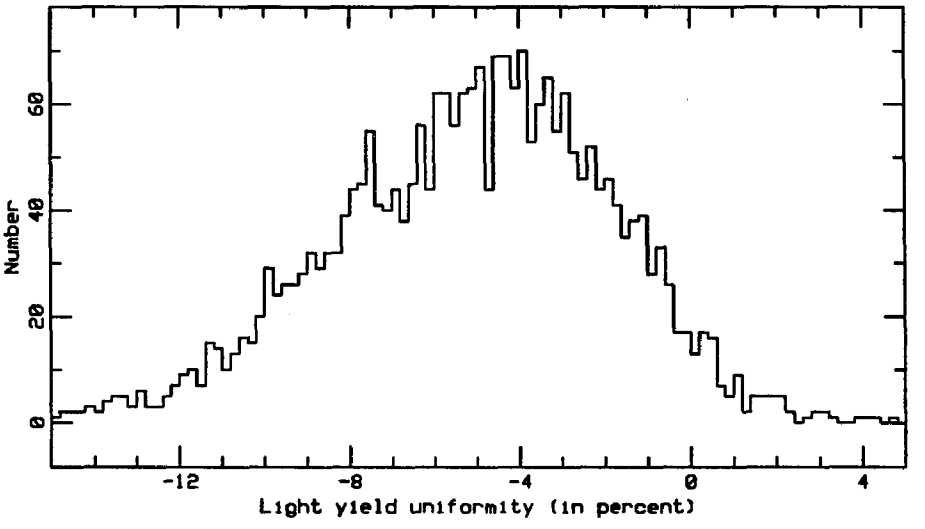

Figure 7.6: Light collection uniformities as measured by the cosmic ray bench

opment of the electromagnetic showers in BGO, we rely on Monte Carlo simulation to determine the shower profiles.

For the light collection curves, we use a set of data obtained by M. Chemarin, of Lyon, with a 'cosmic ray bench'[59]. The cosmic ray bench consists of a set of wire chambers and several scintillators for triggering. The crystals are placed between the wire chambers and connected to an ADC system to read the light output from each crystal. The crystals have their final photodiodes attached, but the ADC system is not the one used in the final readout system. Each time the scintillators indicate the passage of a cosmic ray muon through the setup, the track of the muon is measured simultaneously with the light output from the BGO crystals. The track length of the muon inside the struck crystal and the position of incidence along the crystal are calculated. The energy deposited is proportional to this length. The data were the analysed to find the light output per unit energy deposited as a function of longitudinal position along the crystal.

The light output curves have been parameterized in two different ways. The simpler parameterization assumes that the light output varies linearly with position. A linear variation is expected from the tapered shape of the crystals and is a fairly good approximation to most of the measured light yield curves. Since the normalization of the curve is absorbed in the calibration constant, a single param-

eter suffices to describe the light collection efficiency. For historical reasons, the parameter used is the fractional change in light output between the two ends of the crystal, $R = [s(0) - s(L)]/s(L)$, where $s(x)$ is the relative light output at a distance x from the front face of the crystal and L is the length of the crystals. The parameter, R, is referred to as the light yield uniformity of the crystal. A plot of the uniformities for all of the crystals of the second half barrel, as measured by the cosmic bench, is shown in Fig. 7.6. The second parameterization is a third order polynomial fit to the light output per unit energy versus position. This reproduces some of the nonlinear features of the uniformity curves of some of the crystals.

Figure 7.7 shows electromagnetic shower profiles in the central crystals and one side crystal of a 3×3 matrix for centered electrons, (solid lines), and electrons offset by 4 mm in θ and ϕ, (dotted lines). The two shower profiles for the central crystal are almost identical; one can see the dotted line for the shifted impact point only near the peak. For the side crystal, there is a larger, but still quite small, shift. However, there is a significant shift in the shower profile between the central and the side crystals. Since the change in shower profile due to the shifted impact position is significantly smaller than the change as we move from crystal to crystal in the matrix, we chose to examine the effect of the light yield nonuniformity only as a function of position in the matrix and neglect any effect due to shifts in the impact position.

We have calculated the relative light output per unit energy deposited for the central, side, and corner crystals of a 3×3 matrix, for a range of linear light uniformity curves. The result is shown in Fig. 7.8. The shifts are on the order of a few percent. The effect of the shift on the sum of nine crystals is weighted by the relative fraction of energy deposited in the central and the surrounding crystals. Roughly 20% of the energy is deposited in the surrounding crystals, therefore a shift of (in the worst case) 5% produces a shift of 1% in the mean of the sum of nine crystals.

To calculate the energy in 3×3 matrices for the calibration, we have chosen to apply a correction, based on the light yield uniformity of each crystal, to the energy deposited in the side and corner crystals. We have adopted the calibration data at 10 GeV as the reference. The correction is defined to be unity for the central crystal at 10 GeV, since each crystal is the central crystal while it is being calibrated. The correction is calculated assuming a linear light yield uniformity curve using the average shower profile in each crystal in a 3×3 matrix centered at $\theta = 13$. The correction, c_{ly}, is given by

$$c_{ly}(R, d\theta, d\phi, E_{in}) = \frac{E'}{E} = \frac{1 + R\xi(d\theta, d\phi, E_i)}{1 + R\xi(0, 0, 10\ \text{GeV})}, \tag{7.6}$$

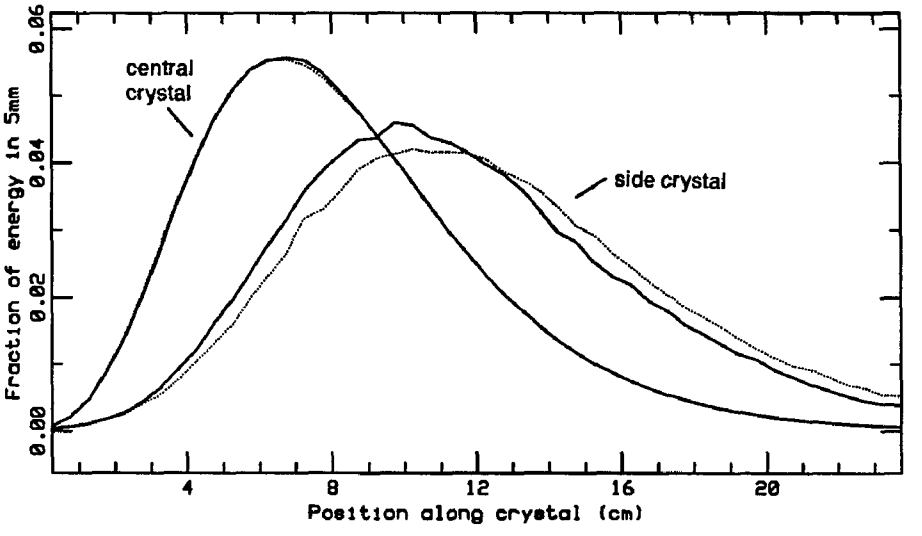

Figure 7.7: Shower profiles in central and side crystals

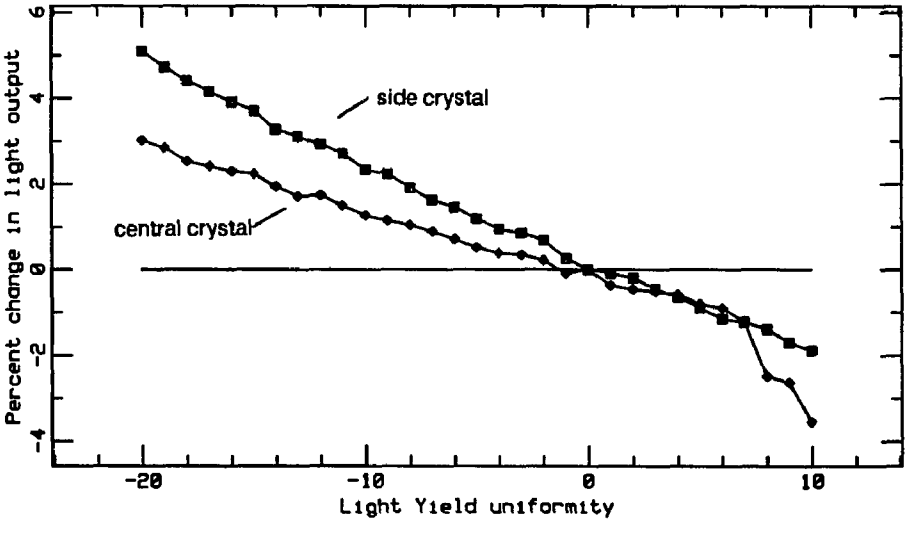

Figure 7.8: Change in relative light output due to change in shower profile as a function of the light yield uniformity, R

$d\theta$	$d\phi$	2 GeV	10 GeV	50 GeV
-1	-1	0.5888	0.4999	0.4265
-1	0	0.6360	0.5448	0.4670
-1	1	0.5829	0.5034	0.4273
0	-1	0.6252	0.5307	0.4566
0	0	0.8080	0.7061	0.6153
0	1	0.6174	0.5305	0.4533
1	-1	0.5598	0.4819	0.4073
1	0	0.6136	0.5210	0.4431
1	1	0.5607	0.4787	0.4036

Table 7.1: Light yield correction coefficients, ξ

where E' is the corrected energy, E is the uncorrected energy, E_{in} is the incident energy, $d\theta$ and $d\phi$ are the position of the crystal relative to the center of the matrix, and R is the light yield uniformity of the crystal, defined above. The output of the Monte Carlo simulation is summarized in coefficients ξ. Values for ξ are given in Table 7.1. The normalization $1 + R\xi(d\theta = 0, d\phi = 0, E_{in} = 10 \text{ GeV})$ makes the correction unity for the central crystal at 10 GeV. A separate correction is required to compensate for energy leakage. This treatment neglects the change in shower profile due to shifts in the impact point of the electron and changes in crystal shape as a function of θ. We have found the correction to be sufficiently accurate, given the accuracy required for the calibration and the uncertainty in the measured light yield uniformities.

For reconstruction of events at LEP it will be necessary to generate Monte Carlo events at a greater range of energies. To cover a continuum of energies, we have found that interpolation of the values of ξ between logarithms of energies is effective. One would expect a logarithmic dependence because the depth of the shower profile scales with $\ln E$.

7.4 The Mechanics of Calibration

Each crystal is calibrated individually. The tedious tasks of positioning the detector, ensuring that an adequate data sample is taken for each crystal, and monitoring the beam conditions are accomplished by programs running on the test beam computer. The high degree of automation speeds the calibration and makes it less prone to human error. The program in control of the flow of the calibration

is MOCA, written by E. Longo of Rome. MOCA's primary task is to position a crystal at the center of the beam, wait until a specified number of events have been collected for that crystal, and then move to the next crystal. It is unlikely that the calibration could have been completed in the allotted time were human intervention necessary to move from crystal to crystal.

To permit each crystal to be moved to the center of the beam, the detector is mounted on a moving table, which was constructed at the University of Geneva. The table can rotate the detector in theta and phi, and move it in x. Motion on these three axes is sufficient to position each crystal at the center of the beam with the long axis of the crystal aligned with the beam. A drawing of the table is shown in Fig. 7.9. Aside from the impressive mechanical achievement of moving the eight tons of detector, its significant parameters for the purposes of calibration are how accurately it can position each crystal relative to the beam line.

The readout for the three coordinates has a precision of 1 millidegree which corresponds to a precision of 10 microns in crystal positions. The table control program accepts an error of up to 10 millidegrees in any coordinate when positioning the table to allow for the difficulty of moving the eight tons of detector. The performance of the table was investigated by comparing the impact point calculated from the beam chambers (to be discussed in the next section) and the table coordinates to the impact point from the center of gravity of energy deposited in a 5×5 crystal matrix[51]. By comparing the results for a given crystal calibrated twice at one energy, we eliminate errors due to deformations within the barrel and in the survey of the table and the beam chambers. The root mean square change in the mean deviation between the center of gravity and mechanical positions was 0.16 mm in both theta and phi of the front face of the crystals. Therefore, the table positions are repeatable to that accuracy.

7.5 The Beam Line

The electron beam is a tertiary beam[52]. The primary protons are accelerated to 450 GeV/c by the CERN Super Proton Synchrotron (SPS). A slow extraction, 2.4 seconds of beam at 14.4 second intervals, of the protons onto a primary target, one interaction length (40cm) of beryllium, produces the secondary H3 beam. The H3 beam shares its target with the H1 beam line, so the selection of target is not under our control. We use the H3 beam in a negative 135 GeV/c mode. The H3 beam is steered onto a secondary target, 750 meters downstream from the primary target, to produce the tertiary X3 beam. By the time the beam reaches the X3 target, it is predominately pions but has a significant electron component from in flight decays.

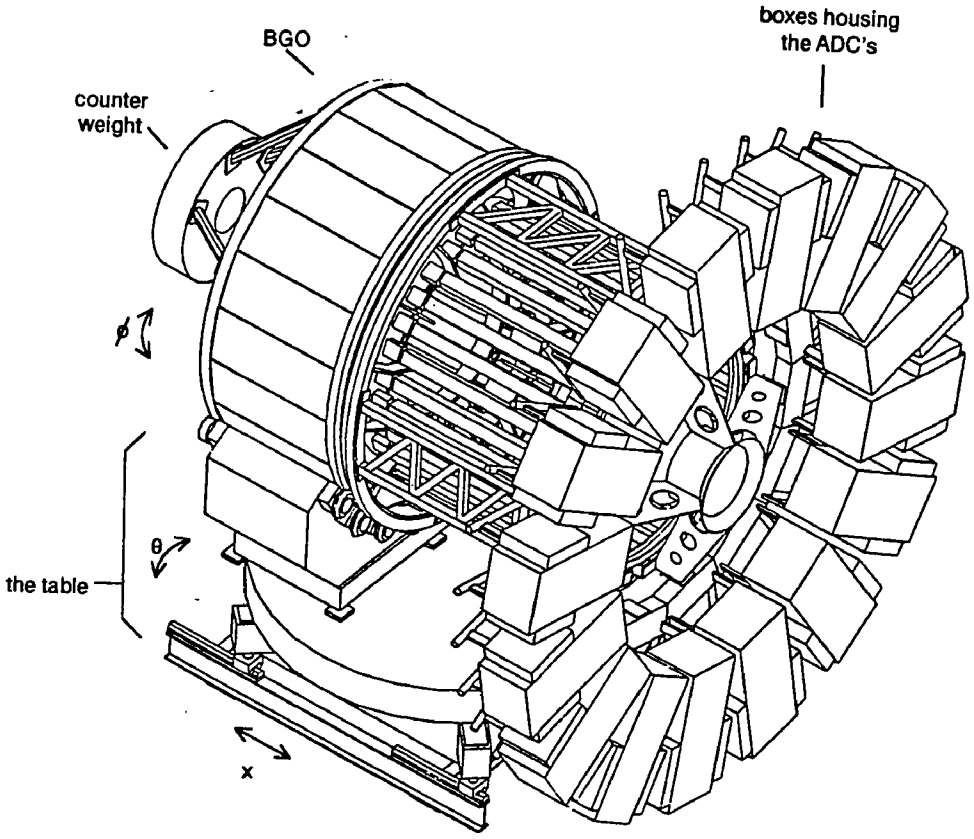

Figure 7.9: The Turning Table

Below 20 GeV, the tertiary electron beam is produced via a one interaction length thick beryllium target. Electrons in the desired energy range are obtained via three physical processes: degradation of the energy of electrons already present in the beam, conversion of bremsstrahlung photons, and conversion of pions (via a few intermediate steps). Most of the electrons at higher energies (upwards of 8 GeV) are degraded electrons from the primary target. Below a few GeV, the dominant process is conversion of negative pions to photons. Negative pions can interact hadronically to produce neutral pions, some of which decay to photons while still in the target. The photons can then convert to produce electrons. The production of electrons from bremsstrahlung photons is a lower rate process. At 10 GeV, the beam has a few percent contamination of negative pions which have degraded in energy or been produced in the target. At 2 GeV, the pion contamination is negligible.

At 20 GeV and above, a one radiation length (5 mm) thick lead target is used. The tertiary beam is obtained by degrading the energy of electrons from the secondary beam. Lead is chosen to maximize the ratio of the coulomb cross section to the hadronic cross section in order to minimize the pion background.

The beam line after the X3 target is shown in Fig. 7.10. The secondary beam strikes the X3 target and is converted into the tertiary beam. The beam is refocussed by the two focussing quadrapoles, Q1 and Q2. The electrons are momentum selected using the pair of collimators and the dipole magnet B1. The selected beam has a $\Delta p/p$ of 1%. The six wire chambers, DWC1-6, together with the dipole magnet, B2, form a highly accurate magnetic spectrometer capable of measuring momentum to 0.3% for 10 GeVelectrons, which will be discussed below. The selected beam is refocussed onto the BGO calorimeter by the quadrapole magnets, Q3 and Q4. Two helium filled čerenkov counters, C1 and C2, are used for particle identification. The čerenkov counters are useful in distinguishing electrons from pions at energies less than 20 GeV . They are too short (Č1 is 8 meters and Č2 is 5 meters) to be effective at higher energies. Fortunately, the beam has very few negative pions below 4 GeV, so at low energies momentum selection via B1 very effectively eliminates the pion contamination. Therefore, the čerenkov counters are not necessary and are evacuated in order to reduce multiple scattering.

The scintillators, S1, S2, S5, SH, and the A's, are used for triggering[53]. The main trigger is the coincidence of the two 5 cm × 5 cm counters, S1*S2. Scintillator S1 is upstream of the magnet B2, and S2 is downstream of the magnet and the čerenkov counters, therefore a particle which strikes both scintillators must have passed through the entire beamline. To veto halo particles, there are four scintillators used in anti-coincidence, AL, AR, AT, and AB, surrounding S2. To define the narrow beam used for calibration, there is a one 1 cm × 1 cm scintillator, S5. It is mounted, with DWC6, as close to the calorimeter as possible. Also for use in

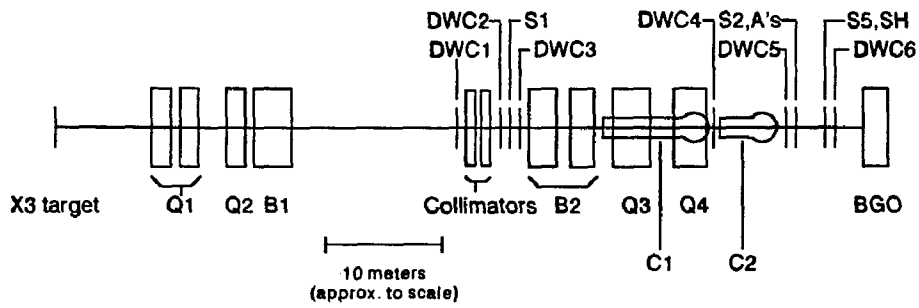

Figure 7.10: The X3 Beam Line

the calibration, there is a 16 cm × 16 cm scintillator, SH, mounted near S5. This counter is used to eliminate double electron events. It acts as a veto if hit before the main trigger and is used offline to flag events in which a second electron arrived after the main trigger but before the BGO electronics had stopped integrating the signal. SH is large enough that no particle can miss it and still shower into the region being calibrated.

The trigger used for 50 GeV and 10 GeV running is coincidence in S1, S2, S5, Č1, and Č2 together with anti coincidence in any of AL, AR, AT, or AB. For 2 GeV running the čerenkov counters are evacuated and are not required in trigger. Typical beam conditions for each of the three energies are given in Table 7.2. The trigger rate is never more than half the selected beam rate because the synchronous nature of the calorimeter ADC electronics imposes a fifty percent duty cycle. The loss of events at higher beam rates is due to approximately 5 ms of dead time required to read out the test beam instrumentation. The BGO readout system is fully buffer and never required more than 1 ms of dead time for an event.

The magnet B2 and the wire chambers DWC1-6 are used to measure momentum. The wire chambers are a standard CERN item[54]. The wires are readout via an analog delay line, which converts the position information from the chamber to more easily handled timing information. The chambers have a resolution of about one millimeter. The chambers have been placed to maximize the lever arm for the momentum measurement, within the confines of the test beam. To focus the beam there are two downstream quadrapoles. Because of the limited space in the test beam, the quadrupoles are within the spectrometer. This increases the uncertainty in the momentum calculation. The three downstream chambers are also used to

Energy	2 GeV	10 GeV	20 GeV	50 GeV
Target	Be	Be	Pb	Pb
Hits per burst				
S2	16804	9240	16589	11132
S1*S2	4273	2077	2408	2283
Č1	-	3041	6031	3930
Č2	-	6811	15198	16442
Selected beam	468	966	1235	1316
Triggers	209	370	379	367
Geometric vetoes				
A veto	1.8%	1.0%	2.8%	2.7%
SH veto	0.0%	0.0%	0.0%	0.0%
S1*S2*S5/S1*S2	11.1%	49.8%	52.8%	60.1%
Particle id				
C1 pressure	0.04 bar	1.00 bar	1.02 bar	0.99 bar
C1 efficiency	-	0.98	1.00	0.99
C2 pressure	0.04 bar	1.29 bar	1.30 bar	1.28 bar
C2 efficiency	-	0.99	0.99	1.00
Electrons	2510	1909	2302	2175
Pions	-	93	20	13

Table 7.2: Typical beam conditions

find the impact position (in x) at the BGO. There are two additional chambers downstream used to find the impact point in y.

It is crucial for the momentum measurement and the reconstruction of the impact point at the BGO that the wire chambers give accurate position data. There are two components necessary for this. The knowledge of the position of the chamber relative to the other elements of the beamline and to the BGO, and the knowledge of the time to position conversion for each chamber.

The delay wire chambers produce two signals when a particle passes through them. The signals are essentially the time it takes the pulse generated by the particle to pass to the two opposite ends of the delay line. The position is determined by taking the difference of the two times. The sum of the times is used as a check on the quality of the hit. The time to position conversion is found by putting a pulse on selected wires in the DWC. A linear relationship is assumed. There are calibration inputs on the central wire and on wires at ± 30 mm. The chambers were recalibrated several times. We have chosen to use the calibration done by D. Boutigny and Y. Karyotakis in November 1987. In these chambers, the delay time depends on the pulse height and shape. Because the pulse height for particles is not measured during running we cannot correct for this effect. To study this, the calibration has been done with a range of pulse heights. It has been estimated that variations in the pulse height contribute a systematic error of the 0.5 mm to the chamber resolution[55].

The wire chambers have been surveyed several times. We choose to use the survey done on March 1988. The path of an electron on either side of magnet B2 should be a straight line. There are three chambers on either side of the magnet, so we can check the alignment of the chambers by comparing the position of the hit in the center chamber, DWC2 or DWC5, to the position in the central chamber extrapolated from a linear fit to the positions in the other two chambers. We refer to the difference between the actual hit and the fitted position as the residuals. Plot of the residuals for the two sets of three chambers, for the calibration runs during the summer of 1988, are shown in Fig. 7.11. The plots are the difference between the particle position at chamber 2 (5) and the fitted position from chambers 1 and 3 (4 and 6). There are several obvious abrupt shifts in the chamber alignment. At each point there is a shift, we examine plots of the mean position at each chamber in order to determine which chamber moved. A table of the chamber realignments is given in Table 7.3. It is, unfortunately, not possible to check the alignment of the Y chambers, because there are only two of them.

Examination of the difference between the impact position calculated from the energy deposited in the BGO crystals and impact position calculated from the beam chambers and the table coordinates checks the performance of the entire geometrical

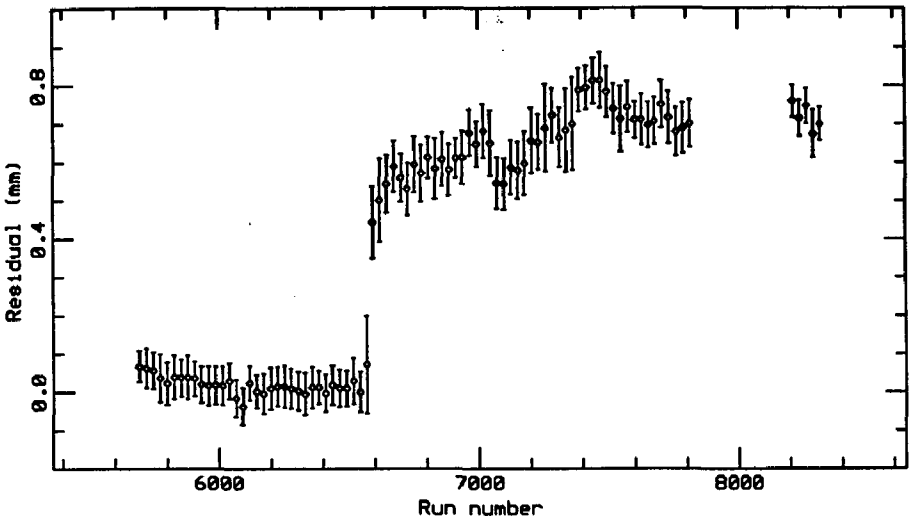

Upstream chambers

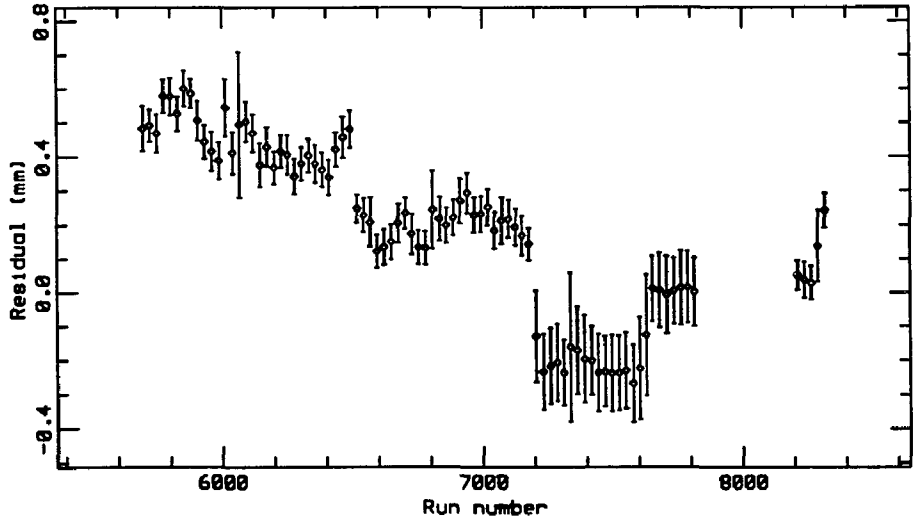

Downstream chambers

Figure 7.11: Wire chamber (mis)alignment

	Shift (mm)					
Run	DWC1	DWC2	DWC3	DWC4	DWC5	DWC6
6576		-0.8	-0.3			
6680	-0.4		0.0			-0.6
7017						0.0
7192	0.0			+1.7	+1.3	
7634				-0.6		

Table 7.3: Chamber realignments

system, the beam chambers, the table, and the survey of their relative positions. The single crystal calibration method is particularly sensitive to the position of each individual crystal. As discussed in Section 7.2, the algorithm determines the impact position where the electron deposits the maximum energy, *i.e.* the center of the crystal. Figures 7.12 and 7.13 show the x and y coordinates of the centers of the crystals as functions of θ and ϕ. The variations in x and y as a function of ϕ arises because the structure housing the crystals is not exactly concentric with the outer mechanical structure; there is an offset of about one millimeter. The origin of the bimodal structure in y as a function of θ is unknown. The shift in x at low θ is most likely due to a shift in the entrance angle of the beam. For the analysis, corrections have been applied for x and y as functions of ϕ and for x as a function of θ. The remaining position errors are on the order of 1 mm, which is consistent with the error we expect extrapolating the position from the wire chambers.

In addition to the momentum measuring magnet, B2, there are two focussing quadrupoles between the two sets of chambers. This complicates the momentum calculation. To handle the complications from the quadrupoles in a computational efficient way, we use a transfer matrix formalism to reconstruct the trajectories of particles passing through the beamline[56]. In this technique, the path of each particle is calculated relative to the path of a hypothetical reference particle. This reference particle defines the zero of position in each chamber and has a momentum referred to as the 'central' momentum of the beamline. In calculating the momentum, we consider motion only in the bending plane of the dipole magnet. The trajectory of a particle can be specified by a three component vector, $(dx, d\theta, dp)$, describing the deviation in position, angle, and momentum, from the reference trajectory. The motion of the particle through a beamline element, such as a dipole magnet, a quadrupole magnet, or a few meters of air, can be described by a 3×3 matrix acting on the coordinate vector. To first order in the deviations from the

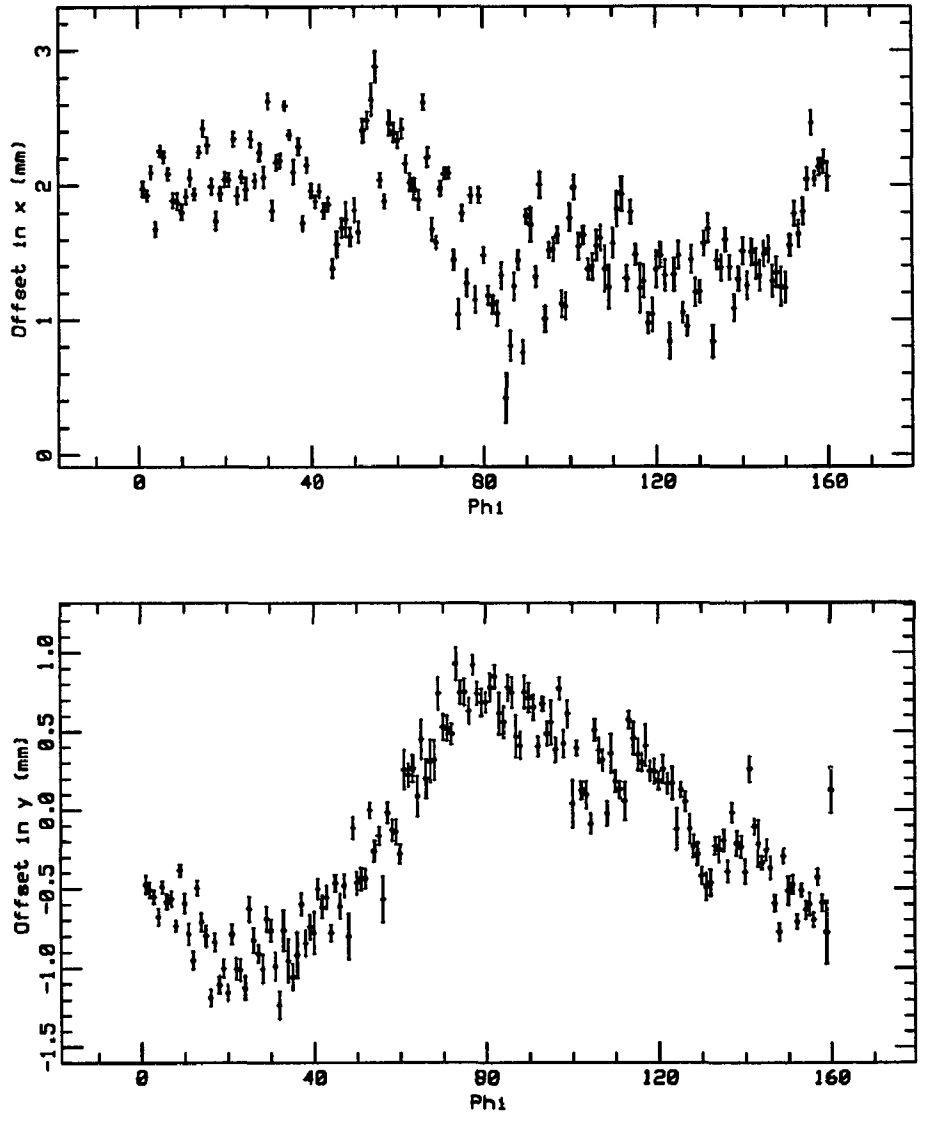

Figure 7.12: Offsets in crystal positions versus ϕ

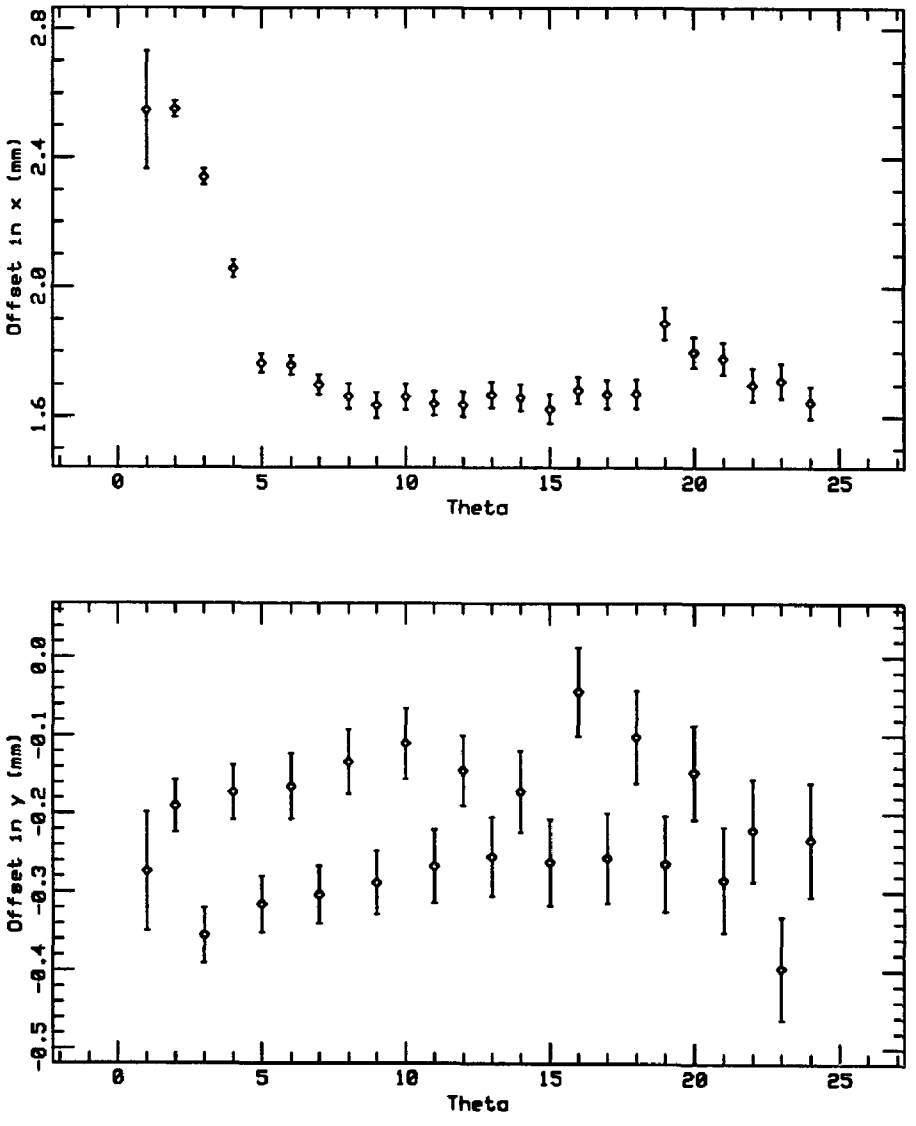

Figure 7.13: Offsets in crystal positions versus θ

reference trajectory, motion through the entire beamline can be reduced to a single matrix which is the product of the matrices for each element in the beamline. We use transport matrices calculated from the measured parameters of the X3 beamline by Y. Karyotakis. The coordinates, $(dx, d\theta, dp)$, at one set of chambers can be calculated from the coordinates at the other set of chambers via the transport matrix. Therefore, we can determine the momentum if at least three of the six coordinates are known. In practice, we require a position measurement in both sets of chambers and an angle measurement in at least one set of chambers.

The momentum dispersion of the beam is very small, less than one percent. Therefore a first order calculation for the momentum should be sufficient. As a check a second order calculation has been performed. A comparison of the first and second order momenta shows that the inclusion of second order terms does not significantly alter the result.

The momentum can be calculated from hits in three chambers with at least one hit in both the upstream and the downstream sets. We have chosen to use only tracks which have at least two hits downstream and a hit in either X2 or X3. To reconstruct the impact point at the crystal, two hits downstream are required, so we lose no events by making the same requirement for the momentum calculation. The downstream track is reconstructed by fitting a line to all the downstream hits. Upstream, if there is a hit in X2, then the fitted position at X2 is used, otherwise the position at X3 is used.

To find the actual momentum of a particle, we must know the central momentum of the beamline. The central momentum used in the calculation of the transport matrices was the nominal beam momentum, *i.e.* 2, 10, 20, or 50 GeV, as calculated by the SPS beamline computers. A correction to this value is necessary because the field strength for the dipole magnet used in the calculations is not entirely accurate. The field of the magnet B2 was accurately measured by D. Luckey and G. Gratta in December 1986. However, their measurements carry a caveat because the equipment was reliable only for fields up to that used for 20 GeV electrons. The value for 50 GeV is an extrapolation of measurements for lower energies and is therefore suspect. In addition, there are small fluctuations in the magnet current and the magnetic field on time scales of several hours. The field strength is continuously monitored via a hall probe, which can be used to correct for these fluctuations. To determine the momentum, we use the central momentum assumed in the calculation of the transport matrices, and correct for the difference between the assumed central momentum and the actual central momentum computed from the measured values of the magnetic field. The corrections to the field strength are on the order of a few percent. We can therefore made a simple linear correction, and neglect higher order terms.

7.6 Temperature

A BGO crystal in our detector is approximately 24 cm long, 2 cm×2 cm on one end, and 3 cm×3 cm on the other end. The crystals are housed in a carbon fiber structure, which surrounds each crystal on five of its six faces. The temperature near the BGO is measured at both the inner and outer edges of the crystal by two types of sensors. At the outer edge, the sensors are glued directly to the BGO. At the inner edge the sensors are glued to the exposed face of the carbon fiber structure. One type of sensors (AD590's) is readout via the normal BGO readout system. A temperature measurement was made after every 1000 events during the calibration, *i.e.* on the order of once a minute during normal data taking. These readings are recorded on magnetic tape with the normal data. They are the set of readings that we use to determine the temperature on an event by event basis for the calibration. The other type of sensors (Pt1000's) is read out by a dedicated system constructed at the University of Lausanne. These sensors are used to monitor the AD-590's.

The light output of the BGO changes by -1.55% per °C of temperature change. To calibrate to an accuracy of a few parts per thousand, we must measure the temperature with an accuracy of a few tenths of a degree. To achieve this level of accuracy, the final temperature calibration is done after the sensors are mounted on the detector in order to take into account self-heating of the sensors. We use two sets of temperature sensor calibration files (one for each half barrel) generated by L. Vuilleumier and T. Boehringer of the University of Lausanne. Since the light output of BGO varies linearly with the temperature, we are in a situation where 'consistency is more important than the truth'.[1] It is extremely important that we have a method for calculating the temperature that gives results which are consistent across the calibration and running at LEP. However, the accuracy with which we know the temperature of the crystals on an absolute scale is irrelevant.

Once we have the temperature from each sensor, we must then reconstruct the temperature inside the crystals. The main factors determining the heat flow in the detector are the thermal properties of BGO and carbon fiber composite material. Thermal conduction in BGO is quite slow, while in carbon fiber it is relatively fast. The rate at which a localized pulse of heat diffuses into a material is given by $\overline{x^2} = 2Dt$, where $\overline{x^2}$ is the mean square extent of the pulse, D is the thermal diffusibility of the material, and t is time. For BGO the thermal diffusibility is 34 cm^2/hour. Were the detector solid BGO, it would take about two hours for heat from external sources to propagate to the middle of the crystals. The thermal diffusibility for carbon fiber is 1300 cm^2/hour. So, if there were no BGO, a heat pulse would require only three minutes to reach the midpoint of the structure. The

[1] A favorite phrase of P. Piroué

difference in conductivities makes constructing a good model of heat flow in the detector rather difficult.

We have taken a special set of data to aid in the understanding the thermal properties of the detector. The BGO is normally at 18°C. For the last few days of running at the test beam, the temperature was raised to 23°C. A plot of the average temperature on the inner and outer sensors is shown in Fig. 7.14. At the time marked $t = 0$ on the plots, the controls for the temperature of the fluid in the cooling system and on the air conditioner that regulates the ambient temperature were both increased by 5°C. The inner sensors measure the temperature on the carbon fiber structure. There is a jump in temperature on the inner sensors on a time scale of a fraction of an hour immediately after $t = 0$. This corresponds to warming of the inner surface of the carbon fiber structure. The inner sensors remain at a higher temperature than the outer sensors, which are attached to the BGO, until equilibrium is reached, after about one day. The time scale for the heating of the entire detector is determined mainly by the heat capacity of the BGO.

We conclude that the temperature in the BGO changes rather slowly. Readings on the order of once an hour should be sufficient to adequately reconstruct the temperature. For the temperature analysis we use the temperatures from each sensor for the first temperature reading of each run. This gives readings on the order of once per hour during normal electron calibration data taking.

The sophistication required in the temperature calculation is proportional to the magnitude of the temperature gradients. If the temperature is essentially uniform, any simple algorithm will suffice. If there are sizable gradients and time variations, we will have to apply heat flow equations to get good results. If there are large gradients, the task is impossible because of the limited number of temperature sensors. A typical temperature profile for the second half barrel is shown in Fig. 7.15. The temperature is uniform to within a few tenths of a degree both spatially over the region being calibrated and temporally over a few hours. However, there are appreciable temperature gradients along the detector in θ. This forces us to take account of the geometry of the detector in calculating the temperatures.

The light output of a crystal is determined by the average temperature inside the crystal. Small gradients at the surface of crystals have little effect on the light output because the temperature inside the crystals is an average of the temperature profile at their surface. The averaging washes out the gradients. We have found that, except for crystals near the edges of the detector, a direct average over the sensor readings gives the same result as an average calculated from a complicated fit to the surface temperature profile. There are no temperature sensors on the exposed faces of the crystals at the edges of the detector at $\theta = 1$ and 24, so we cannot perform an average over the temperature profile on those surfaces. This

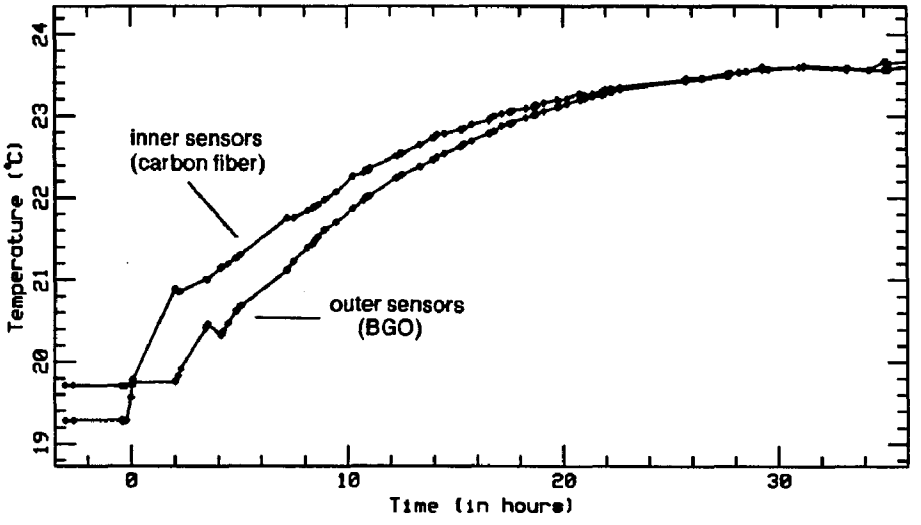

Figure 7.14: Temperature versus time

	Theta							
	2	5	8	11	14	17	20	23
68				18.4				18.7
67			18.4				18.6	
66		18.3				18.6		
65	18.3				18.5			
64				18.4				18.7
63			18.4				18.7	
62		18.4				18.6		
61	18.2				18.5			
phi 60				18.7				18.8
59			18.6				18.7	
58		18.4				18.5		
57	18.5				18.6			
56				18.6				18.9
55			18.5				18.8	
54		18.5				18.7		
53	18.5				18.7			
52				18.6				18.8
51			18.5				18.7	

Figure 7.15: Temperature (in °C) versus position in the detector

has lead us to treat the temperature profile as a function of θ differently from the temperature profile as a function of ϕ.

We calculate the temperature using a fit to the temperature versus θ and an average of the temperature over ϕ and time. To determine the temperature of a particular crystal at a given time, (θ_0, ϕ_0, t_0), we divide the detector into slices of constant θ and take a weighted average over ϕ and time for each slice. A sensor's weight is inversely proportional to the square of the distance from the sensor to the center of the slice. The interval between the time of the reading and t_0 is added in quadrature to the distance, using a time to distance conversion of 8 cm/hour. This is approximately the speed at which a heat pulse diffuses in pure BGO. Only sensors in a 17×17 crystal matrix centered on (θ_0, ϕ_0) and readings made before t_0+ 1 hour are used. The temperature of the crystal is extracted from a linear fit to temperature versus θ performed on the θ slice averages.

As discussed in section 7.1, we define the average temperature of a crystal, T, as the temperature at the shower maximum. The center of slice used to calculate the weights for the averages is a point a few centimeters, 6 cm for 10 GeV, from the front face of the crystal at $\phi = \phi_0$ in the slice. The temperature correction is defined to be zero at 18°, the nominal operating temperature of the BGO crystals when installed in L3. The ADC readings are corrected according to the equation $V' = V[1 + 0.0155(T - 18)]$. The light output of BGO decreases with increasing temperature.

To study the accuracy of our temperature correction we have compared calibration constants produced with 10 GeV data for the same set of crystals in the second half barrel during June and August of 1988. The two sets of data make a good test of the temperature correction because the average temperature changed from 17°C in June to 19°C in August. The percentage change in the calibrations constants is shown in Fig. 7.16. An upper bound on the error in the temperature correction, from the width of the superimposed gaussian fit, is 0.35%. This width also includes the inaccuracy in the calibration procedure. It is particularly heartening that there is no systematic shift in the mean of the calibration constants. The mean of the gaussian is 0.08%, which is below the expected resolution of the detector.

7.7 Pedestals

The BGO electronics is a synchronous system, designed for the synchronous environment at LEP. At the test beam, the arrival times of particles are not known in advance, so the electronics cannot be operated in a synchronous mode. The electronics must be able to accept triggers during a relatively wide gate. This compromises some of the feature of the ADC system and causes a few difficulties,

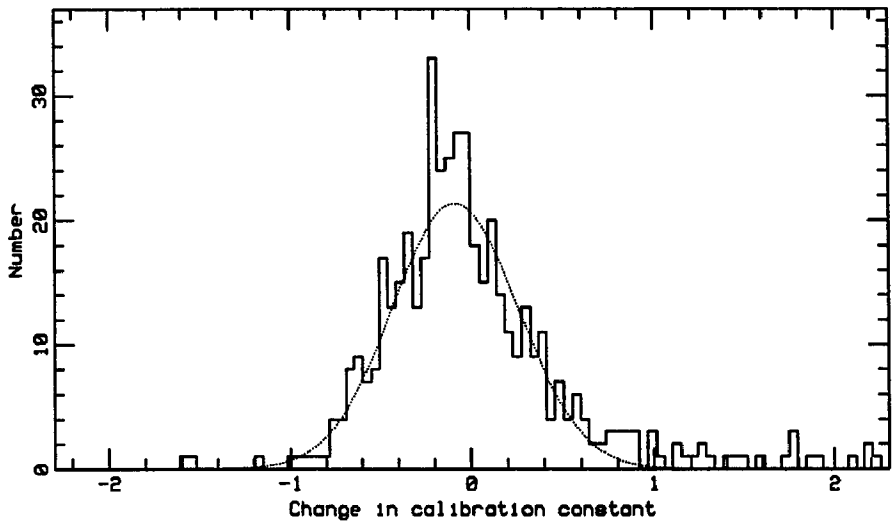

Figure 7.16: Comparison of calibration constants at different temperatures

the most important of which concerns the pedestal of the ADC reading.

A rough sketch of the electronics up to the ADC is shown in Fig. 7.17. The charge from the photodiode on the BGO is collected by the preamp, sent through the front-end of the ADC, and fed into the integrator. The integrator and the sample/hold were designed to operate with the synchronous beam timing of LEP. Since the arrival times of particles are known in advance, the entire timing cycle is fixed. The total integration time is always 5.5 μsec and the integration time after the particle arrives is always 5.0 μsec.

However, when running at the test beam the arrival times of particles are not known in advance. The timing cycle must be adjusted for the arrival time of each individual particle. A diagram of the ADC timing at the test beam is shown in Fig. 7.18. The integration time after the particle arrives must be kept fixed at 5.0 μsec because the fraction of the signal integrated depends on the length of the integration after the particle arrives. The pulse being integrated has a time constant of the about one microsecond, so roughly 99% of signal is integrated within 5.0 μsec. We can time the arrival of the particle and adjust the integration time to an accuracy of a few 10's of nanoseconds. This level of precision keeps the variation in the signal, due to fluctuations in the integration time, well below 0.1%.

Since the trigger can occur anywhere in the trigger gate, the length of time from

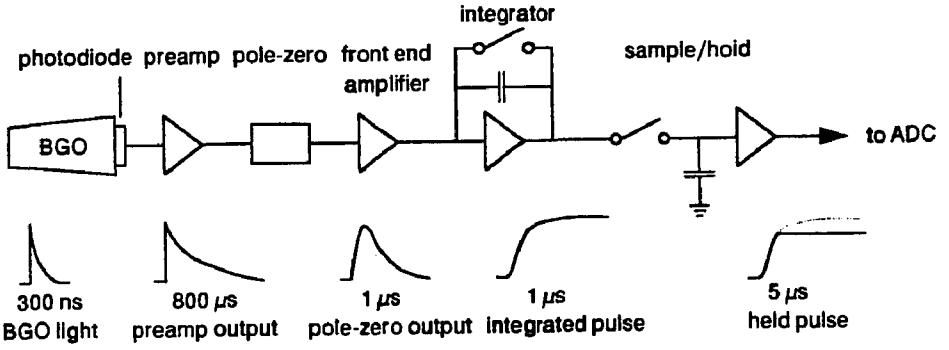

Figure 7.17: ADC front end

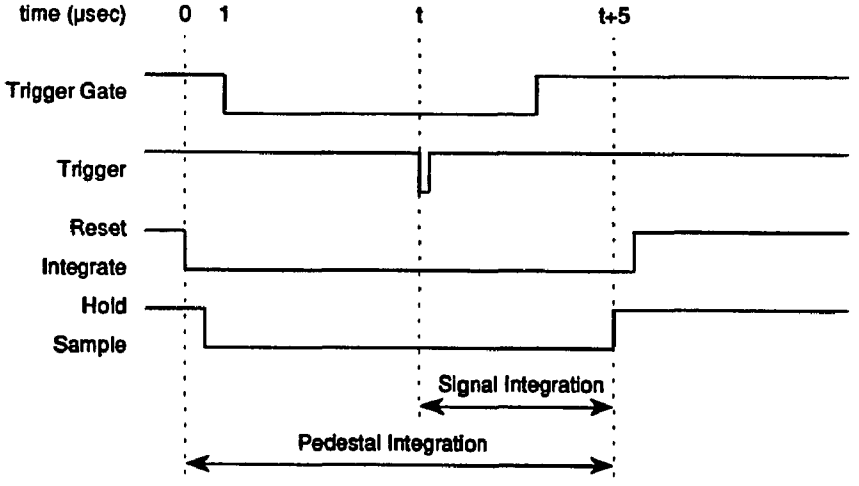

Figure 7.18: ADC timing

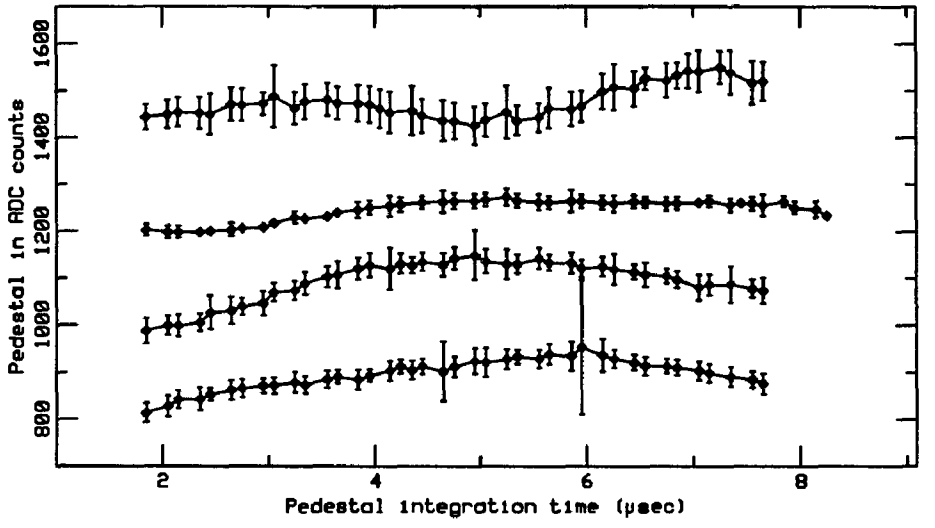

Figure 7.19: Typical pedestal shapes

when the integration begins to when sample/hold switches to hold varies, (we will refer to this time as the pedestal integration time). The value of the pedestal varies as a function of the pedestal integration time. A plots of pedestal value in ADC counts versus pedestal integration time are shown in Fig. 7.19. The width of the pedestal at a fixed point in time is a measure of the noise of the system. The typical noise is of the order to 20 ADC channels, though it varies from 8 to 40 channels. The larger effect is the variation with integration time.

The pedestal is measured during runs taken just before each beam data run. Pedestal runs are taken at least once every few hours. Pedestal triggers are generated by a 'stepping delay box'. This box generates triggers at 150 ns intervals across the range of allowed trigger times, producing 46 data points in the allowed 7 μsec gate. We could use these points directly as the value of the pedestal in each time bin, (perhaps with linear interpolation between bins). This would provide an accurate representation of the pedestal shape. However, pedestal runs are only 2000 events, so each bin has about 45 events. The pedestal subtraction would then be limited by statistics.

A fit to the data, using some appropriately chosen functional form which can reproduce the pedestal shape, is not limited by statistics. It is less susceptible to fluctuations in the readings for individual bins. It also has the advantage that

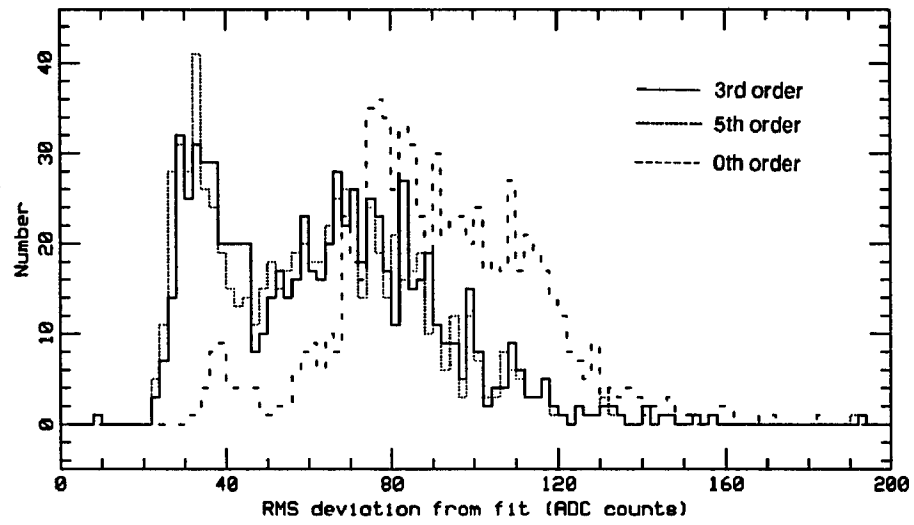

Figure 7.20: RMS deviation of pedestal from fit

the end result is more compact and somewhat easier to use. We chose to use a polynomial fit. The root mean square deviation from the fit for various orders of polynomials are shown in Fig. 7.20. The units in the plot are ADC counts on the first comparator. Adding terms beyond third or fourth order do not significantly improve the fits. For the calibration, we use a third order polynomial fit.

When summing over groups of crystals, the correlation of the errors in pedestal subtraction are important. We have calculated the correlation factor, R, from sums of 9 and 25 crystals. The correlation factor is 0.89 ± 0.09 for the sum of 9 crystals, and 0.87 ± 0.09 for the sum of 25 crystals. The correlation is very high because the pedestal shape is very similar for ADC channels powered by the same power supply. Since the pedestal shapes are similar, the polynomial fits will also be similar, and therefore the deviations of the fits from the pedestals will be correlated. The root mean square width of the sum of the pedestals of 9 and 25 crystals is shown in Fig. 7.21. The units in the plots are ADC counts. A rough conversion to energy units is 10 ADC counts = 1 MeV.

The average width of a sum of 25 crystals is 80 MeV and in some cases is as high as 130 MeV. For all crystals the sum of 9 has a width of less than 40 MeV. The width of the sum of 25 would limit the resolution to about five percent at 2 GeV and about one percent at 10 GeV. The contribution to the resolution of the sum of

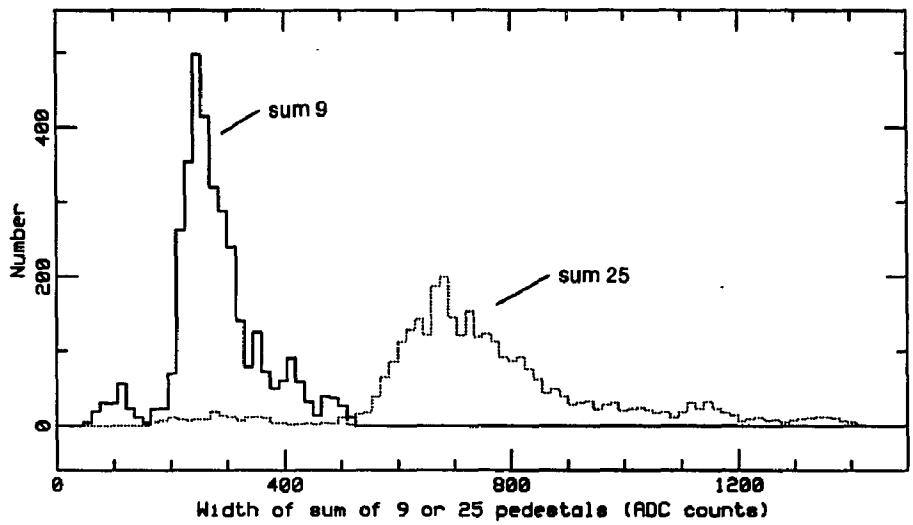

Figure 7.21: RMS width of sum of pedestals of 9 and 25 crystals

less than 2% at 2 GeV and a few tenths of a percent at 10 GeV. We have chosen to use the sum of 9 for calculating the resolution at all energies.

7.8 Electronic Gains

As discussed previously, the ADC on the level-1 cards consists of a DAC, together with a chain of amplifiers and comparators, as shown in Fig. 7.22. The chain of amplifiers is necessary to achieve the large dynamic range of the signal from the BGO using a 12-bit DAC. To improve the performance of the system, particularly with respect to small signals, there are two distinct chains of amplifiers: the low and high energy chains. To know the input pulse height, we must know the 12-bit reading from the ADC and the gain of the comparator used to make the measurement. Uncertainties in the gains adversely affect the energy resolution of the detector. However, since the intrinsic resolution becomes worse as the energy decreases, errors in gains of the low energy comparators have less effect on the overall resolution of the calorimeter.

The 6305s are programmed with a special mode, the 'bump' mode, that is used to determine the relative gains of the successive comparators. In bump mode the first reading is made using the comparator which gives the largest reading that is

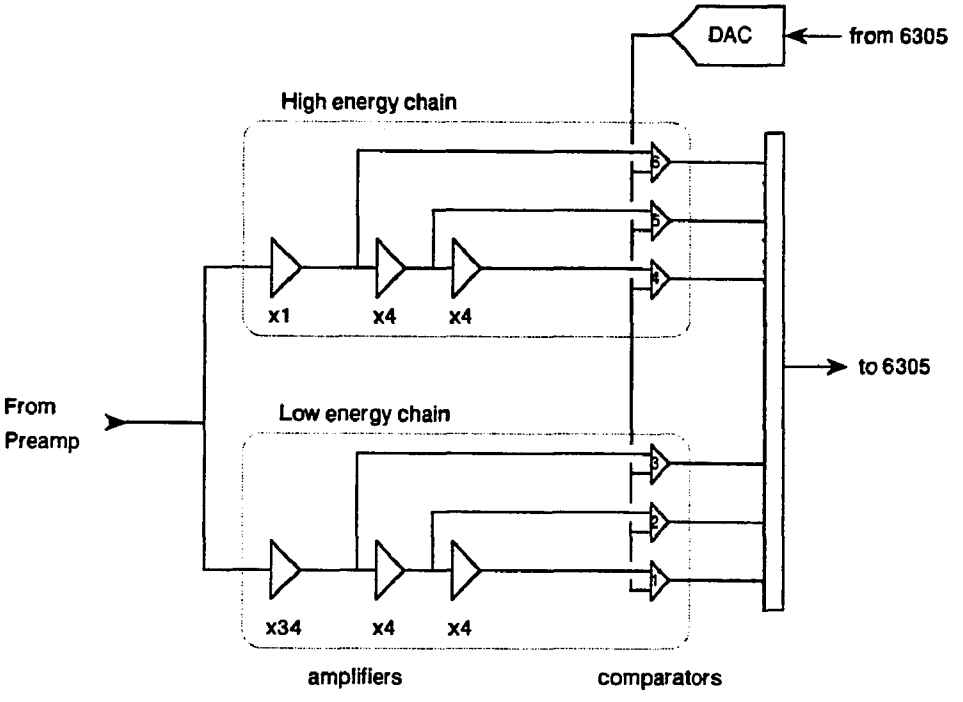

Figure 7.22: The amplifier/comparator chain

less than 7/8 of full scale, and a second reading is made using the next comparator up on the chain. If we plot the raw ADC of the second reading versus the raw ADC for the first reading, the slope of the line will be the relative gain. There will also be a non zero intercept because the op-amps used in the $\times 4$ amplifiers have a non zero relative offset voltage. If the first reading is on comparator three, we measure the voltage at the end of the low energy amplifier chain and the voltage at beginning of the high energy chain. In this case it is necessary to do a pedestal subtraction, since the high and low energy pedestals are completely independent and have different behavior as a function of time. When the gains were determined using the test set, all data were taken at a fixed time, so there was a constant difference between the two pedestals which was absorbed into the intercept of the line.

To accurately fit both the gain and offset it is necessary to have data with ADC readings over the entire range. The data taken with the test set uniformly cover the range of the ADC on each comparator. This allows an extremely accurate determination of the gains and offsets. The beam data provide good coverage on the low energy comparators, but not on the high energy comparators. Calibration data was taken at three energies: 2, 10, and 50 GeV. The 2 GeV data provide no readings on the high energy chain. At 10 GeV there is one set of readings on the fourth comparator for each crystal, from data taken when the beam was directed into that crystal. The fluctuations in the energy deposited in the central crystal are on the order of a few percent, so the readings cover only a narrow range of the possible ADC values. However, the energy fluctuations in the crystals adjacent to the central are a larger fraction of the energy deposited in those crystals. We find that the 10 GeV data provide a good range of readings on comparator two for the adjacent crystals. At 50 GeV, we have the same situation shifted up by one comparator. The net result of the beam data at the three energies is a wide range of readings on the three low energy comparators and a narrow band of readings on comparators four and five.

To determine the gains from the beam data, we made a linear fit to the second reading versus the first reading. We used the raw ADC readings directly and made no pedestal subtraction. Since we did not do a pedestal subtraction this procedure was not useful for finding the gain between the high and low energy chains, g_{34}, since the two chains have different pedestals. The g_{34} determination was done separately, and will be discussed later. The data was collected in sixteen bins of the first ADC reading for each comparator. We first attempted to fit for both the gain and the offset. To insure an adequate lever arm, fits were performed only if at least eight bins had more than 25 events. Most crystals had a sufficiently broad range of data on comparators one and two, but not on comparators four and five. To measure the gains on the higher comparators, we chose to fix the offset to the value determined

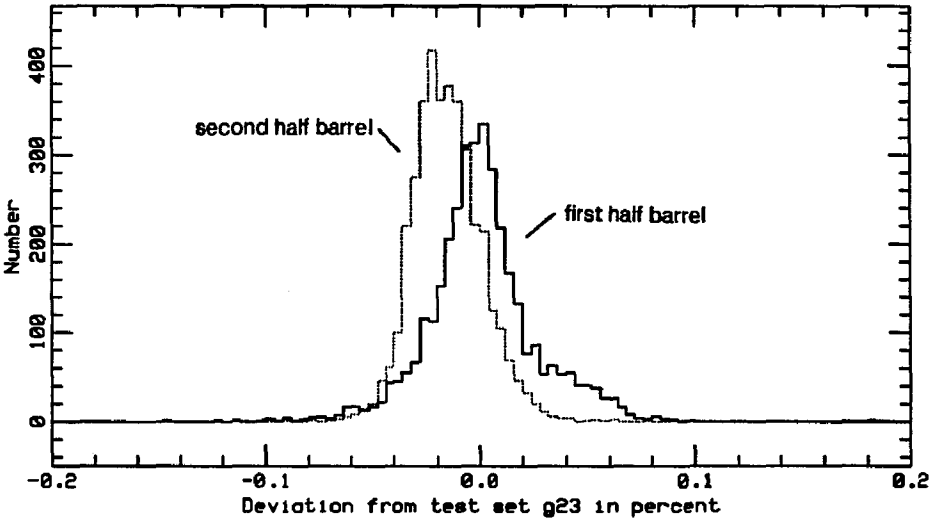

Figure 7.23: Comparison of g_{23} from beam data and the test set

using the test set, and require only three bins for the fit.

The results for the gain between comparators 2 and 3 are presented in Fig. 7.23. The quantity histogrammed is the fractional difference in gain in percent between comparators 2 and 3, Δ_{23}. We define Δ_{ij} for any pair of comparators to be

$$\Delta_{ij} = 100\frac{g_{ij}^{TS} - g_{ij}^{B}}{g_{ij}^{TS}}, \tag{7.7}$$

where g_{ij}^{TS} is the gain between comparators i and j as determined by the test set and g_{ij}^{B} is the gain determined from the beam data. Figure 7.23 has two histograms. One is Δ_{23} for all of the crystals of the first half barrel, (solid line), and the other is Δ_{23} for all of the crystals of the second half barrel, (dotted line). The average and RMS deviation of $\Delta_{i,i+1}$ for all comparators, except number three, are shown in Table 7.4. The Table presents data for both half barrels. The agreement is extremely good: the gains agree to better than a tenth of a percent, and the offsets agree within a fraction of an ADC count.

Determining the gain between comparators three and four is more difficult than for the other gains because we must compare data from the low and high energy chains. The high and low energy chains have different front-end amplifiers, integra-

comparator	fit with offset fixed Δgain in percent		fit with offset free Δgain in percent		fit with offset free Δoffset in adc bits	
	mean	rms	mean	rms	mean	rms
1^{st} half barrel						
1	0.031	0.041	-0.022	0.039	-0.20	0.17
2	0.001	0.025	-0.005	0.026	-0.04	0.09
4	0.014	0.032				
5	0.027	0.035				
2^{nd} half barrel						
1	-0.004	0.021	0.002	0.037	-0.02	0.23
2	-0.015	0.018	-0.009	0.022	0.03	0.12
4	-0.008	0.036				
5	0.021	0.045				

Table 7.4: Difference between gains and offsets from test set and beam data

tors, and sample/hold circuits. This causes the two chains to respond to a given pulse shape in different ways. To circumvent this problem a special pulser was constructed that was designed to simulate as closely as possible the pulse produced by the photodiodes attached to the BGO. The output of the pulser was fed into a preamplifier card connected to the level-1 board. Ideally, the preamp/level-1 pair used would be the pair as in the half barrel. Unfortunately, the same preamp card was used to measure the gains of all the channels for the first half barrel. The determination of the gains for the first half barrel is therefore sensitive to variations in the preamp cards. In addition to the pulse shape problem, the two chains have different pedestals. Since there is a non-negligible variation of the pedestal with time, it is necessary to do pedestal subtraction before calculating the gain.

We used only the 50 GeV data in the g_{34} analysis. A wide range of ADC readings on comparator three are obtained from electromagnetic shower fluctuations in the crystals adjacent to the central. We obtained pedestal ADC reading using the pedestal fits described in the preceding section. Note that we assumed that the g_{12} and g_{23} gains from the test set were correct in order to extrapolate the pedestals for the low energy chain. With the pedestal-subtracted readings we performed a linear fit of the difference, y, between the reading on comparator three, R_3, multiplied by the test set gain, g_{34}^{TS}, and the reading on comparator four, R_4, versus the ADC reading, x, on comparator three,

$$y = g_{34}^{TS} R_3 - R_4 \qquad \text{and} \qquad x = R_3 \tag{7.8}$$

We fit for the slope, s, and the intercept of the line. If the test set gains were correct then the gain adjusted reading on comparator three would always be the same as the reading on comparator four, resulting in a zero slope. A nonzero slope means that the test set gain should be corrected. The beam g_{34}^{B} is given by

$$g_{34}^{B} = g_{34}^{TS} - s \;\Rightarrow\; \Delta_{34} = 100\frac{s}{g_{34}^{TS}} \tag{7.9}$$

We found the average correction, Δ_{34}, to be $-0.37 \pm 0.76\%$ for the first half barrel, and $-0.36 \pm 0.59\%$ for the second half barrel. The error bars are the RMS deviation of the correction. These values are an order of magnitude greater than those for the other comparators.

To estimate the size of the errors on the beam gains, g_{34}^{B}, we used the crystals which had data on more than one tape. In determining g_{34}^{B} for a given crystal we used the data from only one tape. For those crystals with a good data sample on more than one tape, we plotted the fractional difference in percent, e_{34}^{B}, between the two g_{34}^{B} values from the two different tapes.

$$e_{34}^{B} = 100\frac{g_{34}^{B} - g_{34}^{B'}}{g_{34}^{B}} \tag{7.10}$$

where g_{34}^{B} is the gain from the tape with more data for the crystal, and $g_{34}^{B'}$ is the gain from the tape with less data for the crystal. Figure 7.24 is a histogram of e_{34}^{B} for all the crystals with sufficient data from the first half barrel. We estimate the error by finding the RMS deviation from zero to be 0.48%. The main limit to the accuracy comes from the complicated pedestal shape. The ADC range for comparator 3 is restricted to 800 counts. The corresponding second reading, on comparator 4, varies over 400 ADC counts. For a 0.5% error we need to make an error in pedestal subtraction of only four ADC counts in comparator 3 or two ADC counts in comparator 4.

If the same correction is known to apply to a group of crystals, then, to the extent that the pedestal subtraction errors are uncorrelated among the crystals, we can improve the accuracy of our result by the square root of the number of crystals. One possible systematic error in the test set g_{34} determination is in the special pulser and the preamp card used. In determining the g_{34}'s for the first half barrel with the test set only twelve channels of one preamp card were used. In the barrel, each row in phi has one 24 channel preamp card and two 12 channel level-1 cards. Crystals at $\theta = 1 - 12$ correspond to channels 1-12 on one level-1

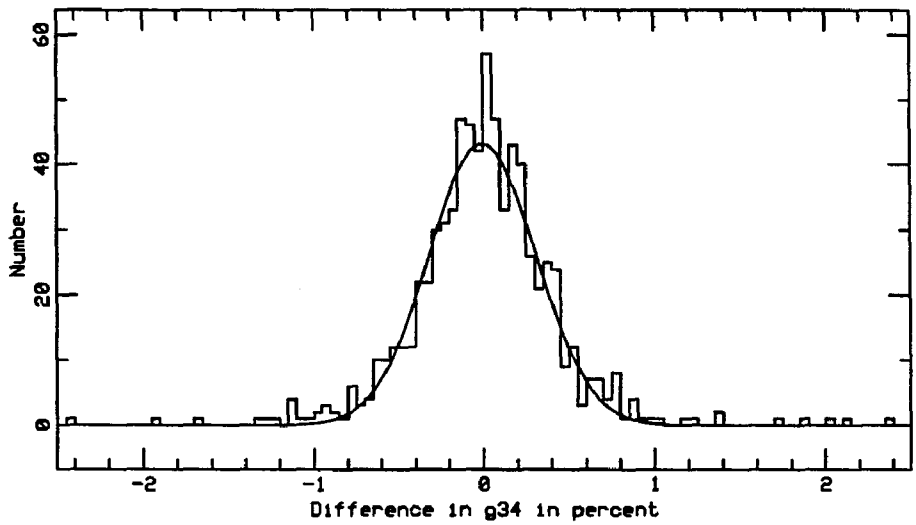

Figure 7.24: Estimate of the error in the test beam g_{34}^{B}

board, while crystals at $\theta = 24 - 13$ correspond to channels 1-12 on the other level-1 board. If there was a bad pulser/preamp channel, we would expect to see a large correction for $\theta = i, 25 - i$. Problems in the layout of the level-1 boards would also give rise to errors with the same mirror symmetry. On the other hand, problems in the layout of the preamp card, or the temperature gradients on either the preamp or the level-1 cards could produce errors that are correlated with theta, but which are not mirror symmetric. Table 7.5 shows the mean and rms of Δ_{34} for the two half barrels as a function of theta.

For the first half barrel, the scatter in the mean values is eight times larger than what we would expect if the gains were not correlated with θ. We conclude that the gains are correlated with θ and choose to correct the g_{34}'s for the first half barrel as a function θ. However, for the second half barrel the scatter is only slightly larger than we would expect for a random sample. Since there is no θ correlation for the second half barrel, we apply a single correction to all the of crystals. After the corrections are applied the rms width of the correction distribution is 0.55% for the first half barrel and 0.59% for the second half barrel. These values are consistent with our estimated error of 0.48% for the accuracy of the test beam g_{34} calculation.

First Half Barrel						Second Half Barrel					
θ	mean	rms	θ	mean	rms	θ	mean	rms	θ	mean	rms
1	0.29	0.37	24	-0.35	0.83	1	-0.31	0.61	24	-0.25	0.57
2	-0.26	0.53	23	-0.58	0.53	2	-0.23	0.64	23	-0.33	0.63
3	-0.53	0.64	22	-0.97	0.63	3	-0.23	0.55	22	-0.30	0.60
4	-0.08	0.55	21	-0.40	0.50	4	-0.16	0.60	21	-0.30	0.62
5	-1.60	0.74	20	-1.74	0.84	5	-0.36	0.59	20	-0.42	0.60
6	-0.36	0.57	19	-0.43	0.47	6	-0.26	0.52	19	-0.43	0.60
7	-0.17	0.56	18	-0.04	0.45	7	-0.39	0.64	18	-0.42	0.64
8	-0.20	0.70	17	-0.11	0.66	8	-0.38	0.60	17	-0.44	0.59
9	-0.42	0.68	16	-0.20	0.63	9	-0.31	0.63	16	-0.48	0.56
10	-0.02	0.34	15	0.23	0.49	10	-0.36	0.61	15	-0.50	0.58
11	-0.41	0.61	14	-0.04	0.48	11	-0.43	0.54	14	-0.50	0.53
12	-0.33	0.62	13	0.06	0.48	12	-0.38	0.55	13	-0.46	0.57

Table 7.5: Correction to g_{34} in percent, (Δ_{34})

7.9 Data Reduction

The data taken during 1988 consumed more than 1000 magnetic tapes, each with 2400 feet of tape at a density of 6250 bits per inch and an average of roughly 100 megabytes of data. A total of approximately 10^{12} bits of data were recorded. Rather than dealing directly with this excess of data, data summary tapes (DST's) were made. The complete set of DST's used for this thesis occupy 25 magnetic cartridges, (IBM type 3840), each with 200 megabytes of data. The data compression is on the order of 20 to 1.

The DST's contain seven files for each raw data tape, with an additional file for the 50 GeV data tape and additional files for each pedestal run. Three of the files are in text format and give diagnostic information. Three of the files contain data for monitoring the leakage current from the photodiodes, determining the temperature inside the detector, and measuring the gains of the amplifiers in the ADC. The extra file for the 50 GeV data is used to measure the gain between the high and the low energy channels of the ADC.

The file of primary interest for the calibration is the beam data file. For each calibration event which passes some minimal cuts, the ADC readings of the 25 crystals in a 5×5 matrix, the wire chamber information, and the ADC timing information is recorded. A file containing the results of a third order fit to the low energy pedestal

and a linear fit to the high energy pedestals is written for each pedestal run.

The sophistication of the data processing in the DST production was purposefully kept to a minimum, in order to reduce the chance of error. In most cases, the selected raw data is merely repackaged into a more compact form. Keeping the data in raw form has the advantage that we do not need to use a set of conversion factors, particularly the ADC electronic gains, which were not known precisely at the time, and also the raw data format tends to be very compact. The disadvantage is that the processing to convert the raw data into a useable format must be done each time the DST is used. The only conversion factors used in the DST production were the wire chamber calibrations. These had been measured in advance of the calibration, and were well understood since the same set of chambers has been in operation for several years.

The first stage of the analysis is to pick which events will be used. For this analysis, several levels of cuts were used in the event selection. The first level is the tests applied before writing the events to the DST. If it were necessary to relax one of the criteria used for this level of cuts we would need to redo the DST production, so these criteria are purposely very weak. The events are required to meet the following criteria.

- There must be no detectable errors in the raw data format.
  A significant portion of the raw data carries redundant information used to check the integrity of the data, (*i.e.* checksum and parity bits). All of this information is checked and required to be correct.

- There must be hits in the scintillators S1,S2, and S5.
  This allows only events with a particle passing through a 1 cm $\times$ 1 cm area in front of the crystal being calibrated.

- There must be hits in both Y wire chambers.
  This is necessary to reconstruct the y coordinate of the impact point at the crystal.

- There must be a hit in X chamber closest to the BGO, and in at least one other downstream X chamber.
  This is necessary to reconstruct the x coordinate of the impact point at the crystal.

- There must be a hit in at least one of the upstream chambers after the collimator.
  This is necessary to reconstruct the momentum.

- The ADC reading for the central crystal higher than the first comparator. This eliminates events in which no electron entered the central crystal. In this case either the electron missed the central crystal, a pion passed through the beam line, or the electron was absorbed between the last wire chamber and the crystal. This test is useful primarily for crystals at $\theta = 1$, where half of the particles strike the structure supporting the detector.

The next level of cuts is made in the analysis program before the data for each crystal have been grouped together. The cuts made in the DST production are refined by using the improved knowledge of the wire chamber and crystal positions gained during the analysis, as described in section 7.4. Also, a selection is made on ADC pedestals. This pedestal information was produced at the same time as the DST and therefore not available for the previous set of cuts. At this stage the events must pass the following tests.

- The reconstructed tracks must have a $\chi^2 < 2.0$.
- The electron impact point must be within 1 cm of the center of the crystal.
- The electron momentum must be within 5.0% of the beam line central momentum.
- All crystals used in the analysis for this event have pedestals higher than 10 ADC counts.

The final set of cuts are made after a preliminary fit to the data has been made. The exact cuts made are different for the single and multiple crystal calibration algorithms. The number of events per crystal passing each set of cuts for the single crystal algorithm is shown in Fig. 7.25. It is important to note that even though the second level of cuts is nominally the same for both calibration methods, the multiple crystal method requires pedestals on more crystals and therefore, the number of events passing the second level of cuts is greatly reduced. Overall, the two methods tend to accept roughly the same number of events for each crystal. The complete set of cuts gives us an efficiency of roughly 50%.

We have already discussed the main part of the event processing: conversion of the ADC reading to a voltage, pedestal subtraction, and reconstruction of the track of the incident electron. The quantities used in the calibration are the set of pedestal-subtracted ADC readings normalized by the incident momentum and the impact position of the electron at the detector.

For the single-crystal method, the only remaining step is to determine the ADC reading for electrons incident at the center of the crystal. Since, the positions of

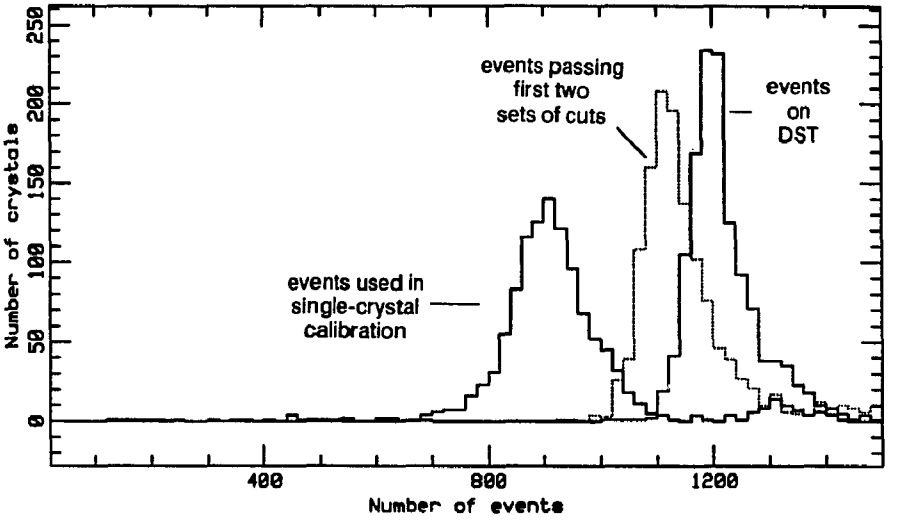

Figure 7.25: Number of events accepted at each stage in the analysis

the crystals were not known to sufficient accuracy before the calibration, we must also find the center of each crystal. The technique, as discussed in Section 7.2, is to fit the data to a function of the form

$$V(x,y) = V_0(1 - \beta[(x - x_0)^2 + (x - x_0)^2]) \tag{7.11}$$

We fit for four parameters: the ADC reading for a perfectly centered electron, V_0, the x and y of the center of the crystal, x_0, y_0, and a parameter that describes how quickly the energy falls off with distance from the center of the crystal, β. Rather than fitting the data directly, we collect events in 1.5 mm $\times$ 1.5 mm bins of the impact point. Events with ADC readings more than three standard deviations from the mean of all the data for the crystal are rejected. We perform the fit using the mean value of each bin weighted by an error equal to the RMS spread of the data points in the bin divided by the number of data points in the bin. The data for crystal ($\theta = 11, \phi = 71$) is presented in Table 7.26. The table shows the average ADC reading and number of events for each impact position bin. The center of the crystal was determined to be at $x = 5.3$, $y = 4.3$, where the units are defined in terms of the bins.

First, a preliminary fit is made using all of the bins with more than five events. This fit gives us an approximate position for the center of the crystal. We then

9	0.0	0.0	0.0	0.0	402.5	0.0	0.0	401.2	377.0	0.0
	0	0	0	0	1	0	0	2	1	0
8	0.0	357.2	401.6	393.7	405.8	404.0	406.5	405.9	392.6	391.1
	0	1	5	2	5	4	6	5	7	1
7	0.0	409.0	399.6	411.7	410.7	410.1	409.7	405.5	398.3	380.4
	0	2	5	10	15	21	21	18	18	5
6	0.0	408.9	409.0	414.4	415.5	419.4	416.4	414.6	409.3	393.8
	0	6	20	21	36	27	45	23	43	8
5	0.0	400.2	411.5	421.5	420.4	417.4	417.4	415.1	409.2	405.4
Y 0		10	23	18	32	40	52	41	58	12
4	0.0	408.0	409.3	421.7	419.5	422.5	422.7	415.6	410.6	397.0
	0	8	21	23	37	40	47	48	39	3
3	0.0	407.6	417.4	420.2	424.1	421.8	421.1	415.9	409.7	402.9
	0	6	7	14	22	28	14	20	17	3
2	370.7	400.0	420.9	408.0	425.1	420.0	419.3	414.9	402.3	407.0
	1	3	2	5	5	10	16	7	7	2
1	0.0	0.0	401.9	410.2	411.1	420.2	402.4	376.8	393.1	385.5
	0	0	4	4	1	3	7	3	2	1
0	0.0	401.6	0.0	0.0	426.3	400.9	0.0	366.0	0.0	0.0
	0	1	0	0	1	1	0	1	0	0
	0	1	2	3	4	5	6	7	8	9
						X				

Figure 7.26: ADC reading versus impact position
For each bin we give the average ADC reading in millivolts (the upper number) and the number of events (the lower number).

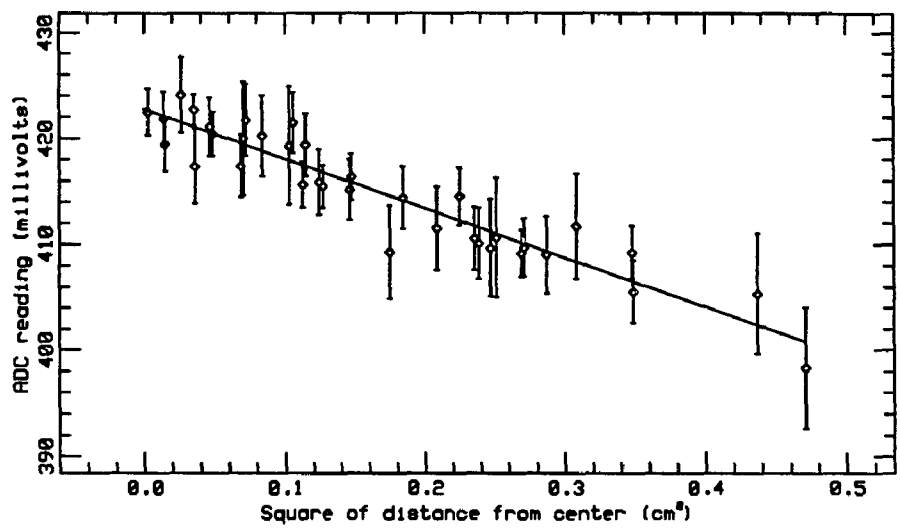

Figure 7.27: ADC reading versus distance from center

perform a second fit, using only bins which are within 6 mm of the center of the crystal and have at least eight events. The second fit is not performed unless there are at least seven bins remaining with a total of at least 100 events. The data for the bins passing the final set of cuts for crystal ($\theta = 11, \phi = 71$) are shown in Fig. 7.27. The data points are the mean and error for each bin, with the distance from the fitted center of the crystal. The line is the result of the fit. Note that the graph hides the fact that we must fit for the position of the center of the crystal.

If the fit is successful, then the crystal has been calibrated. The results of the fit, along with the tape number, run number, time and date of the first event, and position of the crystal are written to a file. The tape number and time of the events are recorded so that we can later determine the temperature of the crystal during calibration, in order to apply temperature correction to the calibration constant.

For the multiple-crystal method, the sum of the N×N matrix are calculated. A previous iteration of calibration constants must be used to make the sum, since we must sum energies rather than ADC readings. If a crystal in the matrix has a negative pedestal, it is not used in the sum. A sum is also made of the Monte Carlo expected energy for each crystal used in the sum. After the sums are made, the center of gravity of energy in the BGO is calculated. The event is rejected is the center of gravity is more than 8 mm from the center of the matrix. If the event

is retained, the center of gravity is used to make a correction of the energy. These fractional difference, δ, between the corrected energy and the MC expected energy is histogrammed.

$$\delta = \frac{E_{sum} - E_{MC}}{E_{MC}}, \tag{7.12}$$

where E_{sum} is the sum of energy in the matrix using the previous set of calibration constants and E_{MC} is the energy deposited in the matrix calculated from a Monte Carlo simulation. After all the events are collected, the mean of the histogram is calculated and used as a correction to the previous calibration constant.

7.10 Energy Reconstruction and the Calibration Constants

The distribution of temperature-corrected calibration constants determined with the single-crystal method on the 10 GeV data for the second half barrel is shown in Fig. 7.28. The calibration constants have a mean value of 0.0193 GeV/mV with an RMS deviation from the mean of 0.0014 GeV/mV. The RMS variation divided by the mean is 7%. This is much smaller than the factor of two variation in calibration constants found for crystals in prototypes of the detector[60]. The improved uniformity is most likely due to the improvements in the surface treatment and painting of the crystals.

A comparison of the results of the single and multiple crystal methods is shown in Fig. 7.29. The two sets of calibration constants use the same set of 10 GeV data for the first half barrel and have no corrections for temperature or light yield variations in the crystals. We define cc_1 to be the calibration constant from the single crystal method and cc_M to be the calibration constant from the multiple crystal method. Then fractional difference between the calibration constants from the two methods for each crystal is then

$$\delta_{cc} = \frac{cc_1 - cc_M}{cc_1} \tag{7.13}$$

The plot is a histogram of δ_{cc} for each crystal in the second half barrel calibrated by both methods. The multiple-crystal calibration constants were generated by S. Rosier with a method similar, but not exactly the same, as the multiple-crystal method described above. Rosier's calibration program uses the standard BGO group analysis routines which were developed, primarily by members of the Annecy and Rome groups, completely separately from the routines used in our analysis[61]. Comparison with Rosier's calibration is more interesting than with our own multiple-crystal calibration because different algorithms are used to find

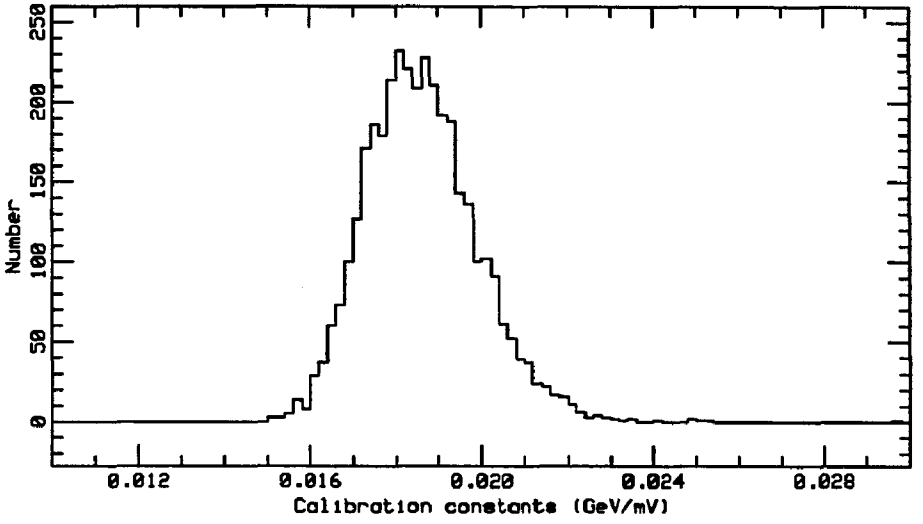

Figure 7.28: The calibration constants

pedestal fits, calculate the momentum, find the peak of the energy distribution, and select events. Our multiple-crystal calibration was used mainly to test the single-crystal calibration method before Rosier's calibration constants were available.

The agreement between the two sets of calibration constants is very good, at the level of 0.3%. The two calibration methods have different primary sources of error. The dominant error in the multiple-crystal method comes from the pedestal subtraction, while the main source of error in the single-crystal method is locating the center of the crystal. The agreement between the two methods shows that both methods can determine the calibration constants to a precision of 0.3%. However, it says little about the absolute accuracy of the calibration because there are two possible sources of error, the temperature correction and the absolute scale of the momentum calculation, which are common to both methods and therefore not tested by the comparison.

The main difficulty in reconstructing the energy of particles in the detector comes from energy leakage and light collection nonuniformities in the crystals. We have no experimental data on the leakage of energy from the BGO crystals, so we rely on Monte Carlo simulation. The simulation and the results have been described in section 7.3. We have studied the energy leakage as a function of the incident energy and of the crystal size. Both of these corrections are small, on the order of

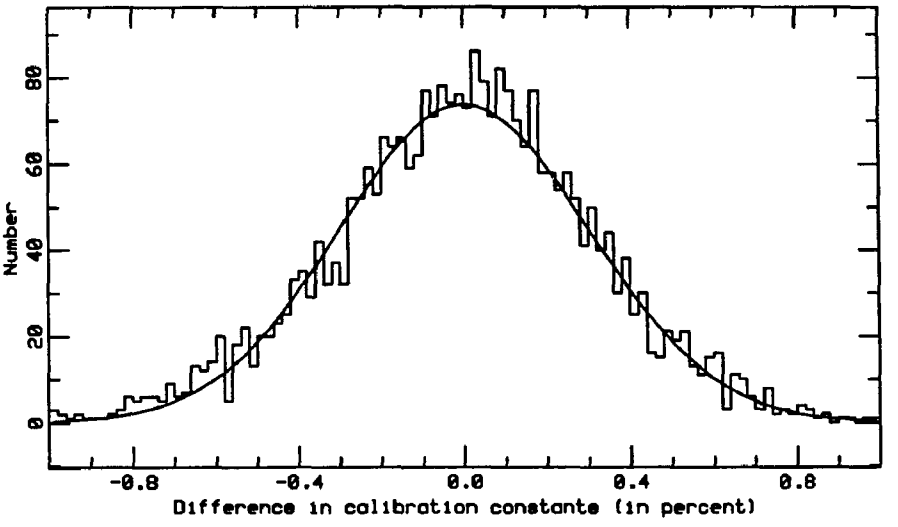

Figure 7.29: Difference between calibration constants from the two methods

a few percent, so we have chosen to treat them separately in order to simplify the analysis.

Figure 7.30 shows δ_{cc}, defined above, as a function of the θ coordinate of the crystal. A sum of 9 crystals cannot be made for $\theta = 1$ and $\theta = 24$, so the crystals at these values of θ are not included in the plot. As discussed in section 7.3, the size of the crystal is determined by its position in θ. The energy deposited in a single crystal is quite sensitive to the crystal size. On the hand, the sum of nine crystals is less sensitive, changing only by about three parts per thousand over the range of crystal sizes. We see, from Fig. 7.30, that the maximum variation in the difference of calibration constants is only 0.1%, which is within the allowed accuracy of the calibration. Since there is no systematic difference in the calibration constants as a function of θ, we conclude that we have adequately corrected for the dependence of the single-crystal calibration constant on crystal size.

Light collection non-uniformities in the crystals lead to a nonlinear light output as a function of energy. We define the calibration constants, without regard to the uniformity, at 10 GeV. When reconstructing events at other energies, we must take account of the light collection non-uniformities. As explained in section 7.3, we parameterize the light collection non-uniformity of a crystal as $R = [s(0) - s(L)]/s(L)$, where $s(x)$ is the relative light output at a distance x from the front

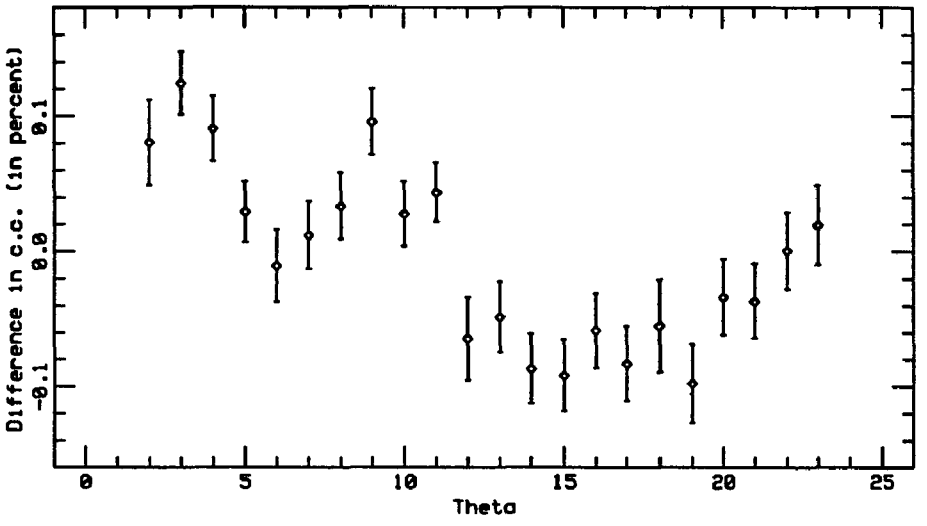

Figure 7.30: Difference between calibration constants from the two methods as a function of θ

face of the crystal and L is the length of the crystals. The parameter R is referred to as the uniformity factor of the crystal. Figure 7.31 shows the effect of light collection non-uniformities on the total light output as a function of energy as observed in the calibration data. We define cc_{10} as the single-crystal calibration constant at 10 GeV and cc_{50} as the single-crystal calibration constant at 50 GeV. The two calibration constants are calculated independently and without regard for the light collection uniformity of the crystal. We divided the crystals of the second half barrel into bins according to their uniformity factor, and calculated the mean and RMS deviation of cc_{10}/cc_{50} for the crystals in each bin. Figure 7.31 is a plot of the mean and RMS deviation of cc_{10}/cc_{50} for each bin versus the uniformity factor, R, of the bin. There is a significant shift in cc_{10}/cc_{50} due to light collection non-uniformity.

To correct for the light yield nonuniformities, we have generated a set shower of energy profiles for the crystals in a 3×3 matrix at the calibration energies. Using these profiles, and assuming a linear light yield curve, we have calculated corrections over the range of light yields for each crystal in the matrix at each energy. The result is a a correction coefficient c_R which depends on the R of the crystal and the position of the crystal within the matrix. The correction c_R was described in section 7.3. Upon including the light yield correction and the temperature correction, we arrive

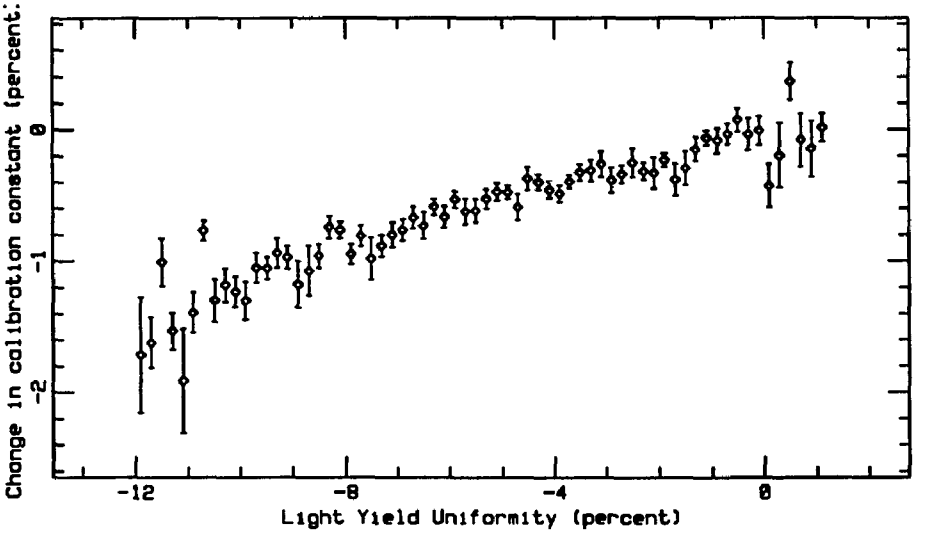

Figure 7.31: Ratio of calibration constants at 10 and 50 GeV as a function of the light yield uniformity

at the equation used to reconstruct energies

$$E = k_{10} V c_{ly}(R, d\theta, d\phi, E_{in})[1 + c_T(\bar{T} - T_{ref})] \tag{7.14}$$

This is the practical version of equation (7.4) described in the Introduction. The reconstructed energy is E, k_{10} is the calibration constant at 10 GeV, V is the ADC reading, $c_{ly}(R, d\theta, d\phi, E_{in})$ is the light yield correction, R is the light yield uniformity of the crystal, $d\theta$ and $d\phi$ are the distance of the crystal from the center of the shower, E_{in} is incident energy, c_T is the temperature correction coefficient $+1.55\%/°C$, and $\bar{T}$ is the average temperature of the crystal, and $T_{ref} = 18.0°C$ is the reference temperature.

7.11 Resolution

The final test of the calibration, and of the detector, is the energy resolution obtained. The overall resolution of the detector has two components: the width of the energy distribution for each individual group of crystals, and the variation in the mean energy between different groups of crystals. We will refer to these two components as the local resolution and the global resolution, respectively. The overall resolution of the detector, i.e. the accuracy with which we can measure the energy of any single event, is the quadrature sum of the two components. The width of individual energy distributions is determined by the energy leakage and the electronic noise. The noise performance of the electronics was not optimized before the calibration data taking, so the results presented here for the individual energy resolution should not be considered final. However, the variation in the mean energy is determined by the accuracy of the calibration and the corrections applied in the reconstruction. Therefore, the calibration results presented here are relevant to the final performance of the detector.

First, we must make an aside on how to parameterize the distribution of energy deposited in matrix of crystals. Due to the nongaussian tails of the energy distribution, the standard deviation of the distribution does not give a good estimate of its width. We have choose to calculate the resolution from the full width at half maximum of a functional form fitted to the distribution. A direct calculation of the FWHM would be dependant on the binning of the histogram and susceptible to errors, particularly at low statistics. Fitting a function is a much more robust method. The function we have chosen to use is a composite of a gaussian with an exponential tail. The function has four free parameters: the area, mean, and width of the gaussian, and the crossover point between the gaussian and the exponential. The parameters of the exponential are determined by requiring continuity in the function and its first derivative at the crossover point.

$$f = \begin{cases} A\exp[-a(x-x_0)^2] & \text{if } x \geq x_c \\ B\exp[b(x-x_0)] & \text{if } x < x_c \end{cases} \tag{7.15}$$

where the free parameters are A, a, x_0, and x_c. From continuity of the function and the first derivative, we find $b = -2a(x_c - x_0)$ and $B = A\exp[a(x_c - x_0)^2]$. The half maximum of the gaussian is at a distance $d = \sqrt{(\ln 2)/a}$ from its center. If the crossover to the exponential tail is below the half maximum point the FWHM is simply that of the gaussian. Otherwise, we have contributions to the width from both the gaussian and the exponential.

$$\text{FWHM} = \begin{cases} 2d & \text{if } x_c \leq x_0 - d \\ d + \frac{\ln 2}{2a(x_0 - x_c)} + \frac{x_0 - x_c}{2} & \text{if } x_c \geq x_0 - d \end{cases} \tag{7.16}$$

The resolution, σ, is defined in the same way that the standard deviation of a gaussian is related to its FWHM, $\sigma = \text{FWHM}/2.35$. The center of the gaussian is used as the reference energy of the distribution.

Figure 7.32 is a histogram of the energy deposited in a 3×3 matrix centered at $(\theta = 11, \phi = 71)$ for 1066 events at 10 GeV. A fit of the gaussian/exponential form described above is superimposed. The peak, x_0, of the fit is at 9.293 GeV. The width calculated from the fit is 0.089 GeV. The fit has a χ^2 of 66 for 100 degrees of freedom. The form fits the data quite well. We have also found that the form provides a good fit to the energy deposited in each individual crystal if, for crystals other than the central crystal, the exponential tail is moved to the high energy side of the gaussian.

We now have the mathematical apparatus necessary to measure the local resolution. We measure the local resolution for each 3×3 matrix by making a histogram of the energy deposited in the matrix, and applying the gaussian/exponential fit to the histogram. The local resolution of the matrix is then given by the FWHM/2.35 of the fit. The average local resolution at the three calibration energies is shown in Table 7.6. Folded into our measured resolution are errors in the momentum measurement, electromagnetic shower fluctuations, and the electronic noise. We have calculated the contribution to the local resolution from each of these three sources, and listed the results in the Table 7.6.

The contribution to the resolution from momentum measurement is calculated from the measured position resolution and alignment accuracy of the wire chambers. The contribution due to pedestal subtraction is calculated as the average deviation of the pedestals from the pedestal fits, as discussed in section 7.7. It is important when calculating the error in pedestal subtraction that we take into account the correlation between the pedestals within the 3×3 matrix. The magnitude of fluctuations in the electromagnetic shower created by the incident electron is calculated

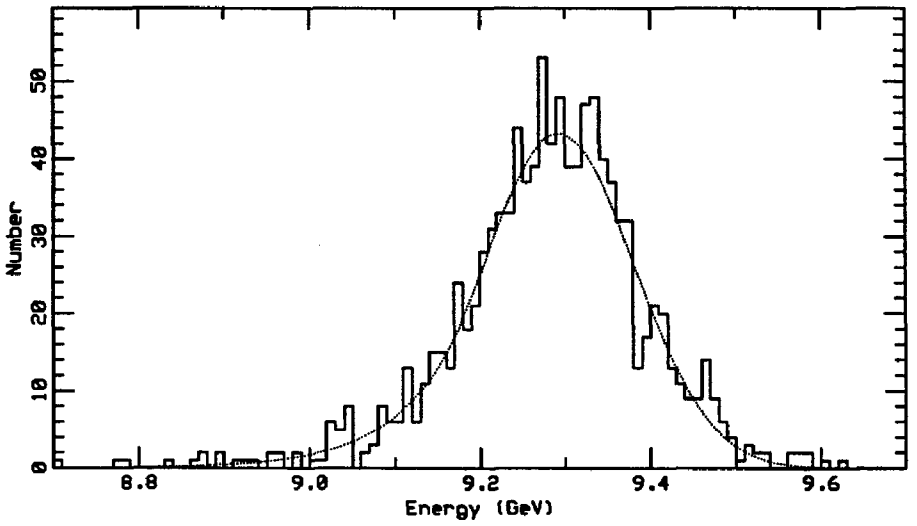

Figure 7.32: Energy in a 3×3 matrix of crystals

via a Monte Carlo simulation, as described in section 7.3. Since we only measure the energy deposited in nine BGO crystals, fluctuations in the electromagnetic shower will vary the fraction of the incident energy which is deposited in those nine crystals.

At energies of 10 GeV and below, the local resolution is dominated by the pedestal subtraction. The pedestals have a complicated shape, as described in Section 7.7, which will not be present in the synchronous environment of LEP. Optimization of the electronic noise, together with the simpler pedestal subtraction, should greatly improve the resolution of the detector below 10 GeV. At 50 GeV we approach the limit on the local resolution of a sum of nine crystal due to electromagnetic shower fluctuations.

To measure the global resolution of the detector, which is essentially a test of the calibration and reconstruction procedures, we examine the distribution of the mean reconstructed energy throughout the detector. The procedure is almost the same as for determining the local resolution. We make a histogram, for each 3×3 matrix of crystals, of the energy deposited in the matrix, and then apply the gaussian/exponential fit to each histogram. To find the local resolution, we calculated the width of the fit. However, for the global resolution we are interested in the peak of the fit rather than the width. The global resolution is given by the RMS spread in peak energies of the fits.

Energy	2 GeV	10 GeV	50 GeV
Measured Local Resolution	2.4%	0.98%	0.47%
Calculated Local Resolution	2.6%	0.92%	0.45%
Momentum Measurement	1.5%	0.35%	0.16%
Pedestal Subtraction	1.7%	0.34%	0.07%
Shower Fluctuations	1.3%	0.78%	0.42%

Table 7.6: Local resolution of 3×3 matrices of crystals

Figures 7.33 and 7.34 show the distribution of peak energy for each 3×3 matrix in the second half barrel. For reconstruction, we used calibration constants derived from the 10 GeV data, and have applied corrections for temperature, light collection non-uniformity, and the momentum of the incident particle. All of the corrections have been discussed previously, see sections 7.6, 7.10, and 7.4.

The distribution of peak energies at the calibration energy gives us an estimate of the accuracy of the calibration technique. Since we calibrate with single crystals but reconstruct with sums of nine crystals, it not a completely trivial procedure. There are errors from the pedestal subtraction and the light yield correction in all crystals except the central crystal. Also, only one calibration constant is used for each crystal. The crystals which were calibrated more than once use different data for calibration and reconstruction and therefore test the repeatibilty of the calibration procedure. The dispersion, given by the standard deviation of the superimposed gaussian fit, at 10 GeV is 0.24%. This is in agreement with the expected accuracy of the calibration procedure.

The distributions at energies other than the calibration energy test the entire reconstruction procedure, including the calibration and all of the corrections. We find the standard deviation to be 0.45% at 50 GeV, and 0.87% at 2 GeV. Table 7.7 shows the estimates of the factors contributing to the accuracy of the reconstruction. The accuracy of the temperature correction was discussed previously. The error of the light yield correction was taken to be the quadrature difference between the global resolution for all crystals and the resolution for a sample of crystals with the same light yield uniformity. Inaccuracies in the electronic gains are important only for ADC readings on the low energy amplifier chain, since the gain at the high/low crossover is the only gain not known to better than a few tenths of a percent. We calculate the error due to the gains by multiplying the average gain error by the fraction of energy on crystals with ADC readings on the low energy chain. The error is larger at 2 GeV because the readings on the central crystal, and therefore

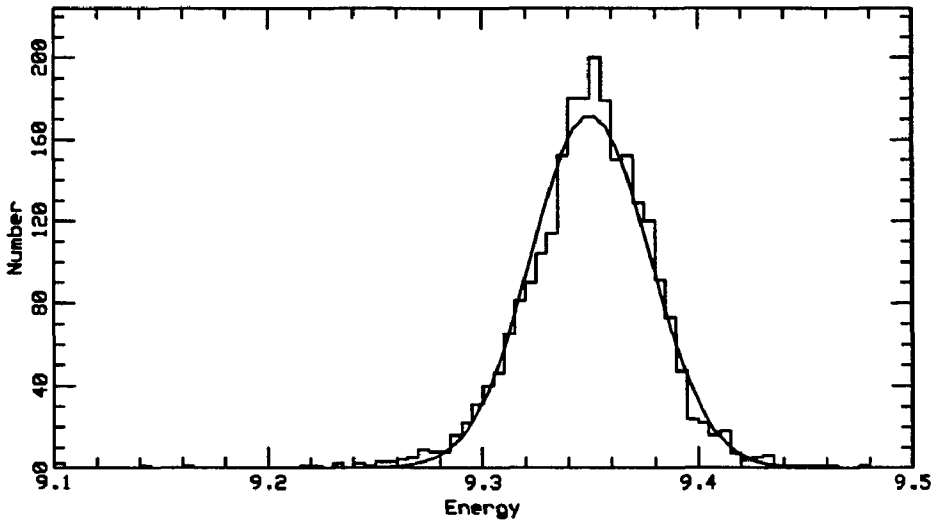

Figure 7.33: Mean energy of the sum of nine crystals 10 GeV

all of the energy, are on the low energy chain.

We chose not to use the local resolution measured during the calibration to calculate the overall resolution of the detector because the electronic noise was not optimized before the calibration and because the momentum resolution of the test beam magnetic spectrometer will not (directly) effect the local resolution at LEP. For a more accurate estimate of the final electronic noise, we use the noise measurements made during low energy tests of calorimeter done after the electronic noise was optimized. The contributions to the local resolution are the electronic noise and the shower fluctuations. The global resolution is the measured global resolution for 2 and 50 GeV. For 10 GeV, the measured global resolution does not include the effects of errors in the temperature correction or the electronic gains, so we add the calculated temperature and gain errors in quadrature. The light yield correction is defined to be the identity at 10 GeV, so we need not include errors in the light yield correction. The overall resolution is the quadrature sum of the local and global resolutions. The results for the overall resolution are shown in Table 7.8.

We conclude that the accuracy of the reconstruction procedure is equal to or better than the local resolution of the detector at all energies. Therefore the accuracy of the calibration procedure is within the limits set by the intrinsic resolution of the detector. It is clear that we have achieved the goal of constructing an electro-

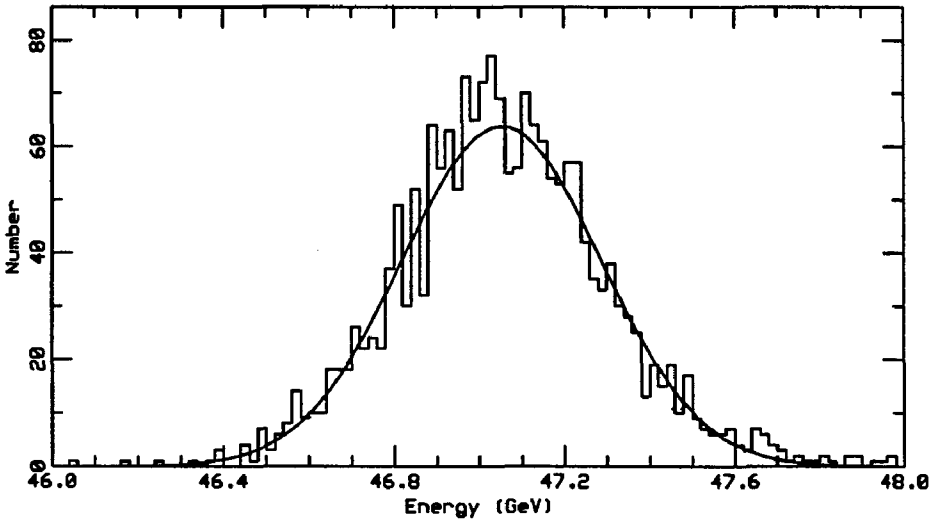

Mean energy of the sum of nine crystals 50 GeV

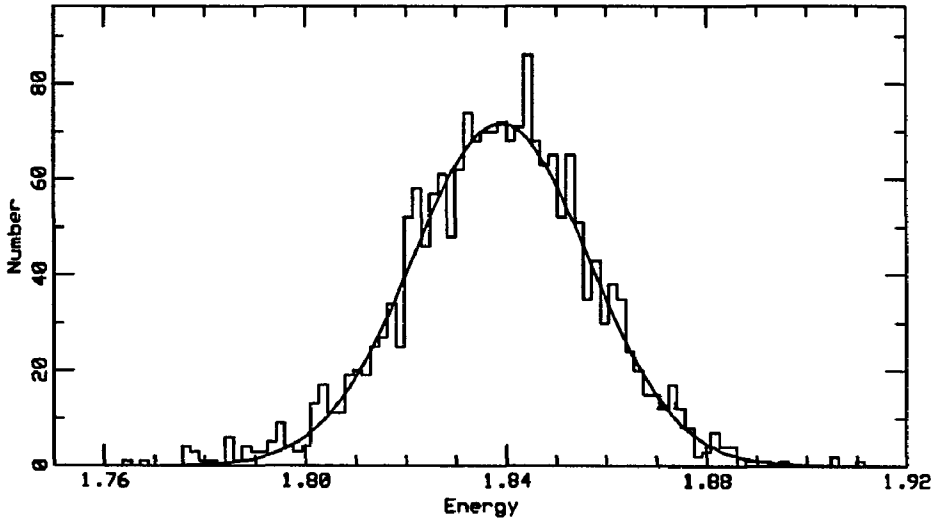

Mean energy of the sum of nine crystals 2 GeV

Figure 7.34: Reconstructed energies

Energy	2 GeV	10 GeV	50 GeV
Measured Global Resolution	0.87%	0.24%	0.45%
Calculated Global Resolution	0.77%	0.24%	0.48%
Calibration	0.24%	0.24%	0.24%
Temperature	0.35%	-	0.35%
Light output	0.4%	-	0.2%
Electronic gains	0.5%	-	0.1%

Table 7.7: Global resolution of 3×3 matrices of crystals

Energy	2 GeV	10 GeV	50 GeV
Pedestal Subtraction	0.35%	0.09%	0.02%
Shower Fluctuations	1.29%	0.78%	0.42%
Local Resolution	1.34%	0.79%	0.42%
Global Resolution	0.87%	0.44%	0.45%
Overall Resolution	1.59%	0.90%	0.62%

Table 7.8: Overall resolution of 3×3 matrices of crystals

magnetic calorimeter with energy resolution of better than 1% for energies greater than 10 GeV.

Appendix A

The ADC cards

In this appendix we discuss a few of the more technical details of the ADC cards. For ease of presentation, some of the information presented in section 6.5 is repeated.

The output of the preamplifier is a voltage pulse with a risetime of 270 ns, corresponding to the decay time of BGO scintillation light, and a decay time of 800 μsec. The maximum peak voltage the preamp can produce is approximately 10 volts.

The front end of the ADC circuit is shown in Fig. A.1. The output from the preamp is fed into the pole-zero, which eliminates the 800 μsec tail. The pole-zero output feeds into three paths: the low energy amplifier chain, the high energy amplifier chain, and the BGO analog output for the trigger. There is an additional RC filter between the pole-zero and the first active component in each of the three paths. The three filters have matched time constants of 0.5 μsec. The RC filters smooth the leading edge of the pole-zero output pulse, increasing the rise time of the pulse in order to allow the relatively slow op-amp in the low energy chain to accurately follow the shape of the pulse. The output voltage of this circuit for an arbitrary input pulse is given by

$$V(t) = N \int_0^t dt\, V_0(t-\tau)(Ae^{-\tau/a} + Be^{-\tau/b}) \qquad \text{(A.1)}$$

where a and b are the two time constants, $a = 0.17$ μsec and $b = 1.0$ μsec. The relative weight of the two components are $A = 2.4$ and $B = -0.4$. We have taken the average pole-zero resistance of 1MΩ. The time constants vary only slightly with the pole-zero resistance. The longer time constant is important because it determines how long we must integrate the signal. The voltage output for an input pulse with a 270 ns risetime and an 800 μsec exponential tail is shown in Fig. A.2. The solid curve shows the output when the pole-zero time constant is exactly equal 800 μsec. The peak voltage, V_p, occurs at 0.7 μsec. The maximum rate of change

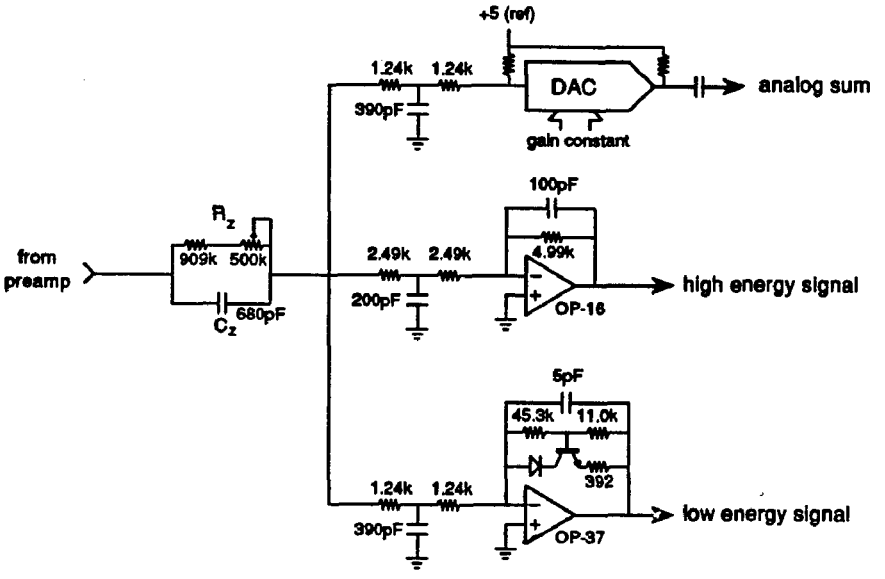

Figure A.1: The ADC front end

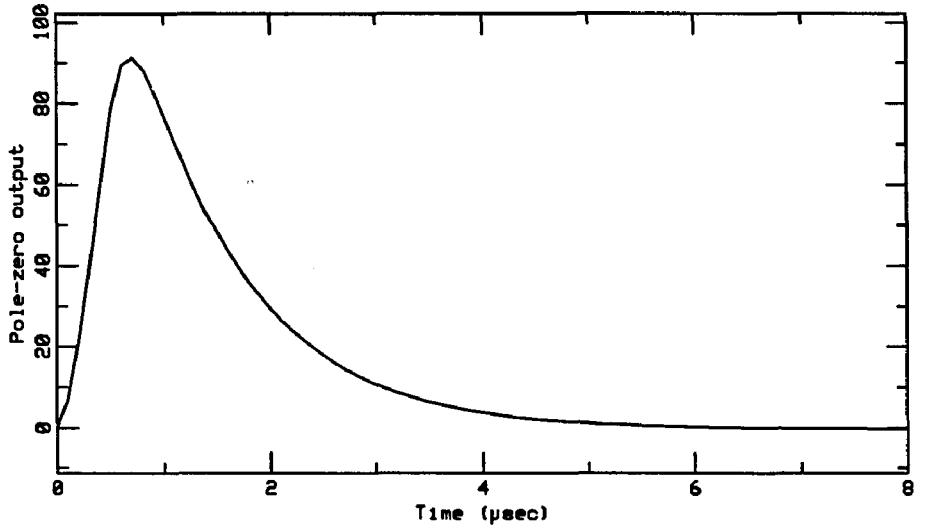

Figure A.2: Shape of pulse at the input to the low energy amplifier

of the voltage is at the beginning of the pulse, where the $dV/dt = V_p/0.5$ μsec. The tail of the pulse has a 1.0 μsec decay time constant.

The resistors in the filters are chosen to split the current from the pole-zero in the ratio 2:1:2 for the (low energy chain):(high energy chain):(analog output). Less current is needed in the high energy chain because it deals only with high current signals. For the circuit to be linear it is necessary that the current sharing between the three chains is always constant. If any of the three chains draws a disproportionate amount of current, the signal in all three chains will be nonlinear.

The low energy op-amp (an OP-37, manufactured by PMI [62]) has a pair of diodes across its inputs which prevent the voltage difference between the two inputs from exceeding 0.6 volts. This voltage limiting is necessary to protect the circuitry of the OP-37. If the voltage difference does exceed 0.6 volts, then current, drawn from the pole-zero, will flow through the diodes. Making the entire circuit nonlinear. In an ideal op-amp the voltage difference at the inputs is kept at zero by adjusting the op-amp output. To accomplish this, the output voltage must be able to change at a rate, the 'slew-rate', equal to the rate of change of the input signal multiplied by the gain of the feedback network. The OP-37 has a slew-rate of 17 V/μsec and the feedback network has a gain of 22.7. Using the maximum rate of change of the input signal calculated above, we find the maximum input signal that the low

energy chain can handle is 1.0 V.

This is above the range of signals that digitized on the low energy chain. However, it is also necessary that there be no voltage limiting on high energy signals, since the nonlinearity also affects the high energy chain. To satisfy this requirement, the low energy amplifier has a somewhat complicated feedback network. The gain for small signals the feedback network operates as a resistance of 56.3 kΩ. However, for signals greater than the voltage drop of the transistor (= 0.6 V), the transistor turns on. This reduces the feedback resistance to a few hundred ohms and reduces the gain to less than unity. Therefore, for signals larger than 0.6 V, the gain is reduced so that the opamp can maintain a zero voltage difference at its inputs.

The small signal gain of the low energy amplifier is 22.7. The amplifier for the high energy channel is a simple inverting opamp. The two 2.49 kΩ series input resistors and the 4.99 kΩ feedback resistor produce a gain of −1.00.

The integrators and sample/hold circuits for the low and high energy chains are very similar. The circuit is shown in Fig. A.3. The circuitry in the dashed box clips large amplitude signals. It is used only on the low energy integrator. It is simply a reverse biased zener diode that goes into conduction when the signal reaches 7 volts. It prevents the low energy integrator from going into saturation on signals large enough for the high energy chain. The signal is integrated on the 470 pF capacitor. The low energy integrator has a 1 kΩ input resistor and a 470 pF capacitor. Their product is a scale factor for converting the time integral of the input voltage pulse to an output voltage. The corresponding values for the high energy comparator are 1.5 kΩ and 470 pF. The ratio of the two *RC* factors gives the low energy chain an additional gain of 1.5 relative to the high energy chain. The overall relative gain between the high and low energy chains is 34. The total gain of the front end is such that a pulse from the preamp with the normal shape (an exponential with an 800 μsec tail) and an amplitude of 10 volts at the peak should produce a signal of 10 volts at the output of the high energy integrator.

The integrator is reset, via the SW7510 FET switch, before each event. This minimizes the integrated noise, and eliminates any 'memory' of previous signals. The timing of the integrator and the sample/hold will be discussed below. When the FET switch is turned off, the 30 volt gate swing across the gate to source capacitance injects charge into the integrator capacitor. This puts the integrator output at around −200 millivolts. The injected charge will be the same on each reset as long as the supply voltages remain constant. Any variation in the injected charge will show up as noise. The opamp for the integrator has a trimmer potentiometer to adjust its input offset voltage thereby controlling the integration slope for zero input signal.

While not digitizing the DAC used in the ADC chain is set to a predetermined

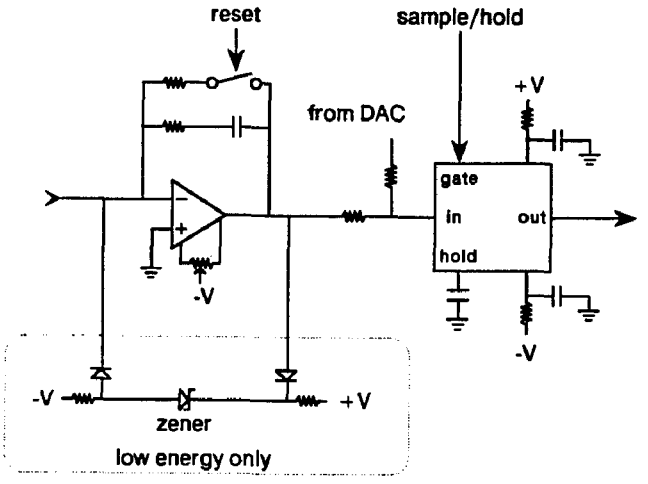

Figure A.3: The resetable integrator and sample/hold

voltage. This voltage is added to the signal at the integrator output via a resistor voltage divider. The same voltage is added to both the high and low energy chains. This allows the level of the pedestal, (the voltage reading for zero signal input), to be adjusted. In principle, it is possible to adjust the pedestals on all channels to be uniform on either the high or the low energy chain, but not both.

The output of the integrator feeds into the sample/hold, an SMP11. The only additional circuitry is the RC pairs on the power lines. In runs taken during the summer of 1986, it was found that there is a shift in the pedestals when many channels powered by the same power supply all have large signals. This shift was traced to a drop in the voltage on the ± 15 volt lines due to the large current drawn when switching from sample to hold. The *RC* pairs filter the power lines. The capacitors are bypass capacitors. The resistors isolate the SMP11's in order to reduce the effect the current surge has on the voltage lines that power the other components in the ADC.

The output of the sample/hold goes to a amplifier/comparator chain as shown in Fig. A.4. The chains for high and low energy are identical. Each chain has three comparators and two amplifiers. The amplifier are non-inverting with a gain of about four. The diodes on the outputs prevents the voltage on the inputs of the comparators from going negative, which could damage the single supply LM339's.

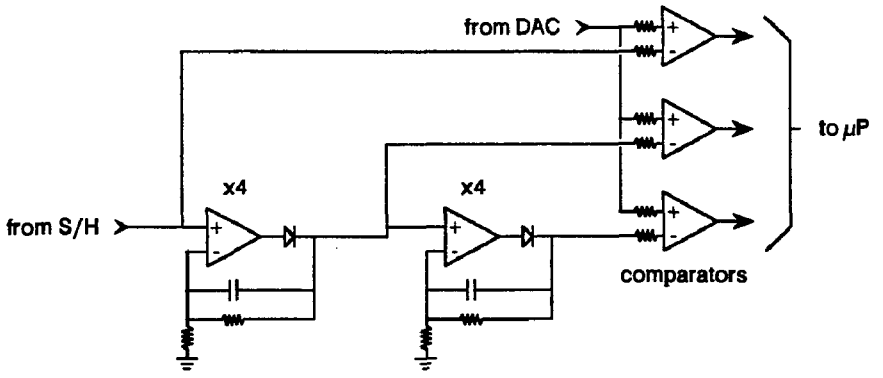

Figure A.4: The amplifier/comparator chain

The other input to each comparator is the output of the 12-bit DAC.

The output signal of the low energy op-amp is feed into an integrator and a DC feedback circuit. The DC feedback attenuates the signal and averages it with a long time constant of about one second. The feedback is designed to balance the leakage current of the photodiode and variation in the preamp FET operating point in order to maintain a stabilize the DC output level at the integrator. The input to the feedback is clipped by a pair of diodes at ± 1 volt to prevent the feedback from large signals from interfering with the normal ac coupling of the signal into the pole-zero. This signal corresponds to roughly 200MeV. Since the feedback voltage is mainly determined by the photodiode leakage current, it is also sent to a comparator so that the voltage can be digitized. This provides a way to monitor the photodiodes.

Separate from the high and low energy chains, there is another DAC used to provide an analog signal for the L3 trigger, see Fig. A.1. The output of the DAC is a current proportional to the input voltage times the value at the DAC digital inputs. After the crystals have been calibrated, the values on the DAC can be adjusted so that each crystal gives an uniform output current for a given amount of energy deposited in the crystal. The 50 kΩ resistor to the +5 volt reference gives the zero reference and brings the DAC into its linear operating region. The resistor at the output keeps the output voltage within the compliance range. The output is coupled via a 1 μF capacitor into the analog sum load. The load is 100 Ω to

ground, so there is a 100μsec coupling time constant. The outputs of the DAC's are summed in groups of three on each ADC card. Since the DAC output is a current, the summation is done by just connecting the outputs.

The last piece of analog circuitry on the ADC card is the test pulse gate. The test pulse is generated on the level-1 driver/receiver board. The gate either passes the pulse to the preamp or leaves an open circuit. The gate consists of two FET switches and two resistor dividers. The input signal is attenuated by a factor of 30 by the time it reaches the preamp input. The large attenuation (and the two FET switches) ensures that no noise feeds into the preamp from the test pulse line when the test pulse is disabled. The attenuation also prevents crosstalk between test pulses on nearby channels.

The digital portion of the ADC card is primarily the 6305 single chip microcomputer. It requires very little digital support circuitry, only a square wave clock produces by a crystal and a pair of CMOS inverters. There are also four CMOS inverters used to buffer signals coming onto the board. The 6305's are connected in a token passing ring with common data bus and handshaking lines. The data bus and handshaking lines are used to communicate with the next higher level of the readout. The token ring is used to determine which of the 6305's is currently in control of the data bus and handshake lines. Details about the token ring and the communication protocol will be discussed later. The resistors on the digital signal lines provide some protection in case of component failures by isolating each 6305 signal line.

The ADC cards are connected with a 40 conductor cable in groups of five to form larger token passing rings. The data bus, handshaking, and control lines are common to all the microprocessors within a ring and a carried on the cable. The token snakes its way through the ring via connections between successive microprocessors. The token passes from card to card via the cable. One end of the cable is connected to a level-1 driver/receiver card. This card is the connection to the upper levels of the readout.

Appendix B

The Data Collection System

The second and third levels of the readout are both housed in VME crates located in the counting room. There are sixteen level2's and one level3 in each crate. The same type of microcomputer is used for both levels, a commercially built Mizar 8115. Attached to each processor is a driver/reciever board that allows the level-2 to communication with level-1, or level-3 to communicate with level-4.

The Mizar 8115 is a VME compatible 16 bit microcomputer based on the Motorola 68010 microprocessor. The board has 512k bytes of 'dual-ported' ram, *i.e.* the ram can be read either by the local processor or from the VME bus. It also has a Signetics 68681 serial inteface/timer chip, an RS-232 serial interface on the front panel, sockets for two EPROMS, and VME bus interface circuitry. The only hardware difference between the level2 and 3 microcoputers is that the level3 runs slightly faster (12.5 MHz clock cycle) than does level2 (10.0 MHz). A complete description of the Mizar boards can be found in the Mizar 8115 reference manual.

We chose the Mizar because it had the lowest cost of the VME microcomputers that satisfied our requirements, and also because it is a 'half-height' VME module. Each slot in a VME crate accepts 'full-height' modules. The half-height Mizar has a connector which can be used to carry the local microprocessor signals to an attached, custom built, half-height card in the same VME slot. This allows us to directly attach the driver/receiver boards to the processors, keeping the pair of boards within one VME slot. The resulting system is compact and elegant. With a standard full-height processor we would have needed another crate to house the D/R boards and another set of cables to make the connection. With the Mizar board we require only the one VME crate and no cable. Eliminating the cable eliminates the related noise problems.

The communication between levels 2 and 3 is done via the VME bus. A standard VME crate is split, both logically and physically into two halves. A full height VME board has two 96 pin connectors mating it to the back plane. All 96 signals on the upper connector are used. They form the standard (or unextended) VME bus. On

the lower connector, 64 of the pins are not bussed, but rather passed directly through to another connector on the rear of the backplane. These 64 pins are intended to be used by modules which need to communicate with external devices. We use these 64 pins for the connection between the level-1 and level-2 driver/receivers. The remaining 32 pins are bussed on the backplane. In normal usage they form the extended (32-bit) VME bus. Since the Mizars make no connection to the extended bus, we chose to redefine the signals for our own needs.

The signals on VME bus are described in the VME bus manual. The bus is completely asynchronous. An asynchronous bus allows each component to operate at its maximum speed, while not requiring any minimum speed. In our system, the three types on modules on the bus have widely varying performance. The level-3 processor has a 80 ns cycle time, the level-2 has a 100 ns cycle time, and the static memory board has a 215 ns cycle time. The asynchronous bus allows the faster components of the system to operate at full speed when the slower modules, notably the memory board, are not in use.

The VME bus has the capability for multiple bus masters, *i.e.* any module on the bus can take control of the bus and request data from any other module. We have found that multiple masters leads to contention problems that greatly reduce the overall performance of the system. The basic point is that in a non-synchronous system, each would-be master will request the bus without regard to the requests of the other processors. After a processor has requested the bus it is essentially stopped until it receives control. With 16 processors each constantly accessing the bus, the delays are enormous. To eliminate bus contention, we allow only the level-3 processor to control the bus. In fact, the backplanes is configured so that if a level-2 attempts to request the bus, the request is blocked and an error light goes on. The large memory of the mizar boards allows us to buffer data in each level-2's memory until the level-3 is ready to accept it. The function of the level-3 is essentially to move the data from the level-2 memory into the fifo which then passes it to the level-4. The level-3 does not buffer any data. Additional buffering would only consume extra clock cycles, since the level-2 buffer is large enough for our purposes.

Another module on the VME bus is the CMOS static ram module, referred to as 'the ram-disk'. The ram-disk uses very low power CMOS static rams with a special stand-by mode. The module is battery backed, so that it retains data even when the power is shut off, emulating the function of a disk drive. The ram-disk is used to store the programs and parameters for the readout system. Initially the code and constants were downloaded from the DAQ computer, via level-4, at the start of each run. This was found to be too time consuming. Instead the code and constants for all the processors in the entire crate are stored on the ram-disk, where they are rapidly accessible. The ram-disk itself is downloaded only when necessary

to change the data taking parameters. Interestingly, the current consumed by the (512k bytes of) CMOS ram is so low that the battery life is not reduced by powering the ram. The internal leakage current of the battery is higher than that required by the ram.

Appendix C

The Driver/Receiver Boards

This appendix presents schematics, and some explanation to make the schematics more useful, of three printed circuits boards which were designed, laid out, constructed, tested, and installed as part of the author's thesis work. Other people who did major parts of the circuit design were Jon Bakken, Peter Denes, and Richard Sumner. The final version of the T-board was laid out by D. Wright.

C.1 Driving and Receiving

Pairs of driver/receiver cards are used to carry signals between the level-1 boards, mounted on the detector, to the level-2 processors, located in the counting room, as shown in Fig. C.1. All of the signals handled by the BGO readout system above the level-1 ADC cards are digital. However, the TTL and CMOS logic used within the level-1 and level-2 circuits is incapable of transmitting signals over long distance. The driver/receiver cards are used to translate standard logic levels to and from the differential signals used to transmit information over the hundred meters of cable between the level-1 boards and the upper levels of the readout. The standard logic levels are fed into high-current drivers which translate the levels into differential pairs of high current signals. Each differential pair is carried along a cable, which is terminated at both ends to minimize reflections, to a receiver. The receiver, after

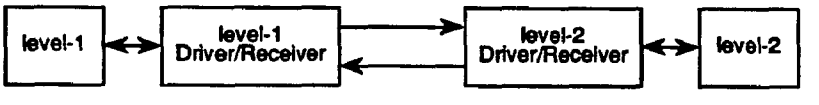

Figure C.1: Position of the driver/receiver boards

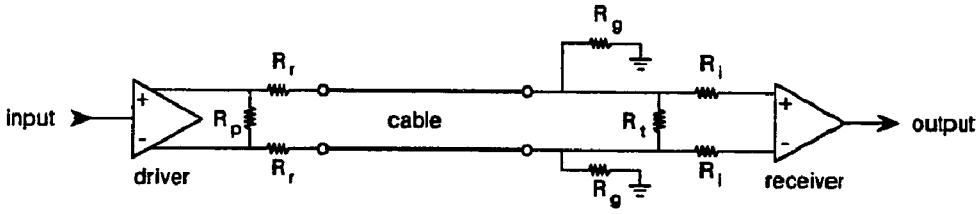

Figure C.2: A differential driver/receiver pair

applying some hysteresis to add noise immunity, translates the pair of differential signals back into a standard TTL logic level. To improve the reliability of the system, all of the differential signals are unidirectional. The bidirectional data bus and token lines are translated into pairs of unidirectional signals for tranmission across the cable.

The average switching cycle of the signals is a few microseconds. The rise and fall times of the differential signals are on the order of a few tens of nanoseconds. There are 32 differential pairs per ring, and 128 rings in the BGO barrel. With 4096 signal pairs on 128 cables packed into a cross-sectional area of 80mm × 160mm carried over a distance of a few hundred meters, there is a large problem with cross talk between the signals. In addition, the long length of the cables and the small cross section of each wire combine to produce a significant resistance along each wire (about 100Ω), complicating the problems of termination.

The transmission of digital signals over long distances is a common problem and there are commercially available integrated circuits designed for the task. We chose to use a differential transceiver pair manufactured by Advanced Micro Devices, the Am26LS31 and Am26LS32. Given this pair of chips, the only remaining question is how to terminate the ends. The most general termination circuit is shown in Fig. C.2. The driver/receiver (D/R) boards have all of the connections necessary to implement this general circuit. The boards were constructed with only R_t, R_r, and R_i. This subset was found to give good performance in tests done by Peter Denes. In case other problems arise, it is still possible to install the other resistors.

The cable has a characteristic impedance of 105Ω , a DC resistance of 67Ω per 100 yards, and a cross-section of 28 AWG. A relatively low cable impedance was chosen to ensure adequate current flow in the wires in order to reduce the sensitivity to magnetically induced voltage differences and capacitively coupled noise. At LEP

we will have 125 meters of cable, at the test beam we had 100 meters. The speed of propagation of voltage differences in wire is about 0.8c, therefore it takes on the order of half a microsecond for the signal to travel the length of the wire. Since this is significantly longer than the rise time of the signals, reflections of the signals at the ends of the cable will be important. To minimize reflections, the termination resistance should be matched to the characteristic impedance of the cable.

The minimal termination necessary is the resistor R_t. Since the 2632 has high impedance inputs, the resistor supplies all of the termination. The value of R_t chosen, 120 Ω, matches the impedance of the cable. The outputs of the driver are nominally at 0.6 and 3.0 volts. Neglecting R_τ and the cable resistance, the current required from the driver is 20 milliamps, which is within the allowed current range of the 2632. Therefore us can use the same termination with a short cable and R_τ. This is the configuration used for driver/receiver circuits that are separated by short distances, notable the T-board circuit discussed below. The actual current in long distance driver/receiver circuits is significantly lower, about 7 milliamps.

The pair of resistors R_τ are designed to prevent reflections at the outputs of the driver. A value of 49 Ω was selected so the series resistance of the pair matches the impedance of the cable. The other set of resistors currently in use, the R_i, protect the receiver inputs. Since the inputs are high impedance the additional resistance has negligible effect on the performance.

We have provided space on the circuits for two other resistors, which we not found necessary to install. The resistor R_p could be used to provide a load for the driver when no cable is plugged in. We have found that the driver has sufficient internal protection that this is not necessary. The resistor is not effective for terminating reflections because of the driver has low impedance outputs. The R_g provide for additional noise immunity. If R_t were removed they could also be an alternate termination scheme on the receiving end.

C.2 The Level-1 Driver/Receiver

The level-1 driver/receiver is a relatively simple board. It basically consists of a large number of copies of halves of the circuit described in the preceding section. There are three completely independent D/R sections on each printed circuit card. The power is separately supplied to each. The cards are mounted in a smaller box on the sides of the boxes containing the ADC cards. There is a short cable connecting each D/R to its associated set of ADC cards, and a long cable that runs from the level-1 D/R to the corresponding level-2 D/R. In addition to the D/R circuit, the cards have additional buffering on the signals sent to the rings, some circuitry to perform loop back tests of the data bus and token lines, and a pulser that supplies

the test pulse to the ADC cards.

Since the 6305 microcomputer on the ADC cards is a CMOS device, the logic levels of the 2631's and 2632's must be translated to CMOS levels. This done with a CMOS buffer chip, the 74HC241. There are two resistors between the CMOS buffer and the ring. The resistor in series prevents the drivers in the 241's from burning themselves out if the output line is accidently shorted to ground. The pull-up resistor brings the lines up to the default state (+5 volts) when they are not begin driven by a 6305.

The loop back for testing the data bus lines is very simple. Since all of the differential signal lines are unidirectional, the data on the level-1 data bus is always send up to level-2, regardless of whether the data bus is being driven by level-1 or level-2. The level-2 microcomputer can test the data bus lines by simply sending down a given value and then checking that the same value is returned. No additional circuitry is necessary on the level-1 D/R to perform the test.

The token loop back is slightly more complicated. In order to test both bi-directional token lines with a single test line, we must switch which token is connected to the test line when the token direction switches. The logic of the token drivers is arranged so that the token line being sent from level-2 to level-1 is always returned to level-2 via the TOKRET line, regardless of the current token direction.

The test pulse circuit is shown in Fig. C.3. The DAC takes its data inputs from the same data bus used by the ADC cards. The data is placed on the bus and then the test pulse strobe is sent, latching the data into the DAC. The DAC will continuously generate the selected voltage. The 1458 op-amps generate +V and -V. The test pulse is generated when the state of the test pulse gate line is toggled making the AD7512 analog switch between +V and -V. Capacitive coupling on the ADC cards converts the voltage step into a current pulse. The voltage pulse is sent to the ADC cards by the LF356 and the LH0002. The two chips act together as a single op-amp. The LH0002 is a high current buffer which can produce 100 milliamps over a voltage range of ± 10 volts. This large current output is necessary to overcome the large attenuation of the test pulse on the ADC cards and the preamp. The test pulse is not precise enough to be a good monitor of the stability of the ADC. It is used to provide a check that each ADC channel is functioning.

C.3 The Level-2 Driver/Receiver

The level-2 driver/receiver (L2DR) is the connection between the level-2 microcomputer and the ADC cards (via the level-1 D/R). In addition to the D/R circuits described above, the board also contains a 68230 parallel interface and timer chip for interfacing to the level-2 microprocessor, circuitry for reading trigger signals off the

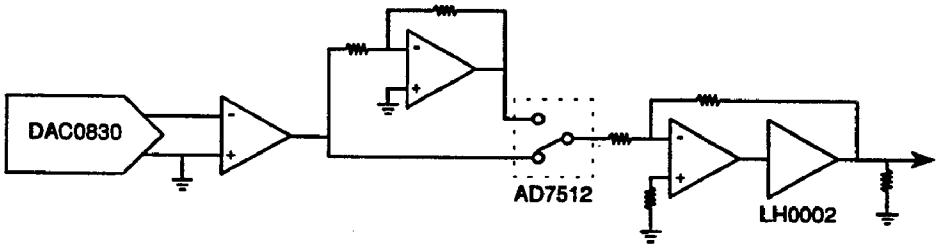

Figure C.3: The test pulser

backplane of the (VME) readout crates, and a circuit that performs handshaking with the level-1 microprocessors in hardware.

The 68230 is a standard part that is commonly used to add I/O ports to the 68010 microprocessors used in the level-2 microcomputer. It provides a large number of programmable output lines and requires only minimal support circuitry. It also provides a programmable timer, which we have wired to cause an interrupt after a programmable delay. Of particular interest on the 68230 are the handshake lines. These can be programmed to perform a variety of handshaking protocols independent of the microprocessor.

We have made use of this feature to do handshaking with the level-1 processors without requiring the attention of the 68010. Prototypes of the readout system polled the handshake lines via software. This wastes significant amounts of time because the level-1 processor cannot send the next byte of data until the level-2 is finished with the previous one. A faster solution is to implement the handshaking in hardware. However, the readout system is very susceptible to noise on the handshake lines. If noise causes the level-2 to read the wrong state on the handshake line, synchronization with the level-1 processors will be lost. This normally results in the rest of the data for that event being lost. The noise observed in the prototypes was mainly very short 'glitches'. The solution to the glitch problem was to always read twice, and accept that the line had changed state only if both reading were the same. This technique is referred to as 'deglitching'. To perform the handshaking in hardware, it is necessary to add the deglitching with hardware.

The circuit of Fig. C.4 implements the same algorithm in hardware. H0 is the input handshake line which is sampled constantly. If it differs from the counter output Q2 then the counter begins to count. The counter must count for four cycles before Q2 changes state. If H0 reverts to it previous state before the count reaches

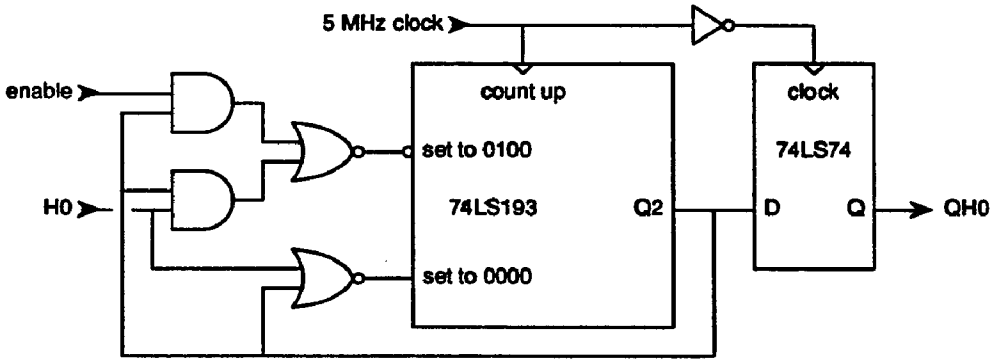

Figure C.4: The hardware deglitcher (PROE circuit)

four, the counter will be reset. If the count reaches four without H0 reverting to its previous state, then the Q2 will take the same state as H0 and the count will stop. Then we have the same situation as we had initially, but with the polarities reversed. Since H0 must remain in a given state for four counts before the counter changes state, glitches will be ignored. The output of the flip-flop QH0 is a 'de-glitched' version of H0. This circuit is referred to as the Princeton Read Out Engine, or PROE.

For flexibility, there is a small amount of additional circuitry on the L2DR that controls whether the handshaking is done in hardware or software. There is also a latch which is necessary to do a loop back test of the data lines between the level-1 and 2 D/R's. The latch is needed because the data bus lines on the 68230 are bidirectional, unlike the D/R lines. To do the test one drives the data bus with level-2, sets the latch, changes the bus direction, and then reads the data off the latch. In normal use the latch is always set to transparently pass data.

As previously discussed the hierarchy of the readout is split at the level-2 crate. The data readout from the ADC cards is sent up through the upper levels, but the trigger and timing signals are connected directly to the main trigger system. The L2DR makes the connection to the trigger system for levels 1 and 2. The trigger and timing signals, and also the BGO readout system event number, are passed along a special bus on the lower half of the VME backplane. A list of these signals is given in Table C.1.

The event number is passed thru standard TTL buffers and is read directly

We use Bn to indicate line n on the lower VME backplane. Pins B2, B12, B22, and B31 are ground connections. Pins B1, B13, and B32 are +5 volt connections. The power connections are defined by the international VME standard. The other pins are used for signal specific to the L3-BGO readout system. Pins B3-B20 (excluding power pins) carry signals EV0 thru EV15 which form a sixteen bit event number placed on the bus by T-board. The signals are positive logic: a high signal indicates a 1 and a low signal indicates a 0. The remaining pins are control signals and are desribed below.

B21 LEVEL3-GO is sent from the T-BOARD to the level-3 processors. It cannot be read by the level-2 processors.

B23 LEVEL2-GO is sent from the level-3 processor to the level-2 processors. It can also be read by the T-board.

B24 LEVEL2-AFU is the OR of the AFU signals from all the level-2 processors in the crate. It can be read by level-3 and the T-board. LEVEL2-AFU = L indicates that at least one level-2 processor is AFU'ed.

B25 6305BUSY is the OR of the busy signals of all of the 6305s connected to level-2 boards in the crate. 6305BUSY = L indicates that at least one 6305 is busy.

B26 68KBUSY is the OR of the busy signals of all the 68K level 2 processors in the VME crate. 68KBUSY = L indicates that at least one level-2 processor is busy.

B27 INTERRUPT is the interrupt line for the level-1 and level-2 processors. The interrupt occurs on the negative-going edge. Level-1 requires a positive-going edge for the interrupt, so this signal is inverted before being sent to the level-1 D/R boards.

B28 TEST PULSE GATE is the gate for the test pulse produced on the level-1 D/R boards. It is passed through the level-2 boards unchanged.

B29 SAMPLE/HOLD is the sample/hold for the analog electronics on the level-1 boards. It is passed through the level-2 boards unchanged.

B30 RESET PREAMP is the reset for the analog electronics on the level-1 boards. It is passed through the level-2 boards unchanged.

			Event number				
signal	pin	signal	pin	signal	pin	signal	pin
EV0	B3	EV1	B4	EV2	B5	EV3	B6
EV4	B7	EV5	B8	EV6	B9	EV7	B10
EV8	B11	EV9	B14	EV10	B15	EV11	B16
EV12	B17	EV13	B18	EV14	B19	EV15	B20

Table C.1: Signals on the level-2 auxiliary VME backplane

by the level-2 processor. The other input signals to the L2DR are buffered and sent to the L1DR. The trigger interrupt is also sent to the level-2 processor. The three other inputs signals are timing signals for the ADC cards. It is important that the circuitry involved with these signals be kept to a minimum in order to minimize timing errors due to variations in propagation delays. The level-2 can neither monitor nor control these signals. The L3OUT signal is buffered and sent to the level-2 processor. It allows a quick means of communication between the level-2 and 3 processors. The busy signals and the error line from all of the level-2's are open collector OR'ed on the backplane. The busy signal and also the interrupt for each ring of ADC cards can be disabled by level-2 microcomputer that controls the ring.

C.4 The T-Board

The trigger and timing signals, as shown in Table C.1, are brought on the auxiliary bus in the readout crate by another card, which, for historical reasons, is called the termination board or T-board. The T-board is basically just another set of driver/receiver chips that interface with an analogous module in the trigger system. The majority of the printed circuit real estate is taken up by driver and receiver chips. Since the two boards are located in the same room, the full termination scheme used on the L1DR and L2DR is not needed. Only the single resistor R_t, see Fig. C.2, is used. The signals are buffered before being put on the auxiliary bus using standard TTL buffers in accord with VME standards.

The signals which indicate that the readout system is busy, however, are slightly more complicated. The busy signals from each readout crate must be OR'ed together. Since differential signals can not be open-collector OR'ed, the busy's make a daisy chain through the readout crates. One readout crate is at the beginning of the chain. The T-board for that crate collects all of the busy signals generated by processors within the readout crate and passes the combined signal to the next crate in the chain. The T-board in the next crate collects the busy signals within the crate and OR's them with the signal from the first T-board. This signal is sent to the following crate in the chain which OR's in its own busy signals. The last T-board in the chain sends the signal, which is the logical OR of every busy signal in every crate, to the trigger system.

The T-board also has circuitry to generate a latched busy and a VME reset pulse. These are both controlled by the level-3 processor via commands on the VME bus. The latched busy is used during system start-up to prevent spurious triggers before the system is initialized. The reset is used to put all of the level-2 (and the level-3) processors into a known state. In addition, the board has a large

number of light emitting diodes to provide diagnostic information about the state of the readout.

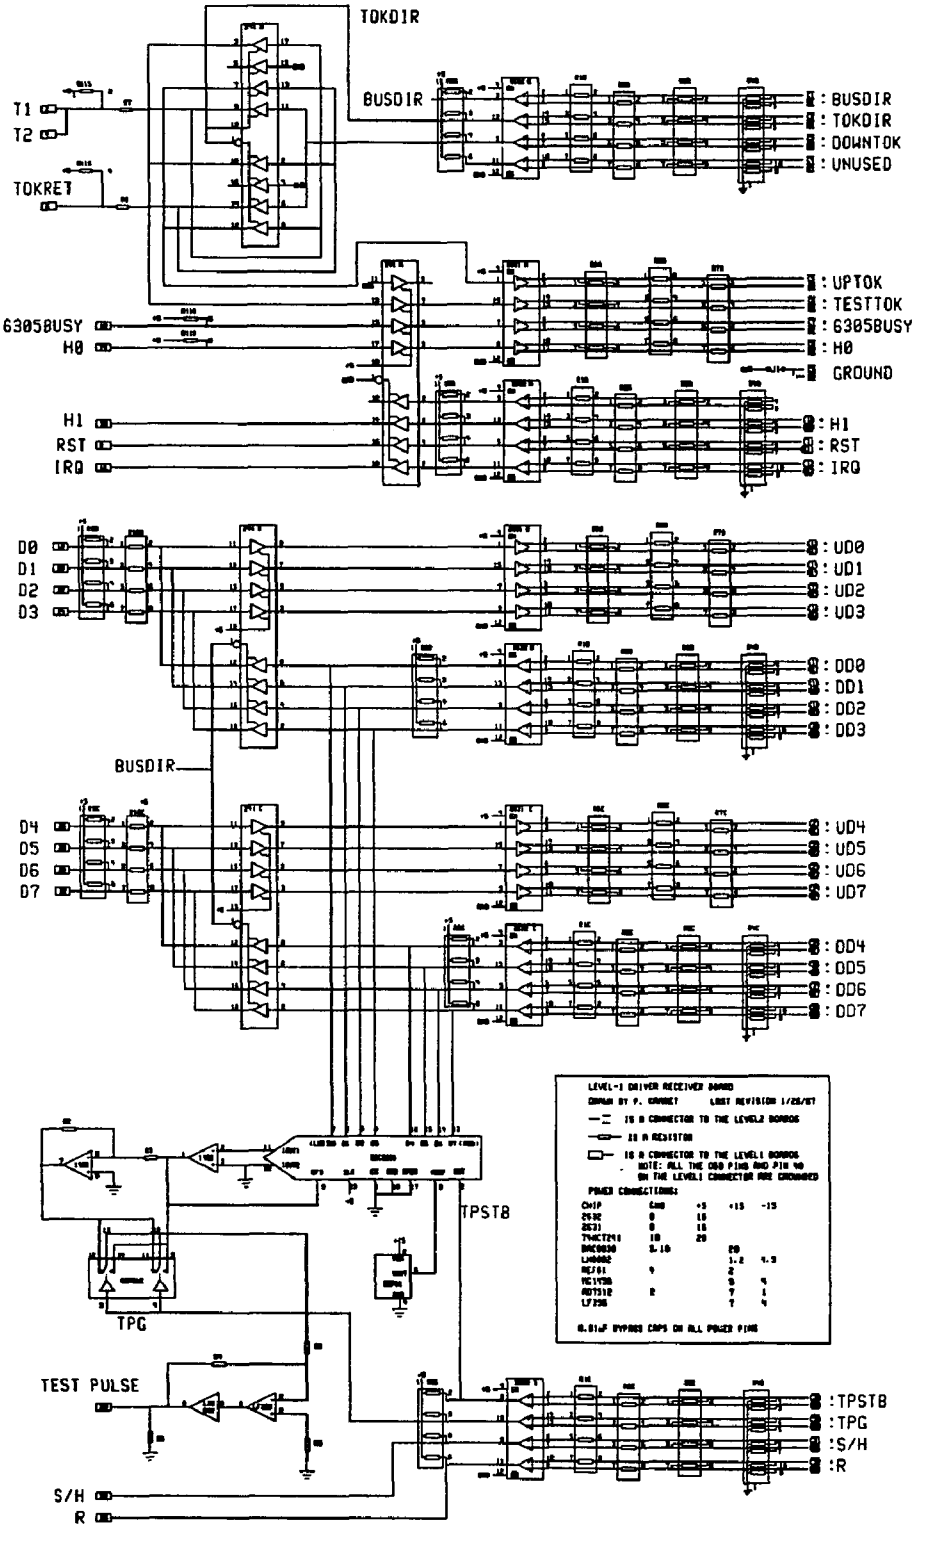
TOKDIR
BUSDIR
T1
T2
TOKRET
: BUSDIR
: TOKDIR
: DOWNTOK
: UNUSED
: UPTOK
: TESTTOK
: 6305BUSY
: H0
: GROUND
6305BUSY
H0
H1
RST
IRQ
: H1
: RST
: IRQ
D0
D1
D2
D3
: UD0
: UD1
: UD2
: UD3
: DD0
: DD1
: DD2
: DD3
BUSDIR
D4
D5
D6
D7
: UD4
: UD5
: UD6
: UD7
: DD4
: DD5
: DD6
: DD7
TPSTB
TPG
TEST PULSE
S/H
R
: TPSTB
: TPG
: S/H
: R

Parts needed for one channel of the Level-1 driver/receiver:

Integrated circuits:

Part	Description	Number	Pins
2631	driver	3	16
2632	receiver	5	16
74HC241	cmos tri-state buffer	4	20
DAC0830	latching DAC	1	20
LH0002	current source	1	10
REF01	10V voltage reference	1	8
MC1458	dual op-amp	1	8
AD7512	analog switch	1	14
LF356	op-amp	1	8

Resistors in SIP packages

Label	Number	Description
Required resistors:		
R1,R2	10	8 pin sip isolated 1K ohm Bourns 4608X-102-102
R3	5	8 pin sip isolated 120 ohm Bourns 4608X-102-121
R8	5	6 pin sip pull-up 1K ohm Bourns 4606X-101-102
R9	2	6 pin sip pull-up 20K ohm Bourns 4606X-101-203
R10	2	8 pin sip isolated 470 ohm Bourns 4608X-102-471
R11	1	8 pin sip pull-up 20K ohm Bourns 4608X-101-203
R6,R7	6	8 pin sip isolated 47 ohm Dale 3-470G
Optional resistors:		
R4	5	10 pin sip pull down 2.2K ohm
R5	3	8 pin sip isolated 120 ohm

Individual resistors:

Label	Number	Description
R1,R2,R6	3	2.5K ohm 1% tolerance
R3,R4	2	5.0K ohm 1% tolerance
R5	1	100 ohm (0.4 inch size)
R7,R8	2	470 ohm (0.3 inch size)

Capacitors:

21 0.01uF bypass capacitors (0.1 inch size)

Connectors:

1 VME type 64 pin connector, male, right angle, solder
1 40 pin, no locking flanges, male, right angle, solder
1 8 pin Molex power connector, locking, right angle
with white plastic spacer mounted underneath

DIP socket for LH0002:

1 10-pin dip socket

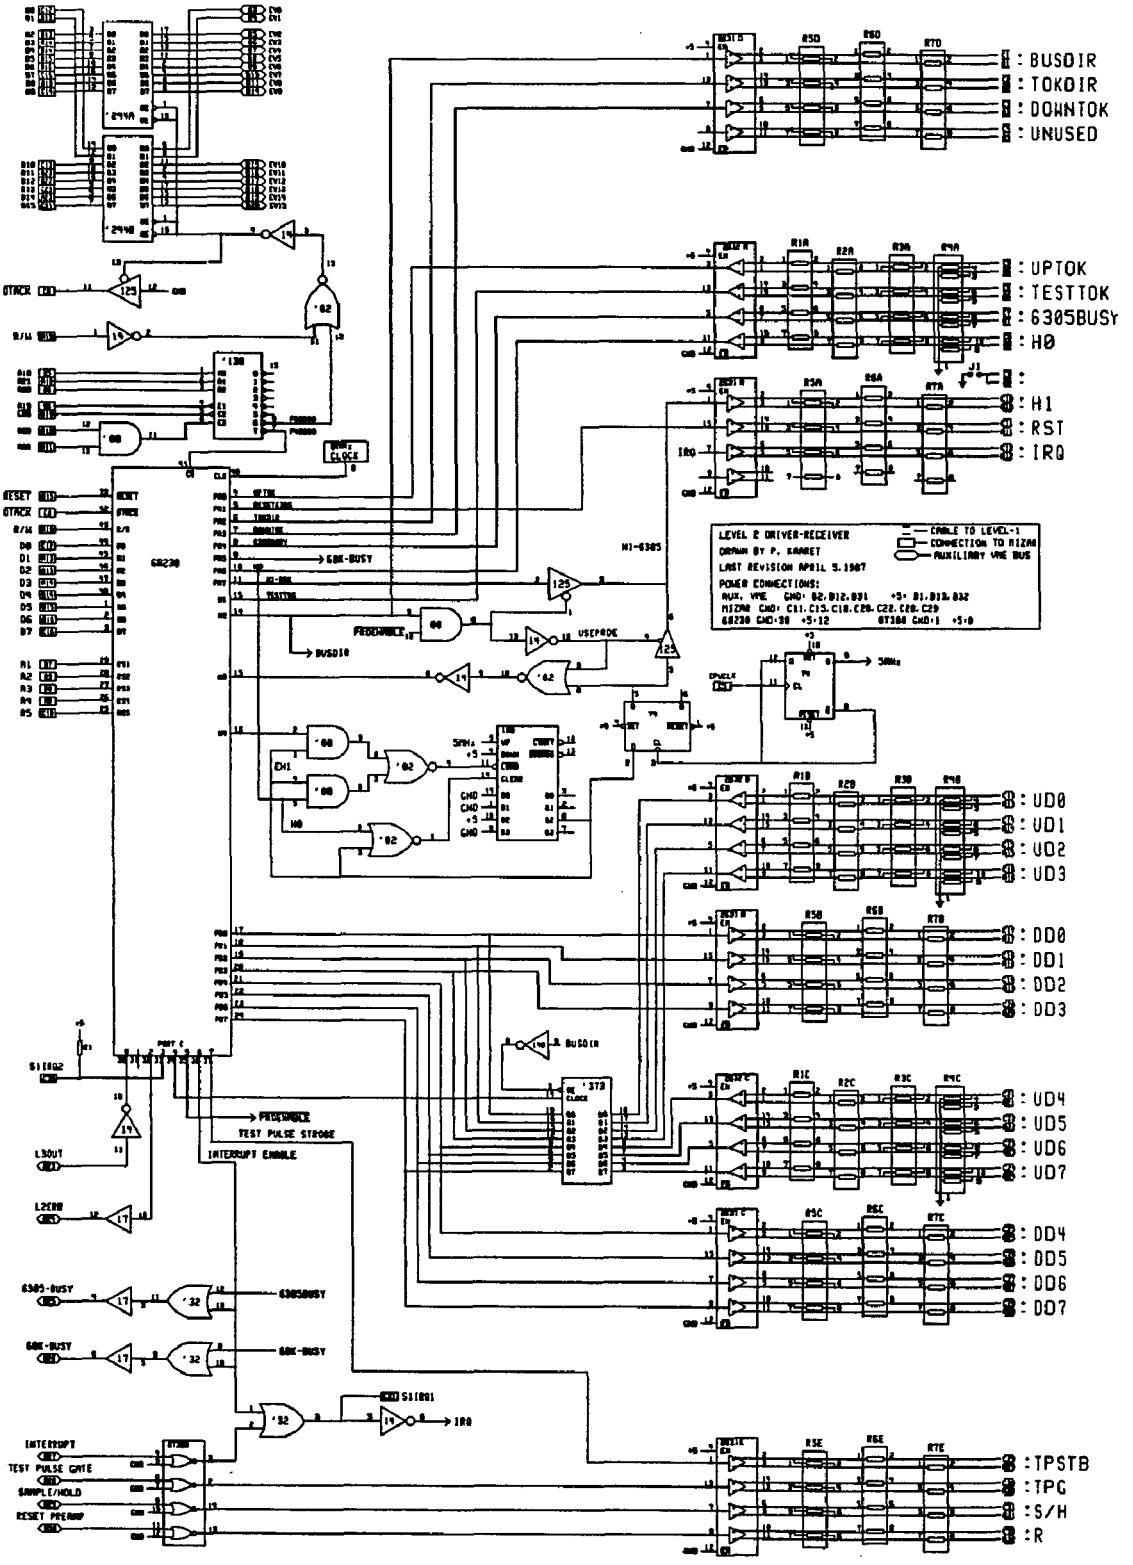
: BUSDIR
: TOKDIR
: DOWNTOK
: UNUSED
: UPTOK
: TESTTOK
: 6305BUSY
: H0
: H1
: RST
: IRQ
: UD0
: UD1
: UD2
: UD3
: DD0
: DD1
: DD2
: DD3
: UD4
: UD5
: UD6
: UD7
: DD4
: DD5
: DD6
: DD7
: TPSTB
: TPG
: S/H
: R
LEVEL 2 DRIVER-RECEIVER
DRAWN BY P. BARRET
LAST REVISION APRIL 5,1987
POWER CONNECTIONS:
CABLE TO LEVEL-1
CONNECTION TO RIZAR
AUXILIARY VME BUS
BUSDIR
TOKDIR
DOWNTOK
UNUSED
UPTOK
TESTTOK
6305BUSY
H0
H1
RST
IRQ
TPSTB
TPG
S/H
R
UD0
UD1
UD2
UD3
DD0
DD1
DD2
DD3
UD4
UD5
UD6
UD7
DD4
DD5
DD6
DD7

Parts needed for one Level-2 driver/receiver board:

Integrated circuits:

Part	Description	Number	Pins
68230	Parallel I/O chip	1	48
2631	driver	5	16
2632	receiver	3	16
74S02	nor gate	1	14
note that the 74S02 must be S			
74LS08	and gate	1	14
74LS14	Schmitt trigger inverter	2	14
7417	high current buffer	1	14
74LS32	or gate	1	14
74LS74	flip-flop	1	14
74LS125	quad tri-state buffer	1	14
74LS138	3 to 8 decoder	1	16
74LS193	counter	1	16
74LS244	octal tri-state buffer	2	20
74LS373	octal latch w/tri-state output	1	20
8T380	bus receiver	1	14
8 MHz clock generator		1	

Resistors in SIP packages

Label	Number	Description
Required resistors:		
R1,R2	6	8 pin sip isolated 1K ohm Bourns 4608X-102-102
R3	3	8 pin sip isolated 120 ohm Bourns 4608X-102-121
R6,R7	10	8 pin sip isolated 47 ohm Dale 3-470G
Optional resistors:		
R4	3	10 pin sip pull down 2.2K ohm
R5	5	8 pin sip isolated 120 ohm

Individual resistor:

RI	1	3.3k ohm 1/8 watt

Capacitors:

C	22	0.01uF bypass capacitors (0.1 inch size)
C2	1	5.0uF bypass capacitors (0.1 inch size)

Connectors:

P2	1	VME type 96 pin, male, right angle, solder
S1	1	VME type 96 pin, female, right angle, solder

DIP socket for 68230:

1	48 pin solder socket

TERMINATION BOARD
DESIGN BY P. KENNEY and D. URICH
LAST REVISION DECEMBER 15, 1987
UPPER VME (P1) CONNECTOR
BUS VME BUS, LOWER (P2) CONNECTOR
CABLE TO CAMAC, LOWER (P2) CONNECTOR
IN - DAISY CHAIN CONNECTOR
OUT - DAISY CHAIN CONNECTOR
POWER:
LOWER (P2) CONNECTOR:GROUND: B2,B12,B31 +5: B1,B13,B32
UPPER (P1) CONNECTOR:GROUND: A1,C3
B13B1: GROUND | +5: B

Parts needed for one T-board:

Integrated circuits:

Part	Description	Number	Pins
2631	driver	2	16
2632	receiver	7	16
74LS08	and gate	1	14
74LS14	Schmitt trigger inverter	2	14
7417	high current buffer	2	14
74LS27	nor gate	1	14
74LS30	8 input nand gate	2	14
74LS123	one shot	1	16
74LS125	tri-state buffer	1	14
74LS138	3 to 8 decoder	1	16
74LS244	octal tri-state buffer	2	20
74LS245	octal tri-state buffer	3	20
8T380	bus receiver	1	14

Resistors in SIP packages

Label	Number	Description
Required resistors:		
RS1	7	8 pin sip isolated 120 ohm Bourns 4608X-102-121
RS2	2	6 pin sip pull-up 1K ohm Bourns 4606X-101-102

Individual resistors:

RI	1	3.3k ohm 1/8 watt resistor

Capacitors:

C	22	0.01uF bypass capacitors (0.1 inch size)

Connectors:

P1,P2	2	VME type 96 pin, male, right angle, solder

Bibliography

[1] J. I. Friedman and H.W. Kendall, "Deep Inelastic Electron Scattering", Ann. Rev. Nucl. Science **22**, 203-254 (1972).

[2] J. H. Christenson *et al.*, "Observation of Massive Muon Pairs in Hadron Collisions", Phys. Rev. Lett. **25**, 1523-1526 (1970).

[3] S. D. Drell and T.-M. Yan, "Massive Lepton-Pair Production in Hadron-Hadron Collisions at High Energies", Phys. Rev. Lett. **25**, 316, 902(E) (1970).

[4] K. J. Anderson *et al.*, "Evidence for Longitudinal Photon Polarization in Muon-Pair Production by Pions", Phys. Rev. Lett. **43**, 1219-1222 (1979).

[5] E. L. Berger and S. J. Brodsky, "Quark Structure Functions of Mesons and the Drell-Yan Process", Phys. Rev. Lett. **42**, 940-943 (1979).

[6] J. J. Aubert *et al.*, "Experimental Observation of a Heavy Particle J", Phys. Rev. Lett. **33**, 1404-1406 (1974).

[7] J.-E. Augustin *et al.*, "Discovery of a Narrow Resonance in e^+e^- Annihilation", Phys. Rev. Lett. **33**, 1406-1408 (1974).

[8] J. G. Branson *et al.*, "Production of the J/ψ and $\psi'(3.7)$ by 255-GeV/c $\pi^\pm$ and Proton Beams on C and Sn Targets", Phys. Rev. Lett. **38**, 1331 (1977).

[9] G. E. Hogan *et al.*, "Comparison of Muon-Pair Production to the Quark-Antiquark Annihilation Model", Phys. Rev. Lett. **42**, 948-951 (1979).

[10] C. B. Newman *et al.*, "Determination of the Pion Structure Function from Muon-Pair Production", Phys. Rev. Lett. **42**, 951-955 (1979).

[11] C. Biino *et al.*, "J/ψ Longitudinal Polarization from πN Interactions", Phys. Rev. Lett. **58**, 2523-2526 (1987).

[12] K. J. Anderson *et al.*, "Production of Muon Pairs by 225-GeV/c $\pi^\pm$, $K^\pm$, $p^\pm$ Beams on Nuclear Targets", Phys. Rev. Lett. **42**, 944-947 (1979).

[13] C. E. Adolphsen, Ph.D. dissertation, U. of Chicago (1985) unpublished.

[14] H. Abramowicz *et al.*, "Determination of the Gluon Distribution in the Nucleon from Deep Inelastic Neutrino Scattering", Z. Phys. **C12**, 289-295 (1982).

[15] J. S. Conway, Ph.D. dissertation, U. of Chicago, (1987) unpublished.

[16] J. Kubar *et al.*, "QCD Corrections to the Drell-Yan Mechanism and the Pion Structure Function", Nucl. Phys. **B175**, 251-275 (1980).

[17] S. Palestini, Ph.D. dissertation, Princeton University, (1984) unpublished.

[18] S. Palestini *et al.*, "Pion Structure as Observed in the Reaction $\pi^- N \rightarrow \mu^+\mu^- X$ at 80 GeV/c", Phys. Rev. Lett. **55**, 2649-2652 (1985).

[19] M. Gluck *et al.*, "Gluon Contribution to Hadronic J/ψ Production", Phys. Rev. D **17**, 2324-2331 (1978).

[20] C. E. Carlson and R. Suaya, "Hadronic production of heavy-quark bound states", Phys. Rev. D **18**, 760-774, (1978).

[21] Y. LeMoigne *et al.*, ",Measurement of Hadronic Production of the $\chi_1^{++}(3507)$ and the $\chi_2^{++}(3553)$ through their Radiative Decay to J/ψ", Phys. Lett. **113B**, 509-512 (1982).

[22] J. Badier *et al.*, "Experimental J/ψ Hadronic Production from 150 to 280 GeV/c", Z. Phys. C **20**, 101-116 (1983).

[23] K. J. Anderson *et al.*, "Production of Muon Pairs by 225-GeV/c $\pi^\pm$, $K^\pm$, $p^\pm$ Beams on Nuclear Targets", Phys. Rev. Lett. **42**, 944-947 (1979).

[24] C. Biino *et al.*, "An Apparatus to Measure the Structure of the Pion", Nucl. Instr. and Meth. **A243**, 323-347 (1986).

[25] A. S. Carroll *et al.*, "Absorption Cross Sections of $\pi^\pm$, $K^\pm$, p and $\bar{p}$ on Nuclei between 60 and 280 GeV/c", Phys. Lett. **80B**, 319-322 (1979).

[26] J. P. Alexander, Ph.D. dissertation, U. of Chicago (1985) unpublished.

[27] J. G. Heinrich, Ph.D. dissertation, U. of Chicago, (1988) unpublished.

[28] S. Weinberg, "A Model of Leptons", Phys. Rev. Lett. **19**, 1264-1266 (1967).

[29] C. Rubbia, "Experimental Observation of the Intermediate Vector Bosons W^+, W^-, and Z^0", Rev. Mod. Phys. **57**, 699-722, (1985).

[30] S. L. Glashow, J. Iliopoulos, and L. Maiani, "Weak Interactions with Lepton-Hadron Symmetry", Phys. Rev. D. **2**, 1285-1292 (1970).

[31] A. Salam, pp. 367 in Elementary Particle Theory (ed. N. Svartholm), Almquist and Wiksells, Stockholm, (1968).

[32] S. L. Glashow, "Partial-Symmetries of Weak Interactions", Nucl. Phys. **22**, 579-588 (1961).

[33] F. Halzen and A. D. Martin, Quarks and Leptons, Wiley, New York, (1984).

[34] E. R. Cohen and B. N. Taylor, "The 1986 Adjustment of the Fundamental Physical Constants", Rev. Mod. Phys. **59**, 1121-1148 (1987).

[35] G. Bardin *et al.*, "A New Measurement of the Positive Muon Lifetime", Phys. Lett. **137B**, 135-140 (1984).

[36] J. E. Kim *et al.*, "A Theoretical and Experimental Review of the Weak Neutral Currents: a Determination of its Structure and Limits on Deviations from Minimal $SU(2)_L \times U(1)$ Electroweak Theory", Rev. Mod. Phys. **53**, 211-252 (1981).

[37] G. Altarelli, "Precision Tests of the Electroweak Theory at the Z_0.", pp. 3-34 in Physics at LEP vol. 1, CERN 86-02, (1986).

[38] P. Baillon *et al.*, "Charge Asymmetry of $\mu^+\mu_-$", pp. 172-176 in Physics at LEP vol. 1, CERN 86-02, (1986).

[39] Bovet *et al.*, "Polarization and Polarimeters at LEP", pp. 58-62 in Physics at LEP vol. 1, CERN 86-02, (1986).

[40] H. Baer *et al.*, "New Particles", pp. 312-332 in Physics at LEP vol. 1, CERN 86-02, (1986).

[41] C. Rubbia and the UA1 Collaboration, "Physics Results of the UA1 Collaboration at the CERN Proton-Antiproton Collider", pp. 634-652 in The Sixth Quark, Proceedings of the Twelfth SLAC Summer Institute on Particle Physics, SLAC-281, (1984).

[42] W. Buchmüller *et al.*, "Toponium Physics at LEP", pp. 203-296 in Physics at LEP vol. 1, CERN 86-02, (1986).

[43] R. D. Pecci, "Physics at LEP", pp. 209-275 in 1986 CERN School of Physics, CERN 87-02, (1987).

[44] The LEP Design Group, Design Study of a 15 to 100 GeV e^+e^- Colliding Beam Machine (LEP), CERN/ISR-LEP 78-17, (1978).

[45] The L3 Collaboration, L3 Technical Proposal, (1983).

[46] The Lyon Group, Front End Electronics for the BGO Detector, CERN/EP/L3 395, (1985).

[47] J. A. Bakken *et al.*, Nucl. Instr. and Meth. **228**, 294 (1985).

[48] R. L. Sumner, BGO Electronics, L3-BGO memo, (1986).

[49] D. P. Stickland, Description of Task to Manufacture BGO Level-1 Circuits, Princeton memo, (1986).

[50] C. Fabjan, "Calorimetry in High-Energy Physics", p. 294 in Experimental Techniques in High Energy Physics (ed. T. Ferbel), Addison-Wesley, (1987).

[51] T. Schaad, Some Practical Hints Concerning the Geneva Table, L3-BGO memo, (1988).

[52] D. E. Plane, The West Area Beams, CERN/SPS 83-22, (1983).

[53] R. Clare, The 1987 L3 Test Beam Trigger, L3-BGO memo, (1987).

[54] A. Manarin and G. Vismara, The Delay Wire Chamber Description, LEP/BI-TA/Note 85-3, (1985).

[55] G. Sauvage, Report on the 1987 Beam Chambers, L3-BGO memo, (1988).

[56] K.L. Brown *et al.*, TRANSPORT: A computer program for designing charged particle beam transport stsyems, CERN/SPS 80-04, (1980).

[57] R. Brun *et al.*, GEANT3 User's Guide, CERN/DD/EE 84-1, (1986).

[58] C. Dionisi and M. Maire, Monte Carlo Simulation for the 1987 BGO Calibration, L3-BGO memo, (1987).

[59] M. Chemarin, and the Lyon Group, A Cosmic Ray Set-up to Study BGO Crystal Characteristics, CERN/EP/L3 394, (1985).

[60] J. A. Bakken *et al.*,"Performance of a Prototype BGO Calorimeter in an Electron Beam from 2 to 50 GeV", Nucl. Instr. and Meth **A254** 535-541 (1987).

[61] Y. Karyotakis *et al.*, HERMES, L3-BGO memo (1987).

[62] Precision Monolithics Incorporated, PMI 1986 Data Book, (1986).

[63] D. H. Perkins, Introduction to High Energy Physics, Addison-Wesley, Reading, Massachusetts, (1982).

[64] M. Roos and the Particle Data Group, "Review of Particle Properties", Phys. Lett. **111B** (1982).

[65] H. A. Bethe and J. Ashkin, "Passage of radiations through matter" p. 166 in Experimental Nuclear Physics Vol. 1 (ed. Segre), John Wiley, New York, p. 166 (1954).