Differential cross sections, charge production asymmetry, and spin-density matrix elements for $D^{*\pm}(2010)$ produced in 500 GeV/c π^- - nucleon interactions.

David Christopher Langs

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

November 2004

ProQuest Information and Learning Company
300 North Zeeb Road
PO Box 1346
Ann Arbor, MI 48106-1346

Abstract

Differential cross sections for the production of $D^{*\pm}(2010)$ mesons produced in 500 GeV/c π^--nucleon interactions from Fermilab experiment E791, are reported as functions of Feynman-x (x_F) and transverse momentum squared (p_T^2) and fitted to standard distributions. The $D^{*\pm}$ charge asymmetry and spin-density matrix elements as functions of these variables are also reported. The asymmetry curve is compared to PYTHIA Monte Carlo models. Investigation of the spin-density matrix elements shows no evidence of polarization. The average values of the spin alignment are $\langle\eta\rangle = 0.01 \pm 0.02$ and -0.01 ± 0.02 for leading and non-leading particles, respectively.

Acknowledgments

This research was supported by the Brazilian Conselho Nacional de Desenvolvimento Científico e Technológio, CONACyT (Mexico), FAPEMIG (Brazil), the Israeli Academy of Sciences and Humanities, the U.S. Department of Energy, the U.S.–Israel Binational Science Foundation and the U.S. National Science Foundation. Fermilab is operated by the Universities Research Association, Inc., under contract with the United States Department of Energy.

In addition, I would like to acknowledge personal financial support from Princeton University, the U.S. Department of Energy, The University of South Carolina, and my parents, Richard and Lenore Langs.

I would like to thank the entire E791 collaboration for providing the data and enormous help in doing this work. I would also like to acknowledge the staffs of Fermilab and all participating institutions for their assistance. I would especially like to thank members of my analysis committee: Milind Purohit, Jeff Appel, Don Summers and Austin Napier.

I'd also like to thank members of my Princeton thesis committee: Dan Marlow, Kirk McDonald, and Igor Klebanov. I'm also grateful to Laurel Lerner for her invaluable assistance in graduating.

I am very grateful to Carl Rosenfeld and Jeff Wilson of the University of South Carolina for making me feel like part of the family. I'd also like to thank Shmuel Nussinov for many illuminating discussions on my work.

And I am most grateful to the three people who more than any others taught me how to be a high energy physicist: Jim Wiener, Nader Copty, and especially my advisor, Milind Purohit, who has been remarkably patient during the long and difficult process of my education.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

In this dissertation, the hadroproduction of the $D^{*\pm}(2010)$ particle is investigated, using data from Fermilab experiment E791, a fixed target, high-statistics charm physics experiment. These results have been published [Ai02], and this document describes in detail the analysis leading to those results. This introductory chapter presents some of the basic definitions and concepts needed to understand the analysis, as well the theoretical framework which motivated the measurements made. Since this analysis predominantly involves hadron physics, a brief introduction to QCD is first presented.

1.1 The Standard Model

Modern elementary particle physics is best described by a theoretical framework known as the "Standard Model". It includes a description of the elementary particles and all the interactions which govern them, with the exception of gravity. These interactions give rise to forces which govern the dynamics of these elementary particles and in addition they can cause particles of one type to transform (or "decay") into

other particles. There are three interactions which are unified in the Standard Model: the weak interaction, the electromagnetic interaction, and the strong interaction.

Particles which participate in the strong interaction are known as *hadrons*. Familiar examples include the protons and neutrons which make up the nucleus of the atom. There are other elementary particles known as leptons which do not participate in the strong interactions. Electrons are the most familiar example. Hadrons and leptons together make up all the known matter in the universe.

According to the Standard Model, hadrons are not elementary particles, but are made up of smaller particles called "quarks". There are six known types or "flavors" of quarks, three with charge[1] $+\frac{2}{3}$: the up (u), charm (c), and top (t) quarks, and three with charge $-\frac{1}{3}$: down (d), strange (s), and bottom (b). Associated with each quark is an anti-quark of equal and opposite charge. Hadrons fall into two categories: baryons which consist of three quarks (or anti-baryons consisting of three anti-quarks), and mesons, which consist of one quark and one anti-quark. In addition to an electric charge, the quarks also carry an intrinsic angular momentum[2] (or "spin") of $\frac{1}{2}$. Thus baryons, which consist of an odd number of quarks, are fermions (particles with half odd integer spin), and mesons, consisting of an even number of quarks, are bosons (particles with integer spin).

When the quark model was introduced [Ge64], existing evidence was consistent with three quarks: up, down, and strange. Strong and electromagnetic interactions preserve quark flavors, but weak interactions are able to change them, allowing hadrons containing a strange quark to decay into hadrons containing only the lighter up and down quarks. Also, the force-carrying particles of the weak interaction come in both neutral and charged varieties. However, strangeness changing weak inter-

[1] All electric charges are expressed in units of e, the absolute value of the electron charge.
[2] Spins are expressed in units of $\hbar$ (Planck's constant divided by 2π).

actions were observed to occur almost exclusively via a charged current interaction. This was a puzzle, until a fourth quark, called "charm", was proposed [Gl70], whose existence introduced a term in the neutral current matrix element almost exactly canceling the one responsible for a strangeness changing interaction. This cancellation would be exact except for the fact that the up and charm quarks differ in mass. The charm quark was predicted to be much heavier than the strange quark explaining the fact that it hadn't been observed. When the J/ψ meson, consisting of a charm and anti-charm ($c\overline{c}$) quark was subsequently discovered [Au74a, Au74b] it provided strong confirmation of the quark model.

1.2 Quantum Chromodynamics (QCD)

The quark model was highly effective at describing the known mesons and baryons, but there was an outstanding puzzle. The baryon Δ^{++} has a charge of +2 and spin $\frac{3}{2}$. It was known to consist of three up quarks, each of spin $\frac{1}{2}$, which must be in a spin-symmetric state in order for the baryon to have a spin of $\frac{3}{2}$. So all three quarks must be in the same quantum state. This, however, is forbidden by the well-established Pauli exclusion principle which states that no two fermions can occupy the same quantum state. To get around this dilemma, a new quantum number, known as "color" was hypothesized [Gr64]. In addition to flavor, quarks possess one of three colors: red (r), green (g), or blue (b). Anti-quarks possess an "anti-color" (anti-red ($\overline{r}$), etc.) The three quarks in a baryon each have a different color, and thus the problem of the Δ^{++} was resolved. The existence of color was confirmed by observing the ratios of hadronic to non-hadronic decay rates, for instance the ratio of the rates of reactions for $e^+e^- \rightarrow$ hadrons vs. $e^+e^- \rightarrow \mu^+\mu^-$ is proportional to the number of colors (N_c), and all measurements are consistent with $N_c = 3$.

Color also turned out to be the key to constructing a quantum field theory for strong interactions. In analogy with quantum electrodynamics (QED) which describes electromagnetic interactions by means of the exchange of virtual field quanta (photons), the theory was named quantum chromodynamics (QCD).

QCD has some similarities and some differences with QED. Unlike QED which has one charge (a negative charge is simply the "anti-charge" of a positive charge), QCD has three charges: the colors, introduced to get around the exclusion principle problems. The field quanta are known as "gluons", and they themselves carry a color/anti-color charge pair, for instance $r\overline{b}$. Because there are three colors (and three anti-colors) we might expect 9 gluons, but the color singlet state $r\overline{r} + g\overline{g} + b\overline{b}$ is excluded, yielding eight. The fact that gluons themselves carry color leads to two important differences between QCD and QED. The photons of QED do not change the charge of the particles with which they interact, whereas gluons, which carry color and anti-color, can change the colors of the quarks with which they interact. Also, because gluons carry a color charge, the gluons themselves can absorb or emit other gluons. This gives rise to an inverse charge screening effect. The further one moves away from a bare color charge, the stronger it appears to be. The result is that hadrons must be colorless (as in mesons) or "white" (as in baryons: rgb), in order to avoid infinite field strengths. This is why free quarks are never observed: they would always carry a bare color charge.

This suggests that the coupling strength for the strong interaction increases as the distance scale increases. When the coupling is mediated by a virtual particle, the uncertainty principle dictates that the distance scale is inversely proportional to the

momentum transferred. The QCD coupling strength is, in fact, given by:

$$\alpha_s(Q^2) = \frac{12\pi}{(11N_c - 2N_f)\ln(Q^2/\Lambda^2)} \tag{1.1}$$

where Q^2 is absolute value of the square of the 4-momentum transferred in the process, $N_c = 3$ is the number of colors, and N_f is the number of quark flavors accessible at the energy scale involved. (At energy scales sufficient to produce charm, but not heavier quarks, $N_f = 4$.) Λ is a constant known as the "QCD mass scale parameter". Coupling constants in general have an arbitrary parameter which must be determined experimentally, and it is an artifact of the renormalization process that it has units of mass in Equation (1.1). Λ is expected to be roughly the mass of a hadron, and is found experimentally to be 216 ± 25 MeV/c [Ha02, p. 95]. For $Q^2 \gg \Lambda^2$, α_s is small, and perturbation theory can be used in QCD calculations, but for $Q^2 \approx \Lambda^2$, $\alpha_s \sim 1$ and perturbation theory is inappropriate. Physically this means that at high energies, quarks behave as free particles, but at low energies they are very tightly confined in hadrons. Note also that gluons can create virtual $q\overline{q}$ pairs, and so a hadron consists of three types of constituents: the permanent quarks that make up the hadron, known as *valence quarks*, the gluons, and virtual quarks and anti-quarks, known as *sea quarks*.

1.3 Hadroproduction of Charm

The production of heavy quarks in high energy hadronic interactions provides a useful testing ground for QCD. Because the mass of the charm quark (believed to be between 1.0 and 1.4 GeV/c^2 [Ha02, p. 29]) is greater than $\Lambda_{\rm QCD}$, perturbation theory should be applicable to charm quark production. The analysis described in this dissertation involves the production of $D^{*\pm}(2010)$ mesons resulting from the interaction

of 500 GeV/c π^- mesons ($\overline{u}d$) with fixed nuclear targets: protons (uud) and neutrons (udd).

Recall that mesons are bound states of a quark and anti-quark. As with other bound systems, they can exist in a ground state or one of several excited states. Ground state mesons are the lightest mesons with a given quark content. Ground state mesons in which the quark and anti-quark are of different flavor can decay only via the weak interaction, the only interaction which doesn't preserve quark flavor. Excited mesons can decay via the strong interaction to lower lying mesons. This gives rise to their extremely short lifetimes. The $D^{*+}(2010)$ meson is an excited meson which consists of a charm and anti-down quark ($c\overline{d}$). The "2010" refers to its mass of 2010 MeV/c^2. The D^{*-} ($\overline{c}d$) is its anti-particle. Unlike the ground state D^+ (mass 1869 MeV/c^2) which has spin 0, the $D^{*\pm}$ is a vector meson (spin 1).

There are several advantages to studying production of D^* particles over the low-lying D mesons. First, their spin causes their production to be favored three-fold over the pseudoscalar mesons at high center of mass energies. Second, they are more likely to have been produced directly in the initial interaction, rather than arising from the decay of an excited meson. Third, the fact that they have non-zero spin allows an investigation of their polarization through an angular distribution study. This last point will be explained in more detail later.

Hadroproduction of D^*'s has been studied experimentally before [Ad99, Al94b, Ba91, Ba88, Ag86], but the E791 experiment, described extensively in subsequent chapters, provides the largest sample yet of hadroproduced D^* mesons, thus enabling measurements with much smaller error bars than previous results.

1.3.1 Measured Quantities

The measurable properties of particles fall into two broad categories: the *intrinsic* properties, common to all particles of the same type, such as intrinsic spin, lifetime, mass, electric charge, etc., and the *extrinsic* properties such as momentum, spin direction, etc. If the particle and anti-particle are considered together, the ratio of the numbers of particles of one charge to the opposite in a sample can be considered an extrinsic property as well. If one wishes to study the physics of the interaction which produced a particle, it is the extrinsic properties which are important. In this analysis these three extrinsic properties: momentum, charge asymmetry, and spin direction, of the D^* particles produced in E791 are studied in Chapters 7, 8, and 9, respectively. First, these quantities will be more precisely defined, then physical models which predict these quantities will be discussed.

Differential Cross Sections

When a beam of particles is directed on a target, the probability of an interaction per unit time is clearly directly proportional to the number of beam particles striking the target per unit time (which is equal to the flux ϕ times the cross sectional area of the beam A), and also to the density of the target material, expressed as number of particles per unit volume n_A, as well as its thickness L, assuming it is thin enough that the beam flux remains approximately constant as the beam traverses the target. In other words:

$$W = \sigma \phi A n_A L \tag{1.2}$$

where W represents the number of interactions per unit time. Clearly, all of the underlying physics describing the interaction is contained in the proportionality constant σ, which has units of area. It is called a *cross section* in analogy with classical

ballistics experiments where the interaction probability depends on the cross sectional area of the target. One can measure a total cross section, or instead study how the cross section depends on other quantities, such as the distribution in phase space of decay products. Measurements of the latter type are called *differential* cross sections, and are denoted $d\sigma/dX$ where X represents some physical quantity or quantities on which the interaction probability depends.

Momentum Distributions

As stated above, the first of the three measurements performed in this analysis is a study of the momentum spectrum of the produced D^*'s. Measurement of the momentum in an arbitrary frame and coordinate system clearly would make comparisons with other experiments difficult, so a standard system has been employed. The beam direction gives a physically meaningful reference direction, and so the momentum is first resolved into components parallel and perpendicular to the beam.

The perpendicular component has both a magnitude and direction, an azimuthal angle measured about the beam axis. The symmetry of the system suggests no interesting dependence on this azimuthal angle, so the square magnitude of the transverse momentum is measured. This magnitude is clearly invariant under boosts along the beam direction, so the value in the lab frame and center of mass frame of the interacting particles is the same. The distribution is expressed as a differential cross section, denoted $d\sigma/dp_T^2$, where p_T^2 denotes square transverse momentum.

The component of $\vec{p}$ parallel to the beam is slightly more complicated. Firstly, it does not have the same value in the lab and center of mass frame. Since the lab frame is arbitrary as far as the underlying physics is concerned, the measurement is boosted to the center of mass frame of the interacting particles. Also, unlike p_T^2, the parallel momentum, even in the center of mass frame, is heavily dependent on the

beam particle momentum. This would make comparisons with distributions obtained from other experiments with different beam energies difficult. For this reason, the momentum is divided by a scale factor, the maximum kinematically allowed parallel momentum, to make the distribution approximately independent of beam energy. This scaled quantity is called "Feynman-x", denoted by x_F:

$$x_F = \frac{p^*_{\shortparallel}}{p^*_{\shortparallel \max}} \tag{1.3}$$

Here, the asterisks indicate the quantities are evaluated in the center of mass frame. Clearly, $-1 \leq x_F \leq 1$. The measured differential cross section is denoted by $d\sigma/dx_F$.

Charge Asymmetry

Asymmetries in the rate of production of a particle and its antiparticle can arise from the quark composition of the interacting particles. When one of the two produced particles has a valence quark with the same flavor as a valence quark in the beam particle, the produced particle is called a *leading particle*.

The bare asymmetry, $\sigma_L - \sigma_{NL}$ is obviously going to be proportional to the total number of particles produced in the experiment, so it isn't physically meaningful. (Here L and NL refer to the leading and non-leading particle respectively.) It is thus customary to divide the asymmetry by this total, to yield the asymmetry parameter:

$$A = \frac{\sigma_L - \sigma_{NL}}{\sigma_L + \sigma_{NL}} \tag{1.4}$$

Clearly, $-1 \leq A \leq 1$, with $A = 0$ when there is no asymmetry.

Angular Momentum

Although the magnitude of the angular momentum is an intrinsic property of a particle (1 for a D^*), the direction of the angular momentum vector can depend on the production physics. Spins might be preferentially aligned with the direction of the interacting particles in the center of mass frame, for instance. If the spin directions are not distributed evenly on the unit sphere, the sample is said to have a non-zero *polarization*.

The polarization can be characterized numerically by means of a *spin density matrix*. The elements of the spin density matrix can be determined by examining the angular distributions of the decay products. Both the spin density matrix, and its effect on angular distributions are discussed in detail in Chapter 9.

1.3.2 The Parton Model

Even before the QCD theory was formulated, evidence existed for a hadron substructure. So-called *deep inelastic scattering* experiments, in which leptons were scattered off nuclear targets, provided evidence that hadrons were composed of point-like constituents. The *parton model* was developed to explain these results.

The original parton model describes nucleons as being composed of point-like subparticles, called *partons*, each of which carries a fraction x of the nucleon's 4-momentum. Furthermore, these fractional momenta were not fixed, but rather governed by *parton distribution functions* (PDFs), $F_i(x)$, where $F_i(x)$ is the probability that a parton of type i carries a fraction x of the hadron's 4-momentum. The subscript i refers to the fact that different parton constituents could have different interaction cross-sections or angular distributions. According to the original parton model, the functions $F(x)$ were independent of the energy scale (the square of the 4-momentum

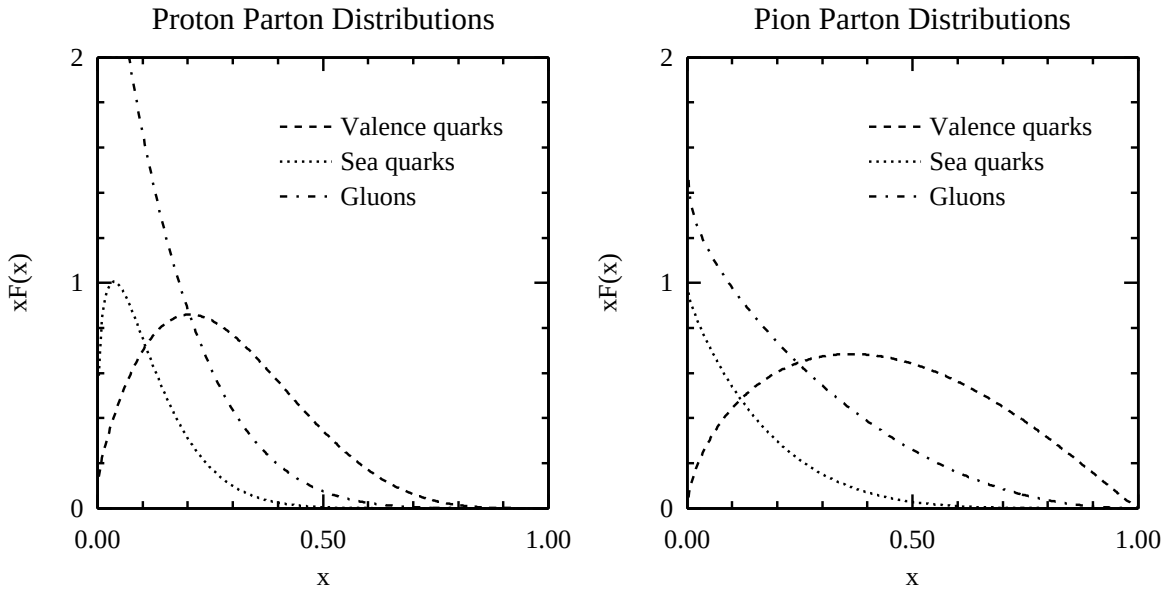

Figure 1.1: Parton distribution functions for protons and pions obtained using PDFLIB package [Pl95]. They were computed at $Q^2 = 5$ (GeV/c)^2 and $\Lambda = 190$ MeV/c. The proton PDFs used the HMRS B data set [Ha90], and the pion used SMRS 2 [Su92b].

transfer, q^2) used to probe the hadron. This is known as the *Bjorken scaling hypothesis* [Bj67]. Experimental results were approximately consistent with this assumption.

QCD accommodated the parton model with a few modifications. The partons were identified as the valence quarks, sea quarks, and gluons which make up hadrons. The fact that lepton-nucleon scattering experiments could account for only about half of the parton constituents of the nucleon provided good supporting evidence for the existence of gluons, which carry no electric charge. The gluons, however, are capable of producing virtual $q\overline{q}$ pairs, and the number of these would increase as the time-scale of the interaction decreases according to the uncertainty principle, $\Delta E \Delta t \geq \hbar$. Since the time-scale is inverse to the energy transferred in the interaction, the number of partons goes up with q^2. This results in a violation of the scaling

hypothesis, with PDFs at high q^2 having more weight at lower x than those at lower q^2. Fortunately, the evolution of PDFs can be calculated exactly by means of the *Altarelli-Parisi* equations [Al77], and so measurements at one value of q^2 are sufficient. The original parton model also did not include any momentum component of the partons transverse to the momentum of the hadron. But Fermi motion, as well as the emission and absorption processes internal to the hadron, result in some transverse momentum. An intrinsic momentum k_T is thus added to the partons, usually according to a Gaussian distribution. Typical widths are a few hundred MeV/c. Sample PDFs for both the proton and pion are shown in Figure 1.1.

1.3.3 Charm Production Cross Section

Because the strong interaction conserves flavor, charm quarks must be produced together with an anti-charm quark. Of course, these quarks will eventually form hadrons which are observed, but the time scale for hadronization is much longer than that of heavy quark production, so it is hypothesized that the two processes can be considered separately. This is the so-called *factorization* hypothesis.

Figure 1.2 depicts our model for the process of heavy quark production in hadron-hadron interactions. The differential cross section for the process is given by:

$$\frac{d\sigma_c}{d^3k} = \sum_{i,j} \int dx_a \, dx_b \, \left[\frac{d\hat{\sigma}_{ij}(x_a P_a, x_b P_b, k, m_c, q^2)}{d^3k} \right] \, F_i^a(x_a, q^2) \, F_j^b(x_b, q^2) \qquad (1.5)$$

Here, the indicies i and j run over the parton constituents of hadrons a and b respectively. k and m_c represent the charm quark momentum and mass, respectively. q^2 is the square of the 4-momentum transferred in the process, typically taken to be the square of the charm quark mass in charm quark production. F_i^a and F_j^b are the parton distribution functions for hadrons a and b and partons i and j, respec-

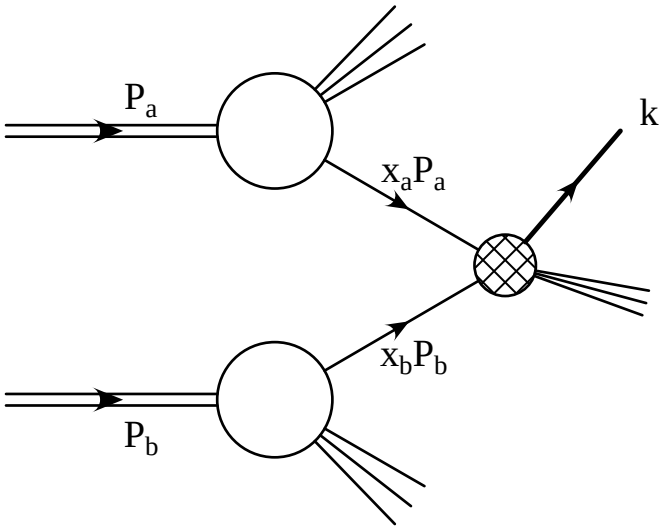

Figure 1.2: The parton picture of heavy quark production in hadron-hadron interactions. Here **k** represents the momentum of the heavy quark produced in the interaction of two hadrons with momenta $\mathbf{P}_a$ and $\mathbf{P}_b$.

tively. The expression in square brackets is the parton-parton cross section for charm production, at a given momentum, k.

As stated earlier, the fact that the charm quark mass is greater than $\Lambda_{\rm QCD}$ suggests that the parton-parton cross section is calculable using perturbative techniques. Because $m_c \not\gg \Lambda_{\rm QCD}$, it isn't clear how accurate these predictions will be. These calculations have been done up to leading order (LO) [Jo78] and next-to-leading order (NLO) [Na89]. There are four diagrams which contribute at leading order to heavy quark production: three gluon-gluon fusion diagrams and one quark–anti-quark annihilation diagram. They are depicted in Figure 1.3. Because they each contain two vertices, they contribute to the amplitude at order α_s^2. Next-to-leading order calculations involve terms up to order α_s^3. Many diagrams contribute at this level. Four gluon-gluon fusion diagrams are shown in Figure 1.4. The upper two are α_s^3 diagrams, and the lower two are α_s^4, but they interfere with leading order diagrams

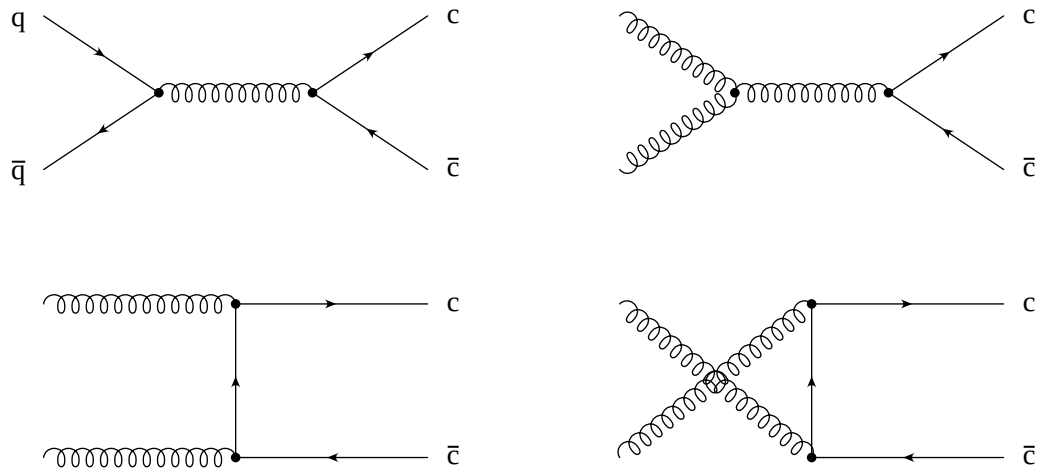

Figure 1.3: The four leading order (α_s^2) diagrams for charm quark production in parton-parton interactions.

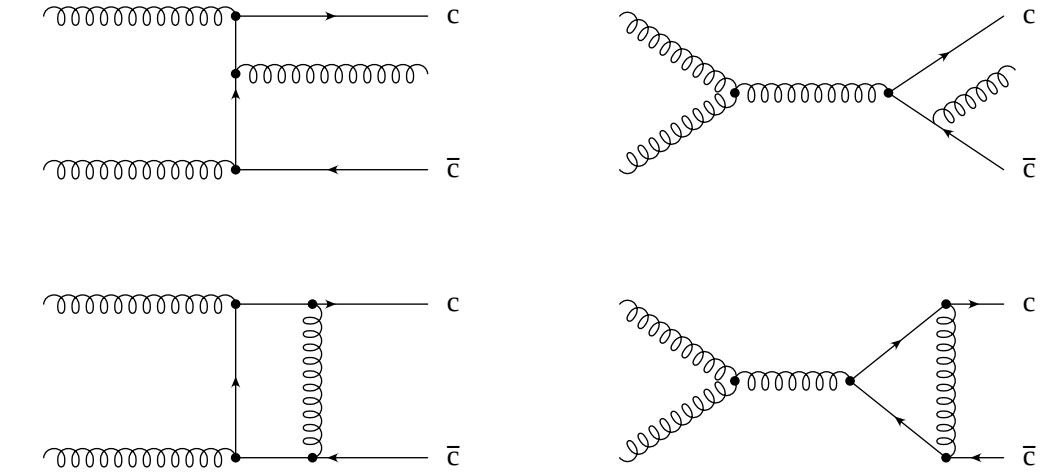

Figure 1.4: Four examples of next-to-leading order diagrams for charm quark production in parton-parton interactions. The upper two are α_s^3 diagrams; the lower two are α_s^4 which interfere with leading order diagrams at the same order as α_s^3 diagrams.

and the interference terms contribute to the probabilities at the same order as α_s^3 diagrams.

1.3.4 Hadronization

The final stage in charm hadroproduction is the formation of hadrons from the quarks, anti-quarks, and gluons which result from the hard parton-parton scattering process. There is generally sufficient energy in the process to result in radiation of soft gluons which create several light quark/anti-quark pairs. Also, the original interacting particles can contribute additional partons to the process. Eventually, all these partons must form colorless hadrons which can subsequently decay or propagate through the detector. The process by which these partons are transformed into hadrons is called *hadronization* or *fragmentation*.

In the context of the factorization theorem, the hadronization process can be represented by *fragmentation functions*, $D_q^h(z)$, which represents the probability that a quark q will become part of a hadron h which carries a fraction z of the quark's momentum. The differential cross section for the production of a charm hadron h with momentum p can be expressed in terms of the charm quark production cross section derived in the previous section by convoluting it with the appropriate fragmentation function:

$$d\sigma_h(p) \propto \int_0^1 d\sigma_c(p/z)\ D_c^h(z)\ dz \tag{1.6}$$

The fragmentation functions can't be computed directly from first principles as the energy scale for hadronization is well below that amenable to perturbative QCD. Fragmentation functions have thus been proposed in order to agree with experimental results.

For heavy quark production, the Peterson function [Pe83] has been quite popular:

$$D(z) \propto \frac{1}{z\left(a - \frac{1}{z} - \frac{\epsilon_Q}{1-z}\right)^2}, \tag{1.7}$$

where ϵ_Q is a free parameter scaling with the heavy quark mass as $\epsilon_q \propto 1/m_Q^2$. The Peterson function describes data from e^+e^- collisions quite well where a heavy quark and its anti-quark initially have equal and opposite momenta. In hadroproduction, however, the heavy quarks frequently combine with remnants of the interacting hadrons and the Peterson function doesn't model this adequately.

Another popular fragmentation function is derived from the Lund String Model. The lines of color force in a color dipole field between a quark and oppositely colored anti-quark or diquark tend to attract one another as a result of the triple gluon vertex of QCD. This gives rise to narrow "flux tubes" between them. The width of the tube is assumed to be independent of length (roughly 1 fm), and so the energy in the field is proportional to distance between the quarks. This is analogous to a stretching string with a spring constant of roughly 1 GeV/fm. In the string model of fragmentation, these color strings join quarks, anti-quarks, and diquarks which together would form a color singlet state. These partons have a momentum, and as they move they stretch the string, losing kinetic energy and adding to the energy of the string until the string has enough energy to create a real $q\overline{q}$ pair. The string thus breaks, with the new $q\overline{q}$ at each end. This process continues until there is not enough energy left to create new $q\overline{q}$ pairs, and the resulting joined pairs form hadrons. This process is expected to be symmetric between the quark and anti-quark ends of the string, and this gives rise to the Lund symmetric fragmentation function [An83]:

$$D(z) \propto \frac{(1-z)^a}{z} \exp\left(\frac{-b(m_h^2 + p_T^2)}{z}\right) \tag{1.8}$$

where m_h is the hadron mass, and p_T^2 its square transverse momentum.

When one end of the string contains a heavy quark, however, the situation is not symmetric and the Lund function is modified [Bo81]:

$$D(z) \propto \frac{(1-z)^a}{z^{1+r_Q b m_Q^2}} \exp\left(\frac{-b(m_h^2+p_T^2)}{z}\right) \tag{1.9}$$

where r_Q is a constant which may depend on the flavor of the heavy quark.

1.3.5 Interpreting Data

Previous experiments have fit x_F and p_T^2 distributions with the functional forms:

$$\frac{d\sigma}{dx_F} \propto (1-x_F)^n \qquad \text{and} \qquad \frac{d\sigma}{dp_T^2} \propto \exp(-Bp_T^2) \tag{1.10}$$

These arose from early models with the x_F form valid only at large x_F and the p_T^2 form at low p_T^2. Neither fit data from high statistics experiments very well, but fits to these forms can facilitate comparison with other experiments.

The functional form:

$$\frac{d\sigma}{dp_T^2} = \left(\frac{C}{bm_c^2+p_T^2}\right)^\beta \tag{1.11}$$

has been shown to fit NLO QCD Monte Carlo predictions very well [Fr94]. If this curve fits our p_T^2 distribution, it would lend support to the validity of using perturbative QCD to model charm quark production.

The software package PYTHIA [Sj94] simulates high energy physics interactions. It uses only leading order QCD to simulate the hard processes, but it also employs a parton shower model to simulate the radiative corrections which occur at higher order. The Lund string model using the Bowler fragmentation function (Eq. (1.9)) is used to model hadronization.

The $d\sigma/dx_F$ and $d\sigma/dp_T^2$ distributions are fit with the above mentioned functional forms and compared to PYTHIA models.

1.3.6 Asymmetry Models

A large number of previous studies of hadroproduction of charm show strong evidence for enhanced production of leading particles at high x_F. (See Refs. [Al94a, Ai96] and references therein.) Many models have been developed to explain this so-called "leading particle effect".

Leading order QCD predicts no asymmetry in the production of c and $\overline{c}$ at any momentum. Next-to-leading order QCD does predict a small enhancement of the leading particle at high x_F arising from subprocesses $qg \rightarrow c(\overline{c}) + X$ and $q\overline{q} \rightarrow c(\overline{c}) + X$ [Na89]. The predicted asymmetry, though, is at least an order of magnitude too low to agree with experiment [Ai96].

Another proposed model is the *intrinsic charm model* [Vo95], which depends on the intrinsic charm content of the beam particle. In particular, the beam π^- is assumed to fluctuate into a $|\overline{u}d\overline{c}c\rangle$ Fock state which is broken in the interaction. Since the charm quarks are close in phase space to the pion's valence quarks, they are more likely to recombine with them, forming leading particles. This model does not appear to be in good agreement with experimental results [Ai96].

The Lund string fragmentation model also predicts a strong leading particle effect. This arises from a phenomenon known as "beam dragging". The charm and anti-charm quarks are produced perturbatively with roughly the same momentum distributions. However, in forming color strings, one of these quarks can form a color string with a beam remnant which has a high forward momentum in the center of mass frame. This beam parton pulls the heavy quark forward thus giving it a large

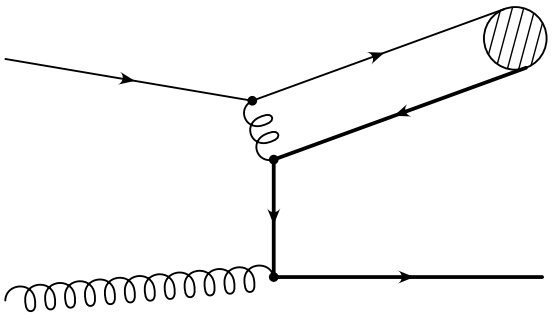

Figure 1.5: An example of heavy quark recombination into a meson. The heavy line represents the heavy quark and the shaded circle represents hadronization into a meson containing the heavy quark.

forward momentum. As a result the asymmetry between leading and non-leading particles increases steeply at large x_F. The default PYTHIA program actually predicts a larger leading particle effect than is seen in data [Ai96]. Parameters were tuned to bring about better agreement in that analysis, but only at the expense of selecting k_T and charm quark masses that might be unphysically large.

Another model is the *recombination model* [Da77]. In this model, some of the $c\overline{c}$ pairs are produced by a short-distance process within the beam hadron. They are thus close in phase space to the beam remnants and may preferentially recombine with them. This is related to the intrinsic charm model in that the charm quarks exist in the beam particle, but instead of existing there in an initial Fock state, they are produced (possibly perturbatively) by interaction with a target parton. (See Figure 1.5.)

In order to test these models, we plot the asymmetry parameter (Eq. (1.4)) for $D^{*\pm}$ as functions of both x_F and p_T^2. Our asymmetry results are compared both to tuned and untuned PYTHIA models, which is a further test of the Lund string

fragmentation model. The x_F results are also compared to a model which incorporates both fragmentation and recombination.

1.3.7 Polarization

Perturbative QCD does not predict any polarization of heavy quark hadrons produced in hadron-hadron interactions. The discovery of momentum dependent polarization in Λ hyperons produced in proton-nucleon interactions thus presented a puzzle. In this experiment [Bu76], it was found that the Λ polarization increased monotonically with transverse momentum. Since then, numerous experiments involving hyperon hadroproduction have observed significant polarization effects. (See Ref. [Tr03], and references therein.)

As there have been no previous attempted measurements of polarization of hadroproduced vector heavy quark mesons, there are few, if any, models specific to that case. Existing models are specific to hyperon hadroproduction or D^* mesons produced in e^+e^- collisions. Both types of models usually involve either the fragmentation process, or a production model other than the simple factorization scheme discussed earlier. An example of a hyperon polarization model is the recombination model for production, together with a Thomas precession model for polarization [An97]. For D^*'s produced in e^+e^-, the Suzuki model [Su86] predicts a polarization arising from processes where hadronization occurs via the emission of a single gluon from the heavy quark. Neither of these models appears to be applicable to hadroproduced charm mesons, however.

Because of the heavy mass of the charm quark, one might expect its spin to be retained in the produced hadron. It has been argued that despite its heavy mass,

fragmentation effects would dominate and hide any polarization arising from the quark production process [Fa94].

In any case, an attempt to measure such a polarization would fill an important gap in experimental data, and may help to clear up some of the outstanding mysteries in this sector.

Chapter 2

Beam Line and Target

Experiment E791, "Continued Studies of Heavy Flavors", was performed at the Fermi National Accelerator Laboratory (FNAL, also known as "Fermilab") in Batavia, Illinois. Its purpose was to collect a high statistics sample of charm particle events produced by a beam of negatively charged pions with an average energy of about 500 GeV striking a fixed target. The detector was located at Fermilab's Tagged Photon Lab (TPL) in the Proton-East experimental area, and recorded data during the 1991–2 fixed target run.

2.1 Accelerator

This section describes the Fermilab accelerator as it was configured during the E791 run, which predated the addition of the Main Injector.

To produce the beam, a primary beam of 800 GeV protons was produced in the Fermilab accelerator. This beam was created in several stages.

Firstly, electrons were added to hydrogen atoms producing negative ions. A

Cockcroft-Walton power supply produced a potential of 750 kV, through which the ions were accelerated, producing a beam of ions with an energy of 0.75 MeV.

The ions then entered a linear accelerator (Linac) approximately 150 m long. The Linac contained very intense electric fields which alternated in direction. The ions moved through drift tubes which shielded them from the field where its direction opposed their motion, and they were accelerated in the gaps between the tubes where the direction of the field increased their energy. Here the ions were accelerated to approximately 200 MeV and then passed through a carbon foil which stripped off the electrons leaving only protons.

The protons entered the booster, which was a synchrotron 150 m in diameter. Here, magnets kept the protons moving in a circular path as RF cavities accelerated them. The protons traveled around the booster 20 000 times where their energy was raised to 8 GeV.

The next stage was the Main Ring, another proton synchrotron, four miles in circumference, and containing 1000 copper coiled magnets. The main ring accelerated the protons to 150 GeV.

Housed in the same tunnel as the Main Ring is the Tevatron, another synchrotron, containing 1000 superconducting magnets. When running in fixed target mode, the Tevatron accelerated the protons to an energy of 800 GeV. The protons were extracted in pulses lasting 23 s containing approximately 2×10^{13} protons.

2.2 Fixed Target Beamline

When the accelerator operated in fixed-target mode, protons extracted from the Tevatron were sent to a switchyard where they were split into three beams which were sent to the meson area, the neutrino area, and the proton area. The proton area beam

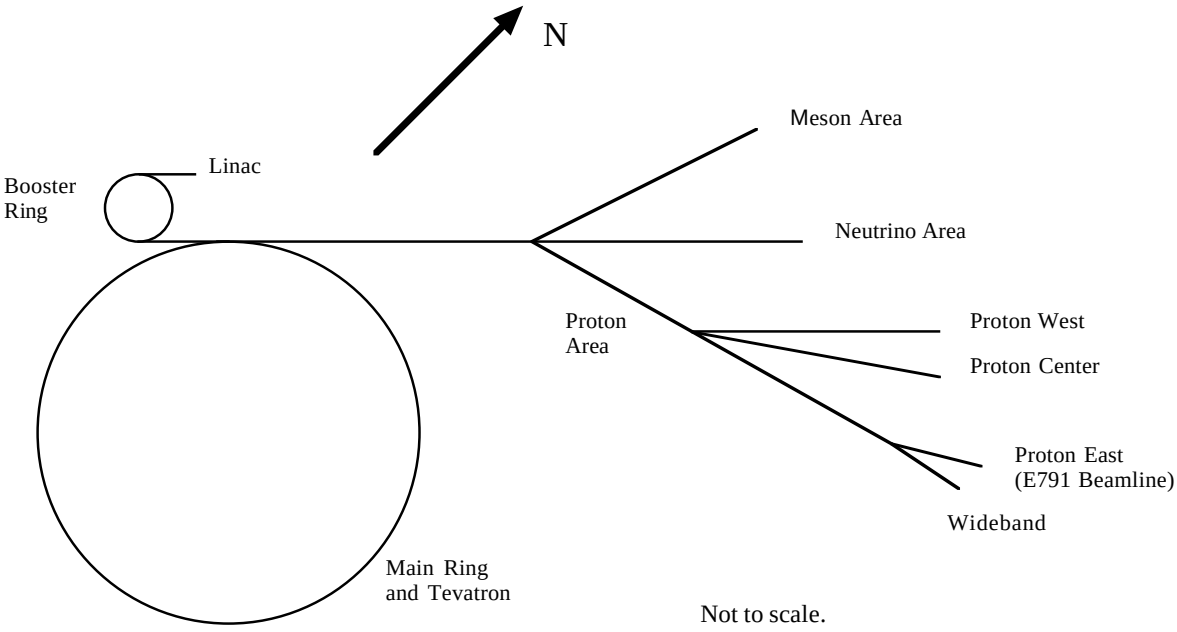

Figure 2.1: Fermilab beam lines (c1991).

was then split into Proton West, Proton Center, Proton East, and Wideband beams. The Proton East beam was directed onto a 30 cm beryllium target. Dipole magnets and collimators were used to select secondary π^- particles with an energy of 500 GeV which were sent to the E791 experimental target. (See Fig. 2.1).

The choice of beam particle was governed by the desire to maximize the size of the charm sample [Su92a]. A photon beam would have produced more charm particles per hadronic interaction[1], but it wouldn't have been possible to get enough photons to generate the amount of charm possible with a hadron beam, owing to both the greater cross section for hadronic interactions and greater luminosity possible with a charged hadron beam. Pions were chosen rather than protons because they have a stiffer gluon distribution [Mc83] (See also Figure 1.1.), and it is believed that charm is produced predominantly by gluon-gluon fusion [Jo78]. Finally, negative pions were chosen as proton contamination would have been substantial in a π^+ beam.

The complete cycle time was 57 s. 34 s were required to ramp up the synchrotron, and the spill lasted 23 s. During each spill, 2×10^{13} protons left the Tevatron, and 2×10^{12} protons were allocated to Proton East. This resulted in roughly 2×10^7 π^-'s being sent to the E791 target.

2.3 Target

Several considerations went into the target design. In fixed target charm physics experiments it is crucial to be able to find clean secondary vertices, and these must be distinguished from vertices resulting from multiple scattering in the target material. This consideration would tend to favor a thin target. One also would like to increase

[1]In the photon-gluon fusion model, all the beam particle's energy goes into the production of a $q\overline{q}$ pair, thus favoring heavier quark production at higher beam energies. In the gluon-gluon fusion model of hadroproduction, only a part of the beam particle's energy (governed by the gluon distribution function) goes into the $q\overline{q}$ production, thus favoring the production of lighter quarks.

Table 2.1: Summary of E791 target information

Foil Number	Material	Z Pos. (cm)	Thickness (cm)	Diameter (cm)	Interaction Length (%)
1	Pt	-8.191	0.052	1.61	0.584
2	C	-6.690	0.157	1.37	0.589
3	C	-5.154	0.157	1.38	0.586
4	C	-3.594	0.153	1.37	0.582
5	C	-2.060	0.158	1.36	0.582

the number of interactions, and this requires more target material. To accomplish these two goals, E791 used five separate target foils spaced roughly 1.5 cm apart.

The thinness of the foils necessitated a fairly dense material in order that a useful fraction of the incident pions interact (roughly 2%). The first target foil was made of platinum ($\rho = 21.4\text{g/cm}^3$), and the remaining four were made of carbon (industrial diamond, $\rho = 3.2\text{g/cm}^3$). The use of two different materials allowed one to study the dependence of charm production cross section on the number of nucleons in a target nucleus. ($A = 195$ for platinum and $A = 12$ for carbon.)

The foils were roughly 1 cm in diameter and were mounted in lexan. (See Table 2.1). The thinness and spacing of the targets allowed one to reconstruct secondary vertices between the foils. One can then require a minimum significance of separation between a secondary vertex and the nearest target foil to reduce background due to multiple scattering. Also, the thinness of the target material allowed one to obtain the z-position of the primary vertices to a resolution of between 250 and 450 μm. (See Fig. 2.2).

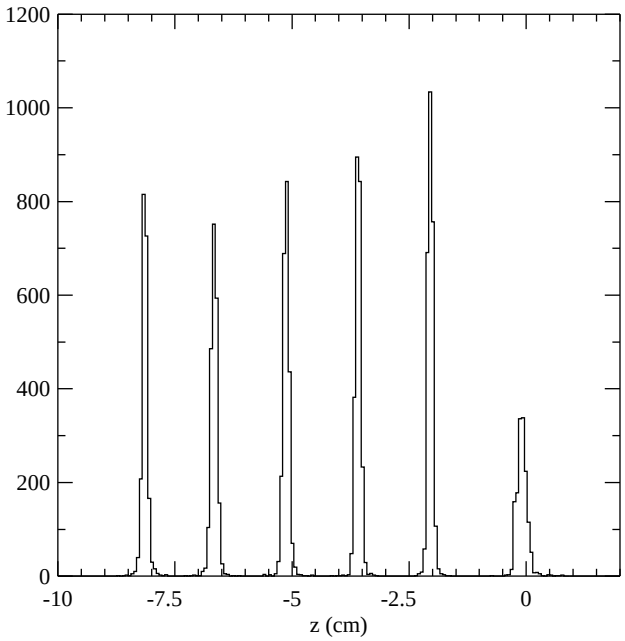

Figure 2.2: Z position of reconstructed primary vertices. First five peaks are at location of target foils. The sixth one is at the interaction counter.

Chapter 3

Spectrometer

3.1 History

Fermilab experiment E791 was housed in the Tagged Photon Lab. The spectrometer was originally built for E516 [Sl85], a charm photoproduction experiment which began in 1979 and ran in the early 1980's. For the following experiment, the highly successful E691 [Ra88], nine silicon microstrip detectors were used for very high vertex resolution. This produced 10,000 fully reconstructed charm events, more than a fifty-fold increase over E516. The next experiment, E769 [Al92], used a hadron beam consisting of pions, kaons, and protons rather than a photon beam.

Several improvements [Su92a] were made in upgrading to E791. The number of silicon microstrip planes was increased to 23. A second muon wall was added, and the data acquisition system was greatly enhanced to handle the high data taking rate necessitated by the open trigger design and high luminosity produced by a hadron beam. Also, for E791 only a pion beam was used.

E791 took data during the 1991 fixed target run (July 1991 to January 1992).

Despite the fact that hadron beams were used since E769, the name "Tagged Photon Lab" remained.

3.2 Spectrometer Introduction

The E791 spectrometer[1](Fig. 3.1) was designed for charm physics. Crucial to obtaining a charm signal is the ability to observe a clear separation between the primary vertex (where a charm particle may be produced) and a secondary vertex (where the charm particle decayed.) Because of the short lifetime of charm particles (on the order of a picosecond or less) good vertex resolution was important. Consider for example a D^+ with a mass of 1.869 GeV/c^2, and lifetime 1.057 ps. A typical D^+ momentum in E791 is 50 GeV/c. This gives a Lorentz factor,

$$\gamma \approx p/mc \approx 27. \tag{3.1}$$

This yields a mean flight length

$$l \approx \gamma c\tau \approx 9\text{mm}. \tag{3.2}$$

Thus, vertex resolutions would have to be on the order of tens of microns to obtain clean charm samples. This is especially true in a hadroproduction experiment like E791 where a large number of particles (typically 10 per event) were produced at the primary vertex. Thus the silicon microstrip detectors were a critical part of the E791 spectrometer. 23 planes were used: 6 upstream of the target and 17 downstream.

In order to measure the charge and momentum of the decay tracks, two magnets

[1]Much of the information in this chapter was obtained from previous E791 dissertations[Wi94, Wi96].

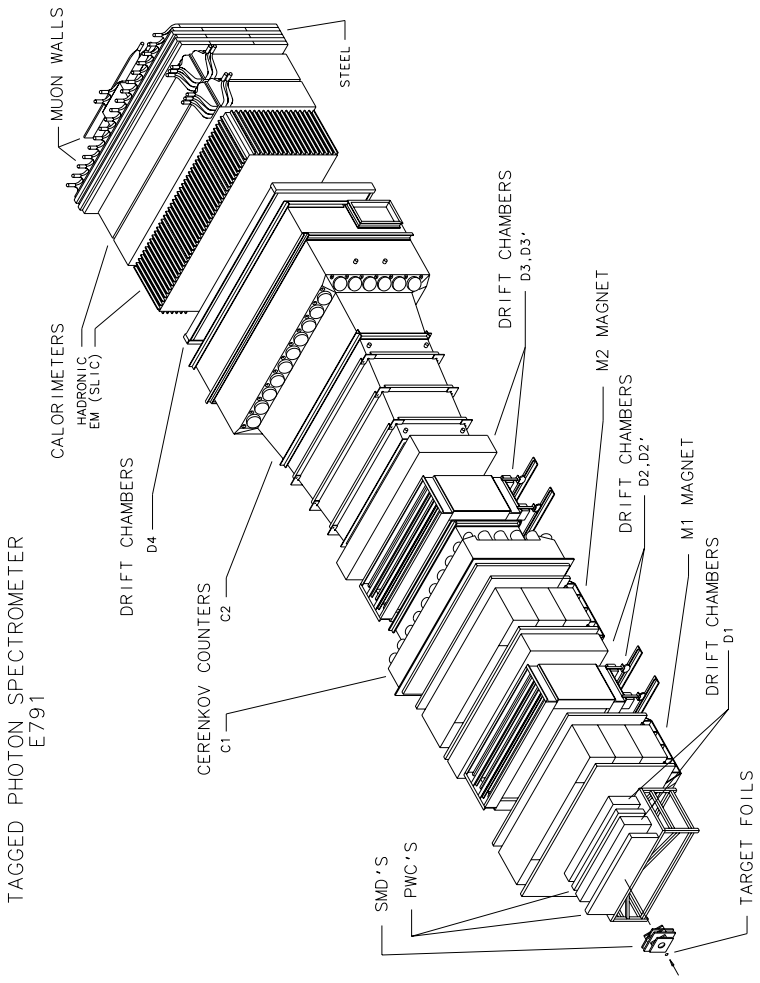

Figure 3.1: E791 Spectrometer. (Source: Fermilab)

were used. Drift chambers were placed upstream, downstream and between the magnets with proportional wire chambers also placed upstream of the first magnet. Two Čerenkov counters provided particle identification. To assist in particle identification and triggering, E791 had two calorimeters, a segmented liquid ionization calorimeter (SLIC) for identification of photons, electrons, and positrons, and a hadrometer for detection of hadrons and muons. Finally two walls of scintillating paddles were used for muon identification.

3.3 E791 Coordinate System

Much of the discussion of the geometry of the detectors will be done in terms of the E791 coordinate system which is discussed here. A right-handed Cartesian system is used with the positive Z axis pointing in the direction of the beam. The Y axis is vertical with the positive axis pointing up. The X axis is roughly in the east-west direction, and right-handedness requires that the positive axis points west.

The origin was chosen on the beam with $z = 0$ at the interaction counter, roughly 1.5 cm downstream of the last target foil. Unless otherwise indicated, coordinate measurements are in centimeters.

Many of the tracking detectors (SMDs, Drift Chambers, and PWCs) were oriented along the X or Y axis, but some were mounted at an angle. Drift chamber wires, for instance, would sag if they were hung horizontally, so in order to do three dimensional tracking, some chambers were hung vertically and some were mounted at an angle (20.5°) with respect to the vertical. For this reason, additional coordinates are defined to make it easier to describe these detectors. The additional coordinates are U, V, W, and W' where the U axis is rotated counterclockwise with respect to the X axis (in the right-handed sense) by an angle of 20.5°, the V axis is rotated clockwise w.r.t.

the X axis by 20.5°, the W axis is rotated counterclockwise by an angle of 60° w.r.t. the X axis and the W' axis is rotated clockwise 60° w.r.t. the X axis.

Note that an X-view detector makes a measurement in the X direction, and thus has strips or wires which are vertical.

3.4 Silicon Microstrip Detectors

E791 used a total of 23 silicon microstrip detectors (SMDs). 6 were upstream of the target foils and were used for beam tracking. The remaining 17 were downstream of the target foils and were used for the vertex tracking. The angular acceptance of the silicon system was roughly 200 milliradians. The beam track information was used to accurately determine the x and y coordinates of the primary vertex, and the remaining tracks were used in constructing secondary vertices. The SMDs used in E791 came from four sources and are discussed below.

3.4.1 E691 SMDs

The E691 SMDs were nine planes inherited from E691, in the form of three X-Y-V triplets, the first having an area of 676 mm^2, and the other two 2500 mm^2.

The planes consisted of silicon wafers approximately 300 μm thick. The strips were spaced 50 μm apart, and were slightly less than 50 μm wide. To form a strip, each side of the wafer was ion implanted: one side with boron (p-type), and the other with arsenic (n-type). Aluminum was deposited over both the n-type and p-type doped layers for good ohmic contact. Thus a p-n junction diode was formed, which was reverse biased at 90 V. The first triplet wafers were manufactured by Enertec, and the other two by Micron.

When a charged particle crossed a strip an average of 23,000 electron-hole pairs

were excited into the conduction band, which corresponds to a charge of 3.7 fC. Signals this small could not travel more than a meter, so they were fanned out to pre-amplifiers. The fan-outs were copper printed on Kapton and G10.

The planes and fan-outs were mounted inside light-tight RF boxes on which were mounted pre-amplifier cages made of silver-plated aluminum for RF shielding. The cages each contained 32 four-channel Laben MSD2 preamplifier cards which had a current gain of 200. These signals were then transmitted over 4 m long shielded cables to the discriminator cards (slightly modified eight-channel Nanosystems MWPC models 5710/810).

3.4.2 E769 SMDs

These SMDs were inherited from E769, the experiment which followed E691 in the TPL. These consisted of two pairs of X and Y planes. One pair was located upstream of the target for beam tracking. The other two were downstream of the target, placed closest to the target (upstream of the 691 planes.)

The beam planes had 384 instrumented strips at a 25 μm pitch, and the downstream planes had the same and in addition 304 50 μm strips in the outer region. Having a finer pitch in the inner region increased the resolution for most vertices, without sacrificing angular acceptance.

The planes were reverse biased at 90 V, and amplified by 4-channel differential output Laben MSP1 preamplifiers. The signals were then sent to 8-channel Nanosystem 5710/810 amplifier-discriminator cards.

3.4.3 OSU SMDs

These SMDs were provided by Ohio State University for E791. They had been originally constructed for Fermilab E653. These were the largest SMD planes in the experiment and were placed furthest downstream. There were six such planes, two each in the X, Y, and V views.

The planes were manufactured by Micron, and were n-type phosphorus doped with a depletion voltage of 45 V. LeCroy HQV810 preamps were used, and the signal was further amplified by cards designed by Ohio State University.

The planes were each 100 cm^2 and had an inner pitch of 50 μm and an outer pitch of 200 μm.

3.4.4 Princeton SMDs

Princeton University provided four new planes for E791. They were placed upstream of the target and used for beam tracking. They consisted of one X plane, one Y plane and two W planes.

The planes had an inner pitch of 25 μm and an outer pitch of 50 μm. The planes were reverse biased at 70 V and the signals were fanned out to MSP1 preamplifiers. The preamp signals were sent over twisted pair ribbon cables to eight-channel Nanometric N339A amplifier-discriminator cards.

3.4.5 Princeton SMD scanners

In order to read out the nearly 16,000 channels from the SMDs ten scanner boards were built by Princeton. Each board contained eight scanners each of which could read 256 channels as one large shift register. Thus 80 channels could be processed in parallel.

The discriminator cards were designed to be read out as a shift register. A clock signal from the scanner was sent to the discriminator which returned the clock signal along with a data signal. A high signal on the data line means the strip corresponding to the clock signal was hit.

Up to eight scanner boards could be placed in one VME crate, and each crate had a controller board which read out its scanner boards and packed the data, by storing only the strip numbers of hit strips rather than the ones and zeros corresponding to each strip. This information was then sent to the event FIFO buffer. (See Section 4.2.2.)

3.5 Proportional Wire Chambers

The E791 proportional wire chambers (PWCs) consisted of a series of evenly spaced wires suspended in a gas mixture of 82.7% Argon, 17% CO_2, and 0.3% CF_4. An electric field was set up by field wires, and when a charged particle passed through, it ionized some of the gas molecules, leaving a train of electron-ion pairs. The electrons were accelerated by the electric field and gained enough energy to ionize more gas molecules. These electrons migrated to the nearest sense wires and the electric charge was amplified and read out. In the E791 system, the planes could be read out in 4 ns. The field was chosen to be in a range in which the output pulse is proportional to the number of primary ion-electron pairs produced. If the field were too weak, the primary electrons would not gain enough energy to further ionize any molecules. If the field were too strong an avalanche condition would result in which the output pulse would no longer be proportional to the number of primary electron-ion pairs.

In E791, PWCs were used for both beam tracking and downstream tracking. Eight beam PWCs were placed well upstream of the target (one set of four was roughly

Table 3.1: Summary of E791 SMD information.

Plane Number	Type†	View‡	Z Pos. (cm)	Dimen. (cm × cm)	Pitch Inr./Outr. (μm)	Number of Instr. Strips	Efficiency (%)
1	769	Y	-80.250	1. × 1.	25/50	384	90
2	769	X	-79.919	1. × 1.	25/50	384	90
3	Pr.	W	-74.529	1. × 1.	25/50	448	90
4	Pr.	W	-33.163	1. × 1.	25/50	448	90
5	Pr.	X	-30.133	1. × 1.	25/50	416	90
6	Pr.	Y	-29.483	1. × 1.	25/50	416	90
7	769	Y	0.670	1.6 × 1.6	25/50	688	83
8	769	X	1.000	1.6 × 1.6	25/50	688	85
9	691	X	1.931	2.6 × 2.6	50/50	512	93
10	691	Y	3.015	2.6 × 2.6	50/50	512	95
11	691	V	6.684	2.6 × 2.6	50/50	512	96
12	691	Y	11.046	5. × 5.	50/50	768	98
13	691	X	11.342	5. × 5.	50/50	768	97
14	691	V	14.956	5. × 5.	50/50	768	94
15	691	X	19.915	5. × 5.	50/50	1000	90
16	691	Y	20.254	5. × 5.	50/50	1000	88
17	691	V	23.878	5. × 5.	50/50	1000	93
18	OSU	V	27.558	10. × 10.	50/200	864	98
19	OSU	X	31.848	10. × 10.	50/200	864	96
20	OSU	Y	34.548	10. × 10.	50/200	864	98
21	OSU	X	37.248	10. × 10.	50/200	864	99
22	OSU	Y	39.948	10. × 10.	50/200	864	99
23	OSU	V	45.508	10. × 10.	50/200	864	99

†See Secs. 3.4.1 thru 3.4.4
‡See Sec. 3.3

Table 3.2: Summary of E791 PWC information.

No. of Planes	Views†	Z Pos. (cm)	Wire Spacing (cm)	Size (cm×cm)
4	X, Y, $X'^{\ddagger}$, W	-3116.0	0.1	6.4 × 3.2
4	X, Y, X', W	-1211.0	0.1	6.4 × 3.2
1	Y	118.5	0.2	53 × 28.8
1	Y	161.1	0.2	53 × 28.8

†See Sec. 3.3
‡X' coordinate offset 1/2 cell from X coordinates.

30 m upstream and the other 12 m upstream (See Table 3.2.)). This long lever arm gave a good angular resolution of ~ 10 μrad. In addition, two PWCs were located downstream of the SMDs but before the drift chambers and magnets to assist in linking the SMD tracks to the drift chamber tracks.

3.6 Magnets

Magnets were necessary for measuring the momentum of charged tracks. When a particle of charge $+e$ passes through a magnetic field perpendicular to its path, it undergoes a deflection through an angle θ, given by [Fe86, p. 327]:

$$\theta \approx \frac{\int B\, dl}{3.333p} \tag{3.3}$$

where p is the particle's momentum, measured in GeV/c, B is the magnitude of the magnetic field in Tesla, and the infinitesimal length element dl is measured in meters. Thus the momentum resolution can be expressed in terms of the angular resolution as:

$$\frac{\delta p}{p} = \frac{\delta \theta}{\theta}. \tag{3.4}$$

Table 3.3: E791 Magnet Parameters

Magnet	M1	M2
Z bend point (cm)	275.08	620.65
Front aperture (cm×cm)	154×73	154×69
Rear aperture (cm×cm)	183×91	186×86
Current (A)	2500	1800
Coils	2	4
$\int B\,dl$ (T-m)	0.711	1.077
p_t kick (GeV)	0.212	0.320

If the magnets bend in the x direction, and r is the lever arm to the x measurement:

$$\delta\theta = \frac{\delta x}{r} \tag{3.5}$$

Substituting into the previous equation we obtain:

$$\frac{\delta p}{p} = \frac{3.333 p\,\delta x}{r \int B\,dl} \tag{3.6}$$

This equation tells us that good momentum resolution depends on having good x position resolution, a long lever arm and a large $\int B\,dl$.

In E791, two large-aperture, copper-coil dipole magnets (Figs. 3.2,3.3) were used for momentum measurement. For most E791 particle tracks the momentum component transverse to the Z axis (p_t) was much smaller than the Z component (p_z), the effect of Equation (3.3) can be approximated by saying a p_t "kick" is imparted to the particle by a magnet. The E791 magnet parameters are summarized in Table 3.3. For both magnets the field pointed in the $-Y$ direction, giving positively charged particles a kick in the $+X$ direction. In track reconstruction, a single-bend point

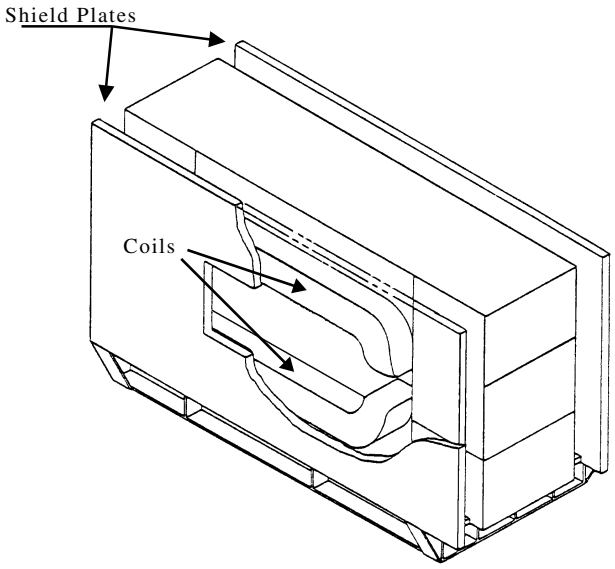

Figure 3.2: First magnet: M1 (Source: Fermilab)

approximation was used, in which the p_t kick was applied at a single point. This was a good approximation to reality and saved reconstruction time.[2]

3.7 Drift Chambers

The SMDs (Sec. 3.4) provided high resolution tracking for reconstructing vertices. To measure the momenta of particles, it was necessary to measure how much they were deflected by the magnets. For this reason, tracking had to be done between and downstream of the magnets. Because these regions are much further downstream of the target, detectors with a much greater area than that practical with SMDs were required to maintain a reasonable angular acceptance. A total of 35 drift chamber planes were used by E791 for this purpose.

A drift chamber is similar in principle to a PWC (Sec. 3.5) but achieves a much

[2]Actually, some of the data was reconstructed using a "full field" model. This is discussed in Section 5.3.4.

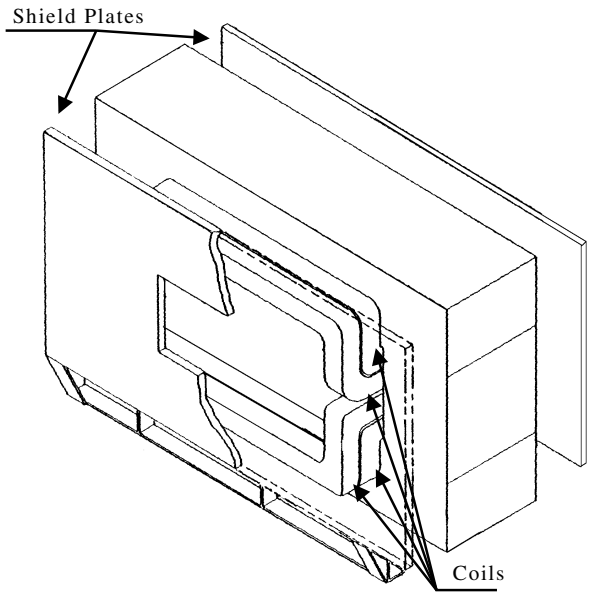

Figure 3.3: Second magnet: M2 (Source: Fermilab)

better resolution by using time information in addition to simply recording which wires are hit. The E791 drift chambers were mounted inside gas-tight rf boxes filled with a mixture of 89% Ar, 10% CO_2, and 1% CF_4. Inside were alternating cathode and sense planes, and the sense planes consisted of alternating sense wires and field-shaping wires. The cathode planes were held at a potential of -2.1 kV to -2.5 kV, with the field-shaping wires 0.4 kV to 0.6 kV higher, and the sense wires were grounded. This created a fairly uniform field. The field strength and gas mixture were chosen so that when electrons were liberated by an ionizing particle, they drifted with a fairly constant velocity, v_d, to the nearest sense wire. By measuring the time required for the signal to appear (Δt), and using the simple relation $d = v_d \Delta t$, one can determine the distance from the sense wire that the ionizing particle passed, and, up to a left-right ambiguity, the particle's coordinate. (Having more than one plane helps to clear

Table 3.4: E791 Drift Chamber Summary

Asmbly	Stat.	Z Pos. (cm)	Views	Size (cm×cm)	Chan.	Res. (μm)
1	D1	142.5 to 148.5	X, X', U, V	126 × 71	672	400
2	D1	177.6 to 183.6	X, X', U, V	126 × 71	864	400
3	D2	382.3 to 385.5	U, X, V	285 × 143	544	300
4	D2	424.9 to 428.1	U, X, V	285 × 143	544	300
5	D2	466.6 to 469.8	U, X, V	285 × 143	640	300
6	D2	497.7 to 500.9	U, X, V	285 × 143	672	300
7	D3	928.2 to 931.5	U, X, V	323 × 143	480	300
8	D3	970.8 to 974.0	U, X, V	323 × 143	480	300
9	D3	1012.5 to 1015.8	U, X, V	323 × 143	480	300
10	D3	1044.3 to 1047.6	U, X, V	323 × 143	512	300
11	D4	1737.8 to 1749.3	U, X, V	511 × 259	416	450

up the left-right ambiguity.) The gas and field strengths used in E791 resulted in a drift velocity of 50 μm/ns.

The drift chambers measured in the X, U, and V views (See Sec. 3.3). No Y planes were used because the horizontally hung wires would sag too much, and degrade the resolution. The drift chambers were divided into four stations, labeled D1, D2, D3, and D4. (See Table 3.4.) D1 consisted of eight planes located before the first magnet (M1) downstream of the SMDs and PWCs. It increased the acceptance of region one tracking.[3] D2 was located between the magnets and consisted of four triplets and was used for region 2 tracking. Downstream of M2, D3 also consisted of four triplets and was used for region 3 tracking. Just upstream of the calorimeters was another triplet (D4). Its long lever arm could have improved momentum measurements, but its poor resolution limited its usefulness.

The drift chamber signals were read out by Lecroy DC201 Nanomaker N-277C

[3]The region upstream of M1 was known as region 1. Region 2 was between the two magnets and region 3 was downstream of the magnets.

amplifier/discriminators. The signals were then transmitted via twisted pair cables to Philips 10C6 time-to-digital converters (TDCs) which were read out by Fastbus Smart Crate Controllers to the DA system.

The TDCs were operated in common stop mode. They started counting when the corresponding sense wire received a signal and were all stopped at the same time by a delayed trigger. The drift time could then be obtained by the following simple formula:

$$t_{\rm drift} = t_{\rm del} - t_{\rm tdc} - (t_{\rm abs} + t_{\rm rel}) \tag{3.7}$$

where $t_{\rm del}$ represents the delay time between the trigger and the TDC stop signal, $t_{\rm abs}$ is an absolute time offset due to detector offsets, and $t_{\rm rel}$ is a relative time offset for each channel resulting from differences in cable lengths and responses for each sense wire channel.

The drift chambers performed well, but did suffer from a problem resulting from the use of charged beam particles. Most such particles did not interact, and passed right through the detector system. The large number of ionizing particles passing through the central regions of the drift chambers caused deposits to gradually build up on the drift chamber wires in those regions resulting in localized losses of efficiency, known as "drift chamber holes". These holes changed in size and shape during the run period, which made acceptance modeling difficult.

3.8 Čerenkov Counters

In reconstructing the momenta of charm particles, it is necessary to know not only the momenta of the decay products, but the rest masses and hence the identities of the particles as well. In many modes, such as the Cabibbo favored $D^+ \to K^-\pi^+\pi^+$,

the charge of the particles can be used to identify them. In a suppressed mode such as $D^+ \to K^- K^+ \pi^+$, however, charge alone is not sufficient. In order to help identify the particles, E791 used Čerenkov counters. [Ba87]

When a charged particle passes through a dielectric medium with a velocity greater than light travels in the medium (c/n where n is the index of refraction for the medium), it emits light which is confined to a cone centered on the particle's path which makes an angle θ with the particle's path where $\cos\theta = c/vn$.

The condition that $v > c/n$ means the particle must have a momentum greater than a threshold momentum p_{th} to emit Čerenkov light. Since $p = \gamma m v$ we have:

$$p_{th} = \frac{\gamma m c}{n} \tag{3.8}$$

And $\gamma = (1 - (v/c)^2)^{-1/2}$ so at threshold $\gamma = (1 - n^{-2})^{-1/2}$. And so,

$$p_{th} = \frac{mc}{\sqrt{n^2 - 1}} \tag{3.9}$$

The number of photons N emitted by a singly charged particle per unit wavelength λ per unit length l is given by [Fe86, p. 181]:

$$\frac{dN}{dl\, d\lambda} = \frac{2\pi\alpha}{\lambda^2}\left(1 - \frac{v_{th}^2}{v^2}\right) \tag{3.10}$$

where α is the fine structure constant. In terms of measured quantities, this becomes:

$$\frac{dN}{dl\, d\lambda} = \frac{2\pi\alpha}{\lambda^2}\left(1 - \frac{m^2c^2 + p^2}{n^2 p^2}\right) \tag{3.11}$$

Thus, with a known index of refraction and a given mass hypothesis, one can predict the number of photons that will be emitted at a given momentum. By measuring

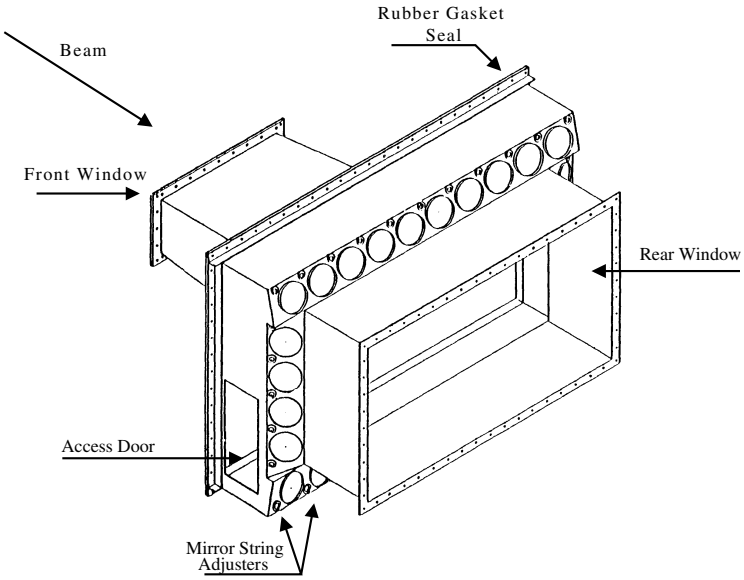

Figure 3.4: First Čerenkov Counter: C1 (Source: [Ba87])

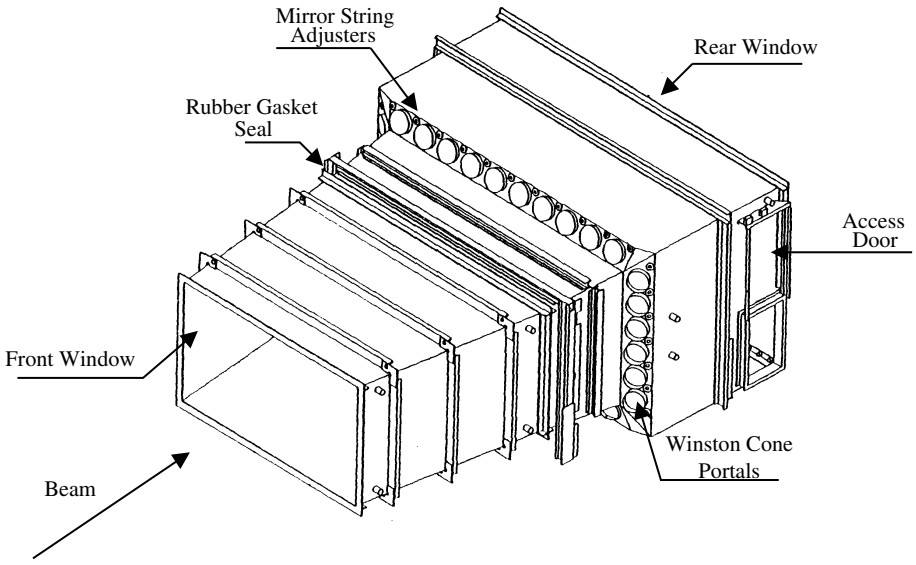

Figure 3.5: Second Čerenkov Counter: C2 (Source: [Ba87])

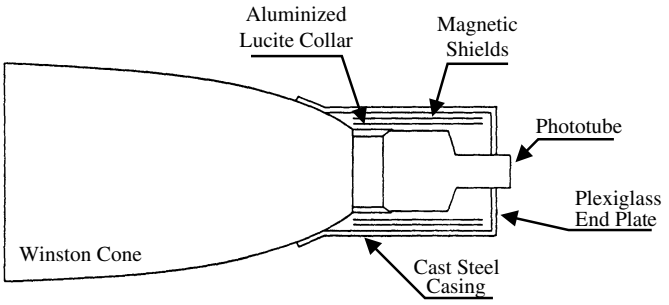

Figure 3.6: Winston Cone used by E791. (Source: [Ba87])

the momentum and number of photons, one can assign probabilities to various mass hypotheses and use this information in particle identification.

E791 used two Čerenkov counters labeled C1 and C2. (See Figs. 3.4 and 3.5). The first was filled with N_2, and the second with a mixture of 80% He and 20% N_2. By using two different gas mixtures, the two counters had different indices of refraction and thus could identify particles in different momentum regions. C1 was located just downstream of magnet M2, and C2 was located between drift chambers D3 and D4.

In order to maximize the number of photons collected, the E791 Čerenkov counters used mirrors and Winston cones to collect the light. A Winston cone is a reflecting surface whose geometry is chosen to collect incident light up to an angle $\theta_{\max}$ w.r.t. the optical axis onto the face of a photomultiplier tube. The E791 Winston cones had $\theta_{\max} \approx 20^\circ$. (See Fig. 3.6).

The photomultiplier tubes used in E791 were RCA 8854 5-inch tubes. They were chosen because of their very good linearity, i.e. the output signal was roughly proportional to the number of photoelectrons. This feature was tested and calibrated using highly attenuated laser light. The output was digitized by LRS 2249 ADCs and single, double, and triple photoelectron peaks were observed with the expected number of ADC counts. In addition to this calibration, each mirror-cone-phototube assembly was calibrated using high-momentum, isolated tracks that were centered on

Table 3.5: Summary of E791 Čerenkov detector information.

	C1	C2
Length (cm)	370	660
# of mirrors	28	32
Gas Mixture	100% N_2	80% He 20% N_2
$n-1$	290×10^{-6}	86×10^{-6}

a single mirror. The momentum of the track was required to be high enough so that dN/dl was nearly constant. That way, the number of photons depended only on the path length of the track through the region where light is expected to be collected by the given cone. By comparing the number of predicted to the number of measured photons the efficiency of each assembly could be determined.

The arrangement of the mirrors in E791's detectors is shown in Figure 3.7 and the characteristics are shown in Table 3.5. The mirrors ranged in size from 6650 cm^2 in the outer regions down to 300 cm^2 in the inner regions where there were a larger number of particles. In C2, light was reflected off a mirror plane into the Winston cones, and in C1, space constraints required double bounce optics as shown in Figure 3.8.

To use the Čerenkov detector for particle identification, measured photons were first assigned to the various charged tracks passing through the detector. Then, a probability was constructed for observing the measured number of photons ($N^{\rm meas}$) when the predicted number of photons is $N^{\rm pred}$. This is given by the Poisson distribution:

$$f_{ij}(N_{ij}^{\rm pred}; N_i^{\rm meas}) = \frac{(N_{ij}^{\rm pred})^{N_i^{\rm meas}} e^{-N_{ij}^{\rm pred}}}{N_i^{\rm meas}!} \tag{3.12}$$

Here, the label $i \in \{1, 2\}$ refers to the counter number and j refers to a mass hypothesis. The joint probability for a given measurement in the two counters is then $f_j = f_{1j} \times f_{2j}$. This gives the probability for an observed measurement given a mass

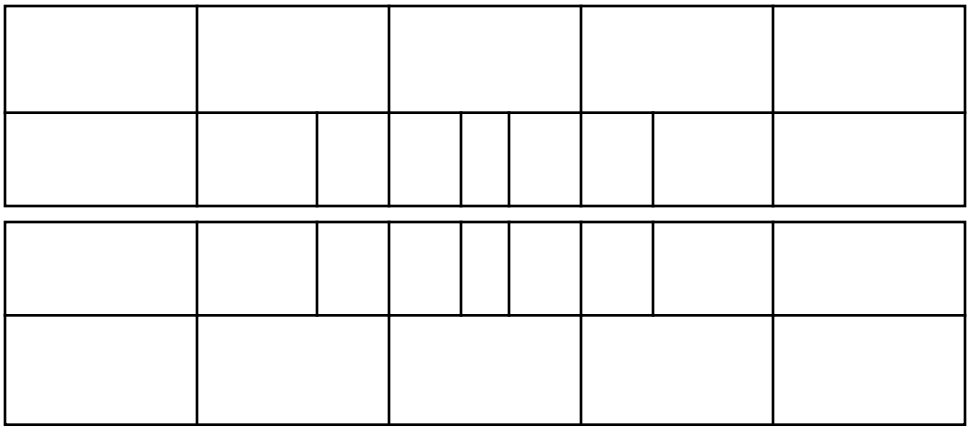

C1 MIRROR ARRAY

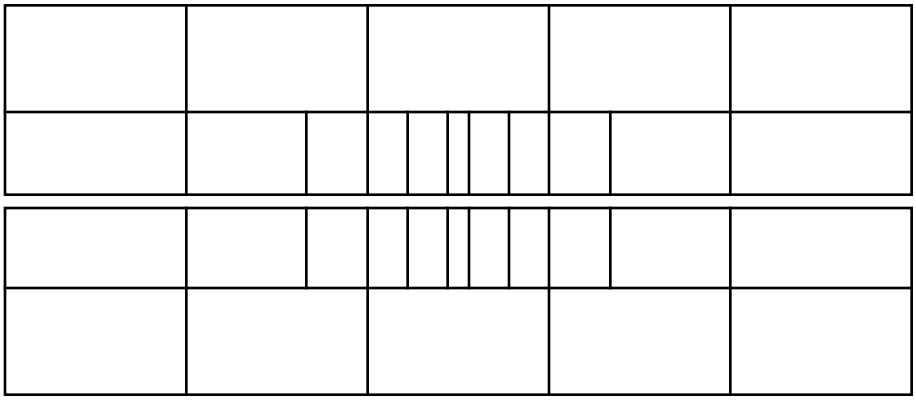

C2 MIRROR ARRAY

Figure 3.7: Mirror arrays for E791 Čerenkov detectors.

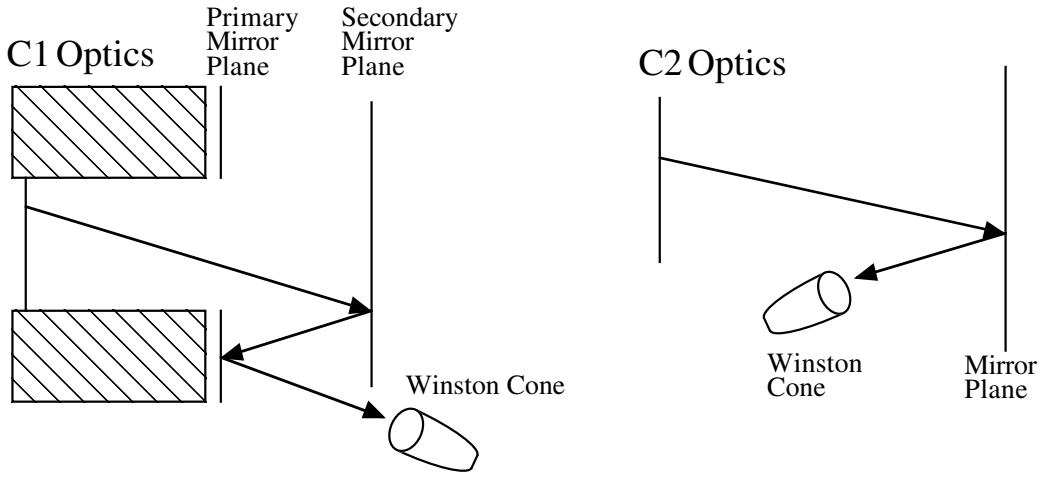

Figure 3.8: Optical arrangements for E791 Čerenkov detectors. (Source: [Ba87])

hypothesis. To turn this around and find the probability for a given mass hypothesis, we must apply Bayes's theorem and include the *a priori* probabilities for detecting each particle:

$$P_j = \frac{f_j \times A_j}{\sum_k f_k \times A_k} \tag{3.13}$$

which gives the probability for a given mass hypothesis j. The E791 Čerenkov detectors were designed to detect electrons, muons, pions, kaons, and protons. The *a priori* probabilities in E791 were $A_e = 0.02$, $A_\mu = 0.01$, $A_\pi = 0.81$, $A_K = 0.12$, and $A_p = 0.04$.

3.9 Calorimetry

The calorimeters served two main purposes in E791. One was triggering which is discussed in the next chapter. The other was to provide identification for electrons and neutral hadrons.

Although the Čerenkov detectors were very good at distinguishing pions from kaons, electrons generally emitted light in both detectors, making them of limited

usefulness in identifying electrons. Also, neutral hadrons and photons don't leave tracks in other parts of the detector, and so calorimetry is quite useful in identifying those particles.

E791 employed both a segmented liquid ionization calorimeter for electromagnetic particles and a steel-acrylic scintillator calorimeter for hadrons.

3.9.1 Electromagnetic Calorimeter (SLIC)

When photons and electrons (or positrons) with an energy above about 10 MeV interact with matter, bremsstrahlung and pair production are the dominant processes. Electrons and positrons emit photons by bremsstrahlung, and photons produce e^+e^- pairs. This results in an exponential growth in the number of electromagnetic particles until the energy of the resulting particles falls below the level at which bremsstrahlung and pair-production are the dominant processes. As a result, the depth of the shower in space is roughly proportional to the log of the energy of the incident particle. This explains why electromagnetic showers are relatively short.

During this shower, energy can be transferred to a scintillating material by Compton scattering or ionization. The light emitted by the scintillator is thus proportional to the energy of the incident particle.

For electromagnetic calorimetry, E791 employed a segmented liquid ionization calorimeter (SLIC). [Bh78, Bh85] The SLIC consisted of 60 layers of corrugated aluminum filled with a mineral oil based liquid scintillator (NE235H). (See Figure 3.9). Each corrugation was 3.17 cm wide and 1.27 cm deep. The aluminum was coated with teflon which has a lower index of refraction than the oil, so that each corrugation formed a totally internally reflecting light pipe for light with an angle of incidence less than about 20°. Separating the corrugated aluminum sheets were 0.63 cm

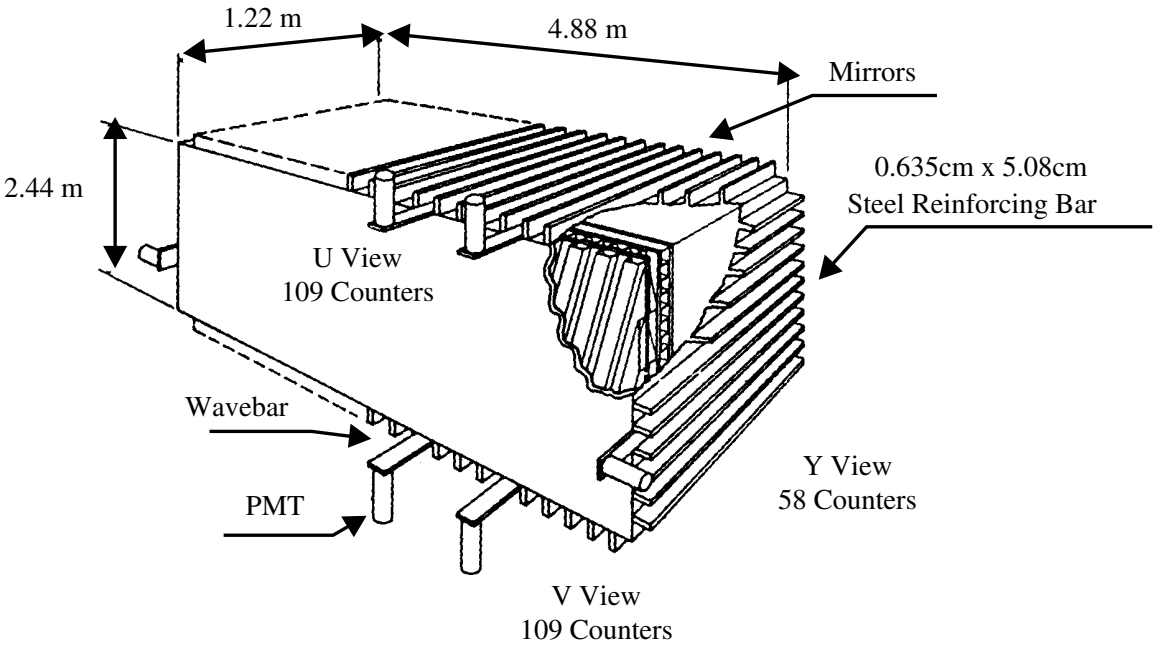

Figure 3.9: The SLIC. (Source: [Bh85])

aluminum-laminated lead sheets, which helped to initiate and maintain the showers. (See Figure 3.10.)

The layers alternated in U, V, and Y views (the U view had corrugations at an angle of 20.5° with respect to the vertical, the V view $-20.5°$, and the Y view had horizontal corrugations.)

Channels were read out individually in the central region and in groups of two in the outer regions. At one end of each channel was a mirror with a reflectivity of 80%, and the other was connected to a waveshifter bar which connected all channels at a given position in a given view. The waveshifter bar was doped with a waveshifting material (BBQ), which absorbed the ultraviolet light from the scintillator and re-emitted it as green light. This light was collected by 2 inch RCA 4902 photomultiplier tubes attached to the end of the waveshifters. The linear pulses from the PMTs were digitized by Fast Encoding and Readout ADCs (FERAs). The SLIC was cali-

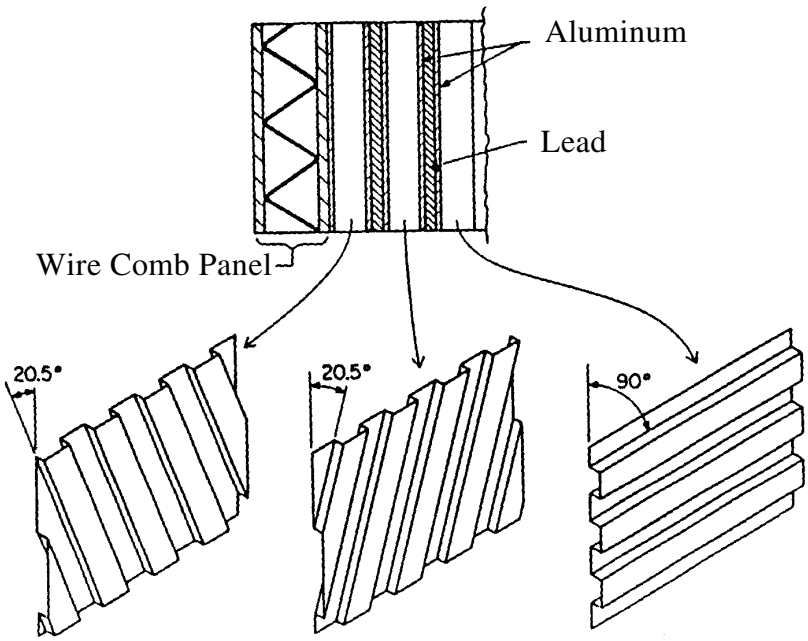

Figure 3.10: Layer construction of the SLIC. (Source: [Bh85])

brated periodically throughout the run period using muon and electron runs. SLIC characteristics are tabulated in Table 3.6.

Particle energies were calculated as follows. [Su85] First, "cells" were found where a cell is defined as a cluster of counters with enough light to be above an energy threshold. Because multiple particles may contribute to a cell, each cell was examined for multiple peaks. "Sectors" were thus formed where a sector corresponded to a single peak. The V and Y positions of the sector were matched with drift chamber tracks to predict a corresponding U, and this was required to match the sector's U position. The energy for the particle candidate was found by minimizing:

$$\chi^2 = \sum_i \frac{(e_i - \sum_j \alpha_{ij}\varepsilon_j)^2}{\sigma_i^2} \tag{3.14}$$

where e_i are the sector energies with errors σ_i, α_{ij} are position dependent energy

Table 3.6: Summary of E791 SLIC information.

Z position (cm)	1866 to 1962
# of U channels	109
# of V channels	109
# of Y channels	116
Usable area (cm $\times$ cm)	490×240
Radiation Lengths	20
Absorption Lengths	1.5

correction factors, and ε_j are the energies of the particle candidates, which were the fit parameters.

The SLIC made it possible to reconstruct energy showers with an energy resolution of:

$$(\frac{\Delta E}{E})^2 \approx (11.5\%)^2 + (\frac{17.4\%}{E})^2$$

and a position resolution of 7 mm.

3.9.2 The Hadrometer.

Hadronic shower processes are different from those of electromagnetic showers. About half the incident energy goes into the production of fast secondaries whose large transverse momentum components give rise to showers of greater lateral extension than electromagnetic showers. The remainder goes into the production of slow pions and other particles.

The E791 hadrometer [Ap86] was located just downstream of the SLIC. Since the SLIC absorbed roughly 99% of the energy of electromagnetic showers, hadrons could be distinguished by the fact that they deposited a significant amount of energy in the hadrometer as well as a wider shower in the SLIC. Neutral hadrons could be identified after subtracting the hadronic energy associated with charged tracks.

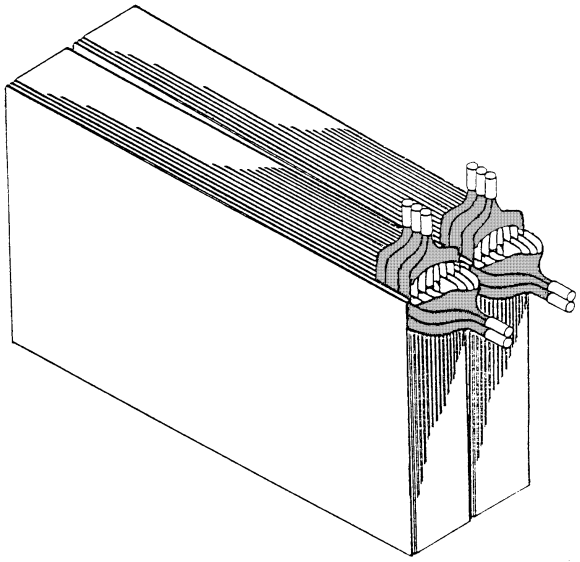

Figure 3.11: The E791 hadrometer. (Source: [Ap86])

Table 3.7: Summary of E791 hadrometer information.

Z position (cm)	1973 to 2131
# of channels	142
Usable area (cm $\times$ cm)	490×270
Absorption Lengths	6

The hadrometer (see Figure 3.11) consisted of 36 layers of acrylic scintillator separated by 2.5 cm thick steel plates. The scintillator layers alternated in x and y views, consisting of vertical (x) and horizontal (y) strips, each 14.5 cm wide. The hadrometer was divided into two modules, and for each module light from strips of the same x (or y) position were collected in a single phototube. 5 inch EMI 9791KB phototubes were used, and the output was digitized by FERAs.

The hadrometer was calibrated using muon runs, where the muons were known to deposit an average energy of 2.4 GeV in the hadrometer, and by minimizing the energy resolution of charged hadrons by requiring the SLIC energy plus the hadrometer

Table 3.8: Summary of E791 muon detector information.

	X wall	Y wall
Z position (cm)	2143	2319
Usable area (cm × cm)	550 × 300	300 × 232
Number of paddles	15 + 4	16

energy to add up to the energy calculated from the measured momentum. The energy resolution of the hadrometer was roughly $\frac{\Delta E}{E} = \frac{75\%}{\sqrt{E}}$.

The hadrometer information is summarized in Table 3.7.

3.10 Muon Detector

Because muons and pions are fairly close in mass, the Cerenkov detector is not effective in identifying muons. The muon detector consisted of two walls of plastic scintillator paddles, downstream of the hadrometer and separated from it by 102 cm of steel (6 absorption lengths). The steel absorbed most of the remaining hadrons, and particles reaching the muon walls were *a priori* assumed to be muons.

The X wall (originally constructed for E516) consisted of 15 plastic scintillator paddles mounted vertically and housed in a light tight box. The 3 center-most paddles were 61 cm wide and the remaining 12 were 41 cm wide, and all paddles were 300 cm long, with paddles in the central region divided lenghwise into three pieces (120 cm, 60 cm, and 120 cm), each read out by their own photomultiplier tube. Light from the scintillator was collected by EMI9791KB phototubes placed at the top of each paddle. No timing information was provided by this system; only a latched signal was recorded. Because the centermost paddle was directly in the beamline, it was the most noisy and inefficient. Therefore, four additional X-view paddles were

mounted on a block of concrete downstream of the X-wall, in a position overlapping the center paddle. The performance of the X wall was rather poor, with an efficiency of approximately 69%.

The newer Y wall was located downstream of the X wall and was smaller. It consisted of 16 NE110 plastic scintillator paddles mounted horizontally to measure the Y coordinate of charged particles passing through. The paddles were 14.5 cm wide and 300 cm long. Phototubes on the east side of the paddles collected the light. TDCs operated in common-stop mode were used to record timing information of the signals; larger counts corresponded to hits closer to the phototubes. An X-view resolution of 66 cm was possible using this timing information. The Y wall was much more efficient than the X wall, with an efficiency of approximately 99%.

The characteristics of the muon system are summarized in Table 3.8.

Chapter 4

Trigger and Data Acquisition System

4.1 Trigger

Data from the detector components were written out to tape in a series of discrete packets called "events". Ideally, each event represents the detector data corresponding to one beam particle interaction. In order to determine when such an interaction has taken place, a set of criteria are established, and hardware is built which can rapidly determine when these criteria are satisfied. This system is called a "trigger", and a signal from the trigger causes data from all detector systems to be read out, packaged into an event record, and written to tape.

The goal of E791 was to collect a large set of events in which a charm particle is produced. Charm events were quite rare, and determining whether or not an interaction produced a charm particle required tracking and vertexing information which was difficult to produce in real-time, especially at the high rate of interaction in a hadroproduction experiment. For this reason E791 elected to use a very "open"

trigger, essentially requiring only a simple interaction, and to do the search for charm events in software.

The trigger was divided into two parts, a "pretrigger" used to determine that some form of interaction took place, and information from the calorimeters, used to determine that a certain amount of energy was deposited away from the beam axis.

4.1.1 Pretrigger

The pretrigger consisted of three pieces of scintillator located near the target. Figure 4.1 shows their position relative to the target, and Table 4.1 summarizes their geometry information. The Beam Spot counter and Beam Halo counter were located upstream of the target foils. The Beam Spot counter was used to determine that a pion was incident on the target. It had a low and high energy threshold: the low threshold was set to require that at least one beam pion passed through it, and the high threshold was set to ensure that at most one beam particle reached the target. The Beam Halo counter was located between the Beam Spot counter and the target. It was larger than the Beam Spot counter and had a 1 cm hole drilled in the center. Its purpose was to reject off axis beam particles, for which the tracking system had very low acceptance. Downstream of the target foils was the Interaction Counter. Its purpose was to ensure that some form of interaction took place by requiring that at least four minimum ionizing particles passed through it. To summarize, the pretrigger required the following three conditions to be true:

1. Beam Spot counter detected one and only one beam pion, AND

2. Beam Halo counter had no signal, AND

3. Interaction counter detected at least four minimum ionizing particles.

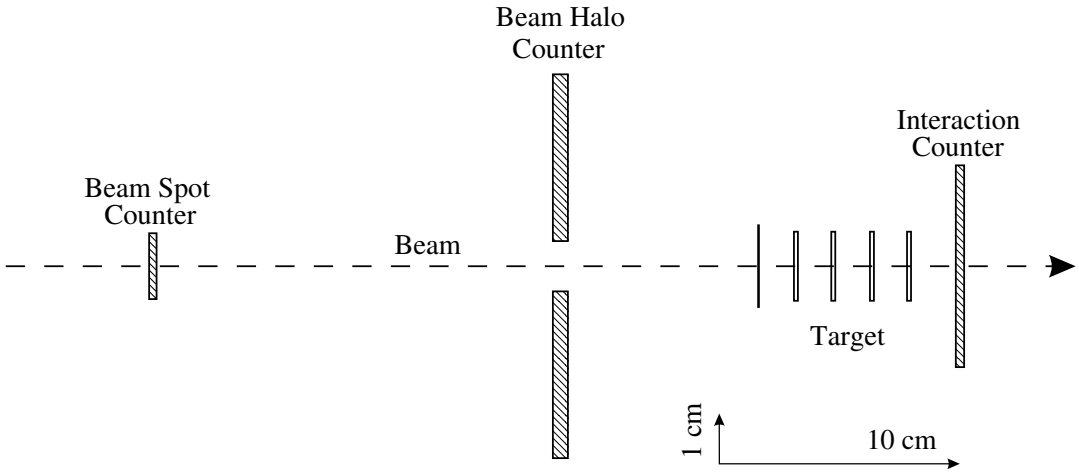

Figure 4.1: Scale drawing of pretrigger scintillator paddles and target foils.

Table 4.1: Summary of E791 pretrigger counter geometry.

	Beam Spot	Beam Halo	Interaction Counter
Area (cm× cm)	1.3×1.3	7.6×7.6	4.0×4.0
Thickness (cm)	0.3	0.6	0.3
Hole dia. (cm)	N/A	1.0	N/A
Z position (cm)	-22.7	-16.3	0.0

The pretrigger took 160 ns to compute. During a 23 second spill (see Section 2.2), approximately 350,000 pretriggers were collected from 2×10^7 incident beam pions, for an average rate of approximately 16,000 pretriggers per second during the spill.

4.1.2 E_t trigger

E791 ran in a few different trigger modes. Some were used for calibration and detector alignment, but the predominant mode, used for physics analysis, was the E_t trigger mode, and it will be the only one discussed here.

Information from the calorimeters (see Section 3.9), was used to ensure a sufficient amount of energy was deposited away from the beam axis. Signals from the last

dynode of each photomultiplier tube were summed together giving greater weight to those channels further from the beam. The resulting quantity is referred to as the "transverse energy" or E_t. The measured E_t was compared to high and low thresholds which could be adjusted. The low threshold was used to primarily to reject diffractive pion events and the high threshold was used to reject multiple beam particle events. The full E_t trigger required the following:

1. Pretrigger was true, AND
2. $E_t > 4$ GeV to reject diffractive pion events, AND
3. $E_{\rm tot} < 700$ GeV to reject multiple beam particle events.

It took 470 ns to collect the calorimetry information and compute E_t. Approximately 180,000 full triggers were taken per spill, about half the number of pretriggers. When a full trigger was registered, data from all detectors were digitized and sent to buffers for packing and recording.

4.2 Data Acquisition System

As described in Section 2.2 pions were delivered to the target during a 23 s spill, followed by a 34 s interspill period. As noted in the previous section, roughly 180,000 triggers are taken per spill, meaning that events had to be acquired by the data acquisition (DA) system at a rate of at least 7,800 per second and written to tape at a rate of over 3,000 per second, as events were buffered and written to tape during both the spill and interspill periods. In order to lose as few events as possible, the DA system was designed to have a dead time of only 50 μs while digitizing and reading out the 24,000 detector channels. Also, a single compressed event required roughly

2.3 kilobytes to store, so the tape system had to be capable of writing at a rate of 9 megabytes per second. This required a highly parallel design with rapid error recovery.

The data acquisition system [Am93] consisted of several parts. The front-end controllers read data from the various digitizers. The Event FIFO Buffers (EFBs) stored data so events could be processed during both the spill and interspill periods. VME crates housed the various boards which read data from the EFBs and compressed and formatted the events into tape records. The Exabyte tape drives continuously recorded the events for later processing. Finally, a VAX 11/780 computer allowed the operator to control and monitor the DA system, and to start and stop data taking runs.

4.2.1 Front-end Controllers

The task of the front-end controllers [Ha87b, Br02] was to read data from the various digitizers belonging to the different detector elements, and to send the data to the EFBs. The controllers had to pack the data into 32-bit longwords, which were then transmitted to the EFBs along 32-bit wide RS-485 lines together with a RS-485 strobe signal. The maximum data rate for a controller was 100 ns per longword. Different controllers had to be built for the different types of detectors, but the interface to the EFBs was common to all controllers. Each data segment sent to the EFBs included a leading word count as well as a 4-bit event synchronization number, which ensured that data for a given event was not mixed up with data from another event. These synchronization numbers were generated by a scaler and sent to all the front-end controllers to be included in their data segments. Since there were a total of 16 front-end controllers, and only 8 EFBs, each controller had to share a data path with

one other controller, and a simple token passing mechanism was used to enable this sharing.

4.2.2 Event FIFO Buffers

The Event FIFO (First-in First-out) Buffers held the event data before it was written to tape. This allowed the tape drives to run continuously, writing data during both the spill and interspill periods. Each of the 8 EFBs consisted of an I/O card, a FIFO controller card, and 5 16 Mb memory cards, as well as a custom backplane. They were housed 2 per crate in 9U by 220 mm Eurocrates.

Input and output was done along 32-bit RS-485 data paths, and each was capable of receiving data at burst speeds of up to 40 Mb/s and concurrently delivered data at several Mb/s. A Near Full status line was used to inhibit the trigger if any EFB was in danger of becoming full.

4.2.3 VME Crates

A total of 6 VME crates housed the electronics which read data from the EFBs, compressed the data and packaged them into events, and wrote them to tape.

Figure 4.2 shows a schematic of the VME part of the DA system. Each VME crate housed 8 Event Buffer Interface (EBI) boards, each of which was connected to an EFB, 9 CPU boards (8 Event Handler (EH) boards and one Boss board), 2 SCSI Magnetic Tape Controller boards (MTC) and a Branch Bus to VME Interface (BVI) board for connection to the VAX.

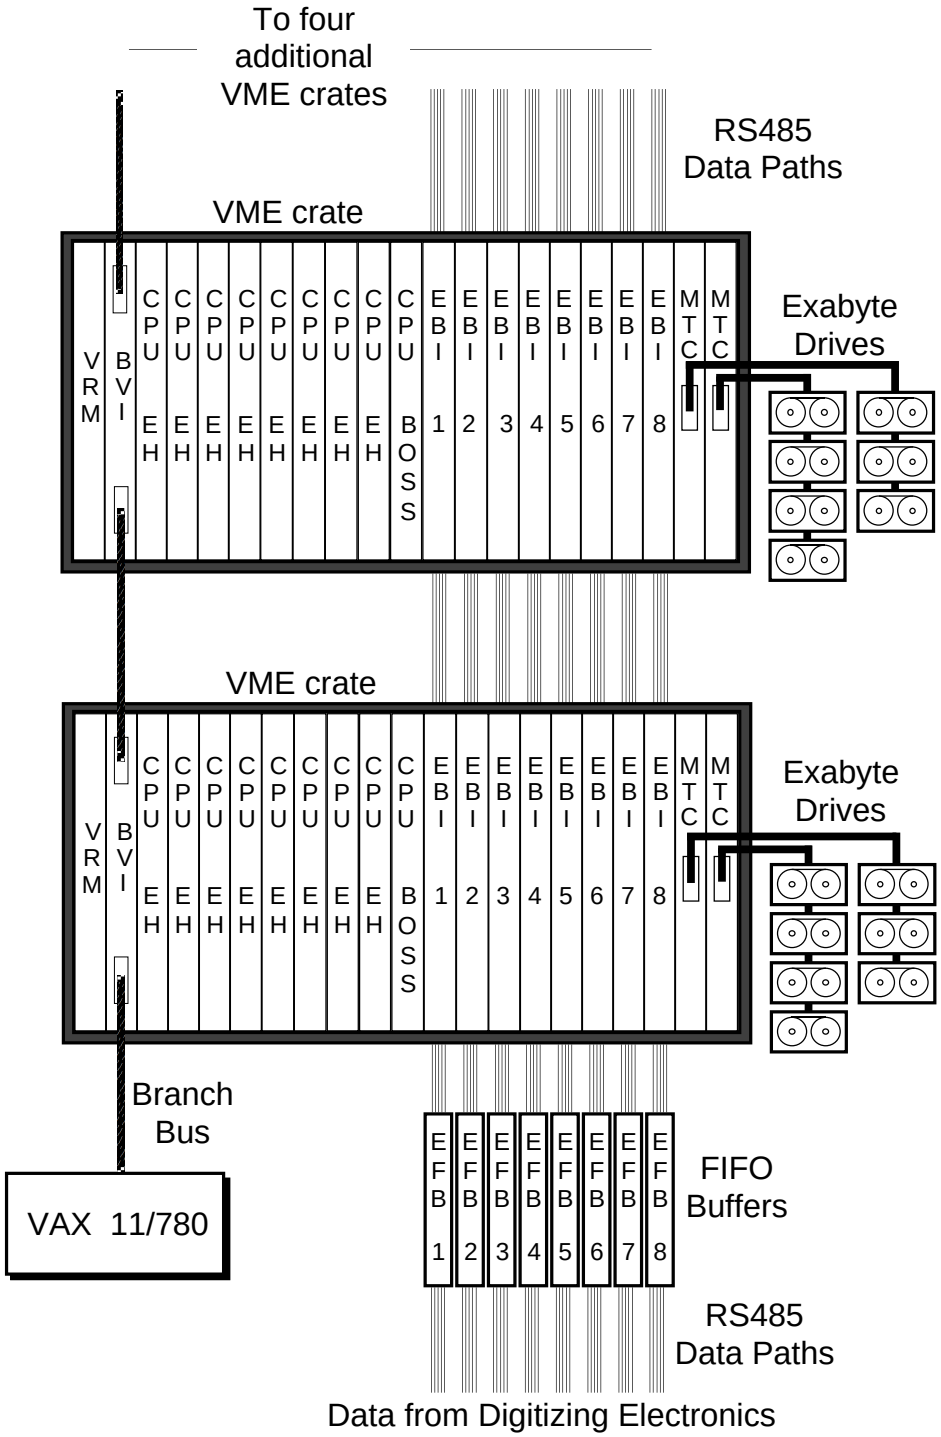

Figure 4.2: Data Acquisition system: VME crates. Two of the six crates are shown. (Source: [Am93])

Event Buffer Interface

The job of the Event Buffer Interface (EBI) boards was to strobe data out of the EFBs and make them available to the VME CPU boards. Data for a single event are spread across all eight EFBs, so each CPU required access to all of the EFBs. Each crate therefore contained 8 EBIs, one for each of the EFBs. Thus, the output of each EFB was connected to 6 crates through the EBIs. The EBIs used a token passing mechanism to establish which of the 6 had control of the RS-485 datapath at a given time. The EBI strobed data from the EFB only when it had the token. When the controlling CPU finished reading data from that EFB, the token was passed to the corresponding EBI in the next crate.

VME CPU Boards

The VME CPU boards (Fermilab ACP 1 nodes) contained a 16 MHz Motorola 68020 microprocessor, a 68881 coprocessor, and 2 Mb of memory. Software for the boards was written in Absoft Fortran, with time-critical parts written in assembly language. Each of the crates contained 9 such boards: one Boss and eight Event Handler (EH) boards. The EH boards could be in one of two possible states: "grabber" or "muncher". The Boss was responsible for controlling the states of the EH's. Each crate had only one EH in the grabber state at any given time; the others were in the muncher state. While in the grabber state, the EH read data from the EFBs in sequence (through the EBIs), collecting the data from whole events, and storing them in its on-board memory. When the grabber's on-board memory became nearly full (approximately 200–300 whole events), the Boss switched the EH's state from grabber to muncher, and assigned a new grabber for the crate.

We can now see how the high degree of parallelism was achieved. At the beginning

of a run, the first crate's EBIs had all the tokens, and the crate had one EH in the grabber state. It began reading the first segment of the first event from the first EFB. Once that segment was read, it passed the token for the first EFB to the second crate and read the next segment of the first event from the second EFB. While it was doing this, the second crate's grabber was reading the first segment from the second event from the first EFB. After this step, the first crate read the third segment from the first event while the second crate read the second segment from the second event while the third crate read the first segment from the third event, and so on. Soon, grabbers from all six crates were busy reading data out of the EFBs.

While reading data, the grabbers checked the word counts and event synchronization numbers to ensure that they were collecting whole events. It would have been disastrous if some missing data caused a piece of one event to be assembled with data from another event, and furthermore this error would propagate forward to all future events in the run. If a grabber detected such an error (usually due to a failure in a front-end controller) it notified the Boss, which notified the VAX, which caused all EFBs to be flushed and the system restarted. This took well under a minute, and only a few spills a day were lost due to such errors.

While in the muncher state, the EH boards compressed the data and assembled events into the E791 standard event format, each event requiring an average of 2.3 kilobytes. Data was packed into tape output buffers in records of 64 kb, and the Boss was notified when a buffer was ready for output. Each muncher had 10 such tape buffers. The Boss also assigned other duties to munchers, such as supplying status information and creating histograms of detector channels for the operator.

The CPU board running the Boss program had a variety of tasks to perform. Its primary duties were controlling the status of the EH boards as grabbers or munchers and managing the tape writing. The Boss periodically polled the EH boards for lists

of buffers ready for writing, and sent commands to the MTC command queues to read these buffers from the EH's and write them to tape. When a buffer was written, the MTC informed the Boss, which in turn notified the EH that its buffer was ready for reuse. The Boss also gathered status information and recoverable error reports to be sent to the VAX. The Boss also queried the MTCs for tape writing error reports. If a non-recoverable error was reported by a drive, it was taken offline. Also, if an EH didn't respond to a Boss request within a given time, it was reset and temporarily taken offline.

Magnetic Tape Controllers and Tape Drives

E791 used Exabyte 8200 8 mm tape drives. Each tape stored 2.3 Gigabytes, and took about three hours to fill. Exabytes were chosen because of their low media costs, which was a major factor considering the large amount of data to be collected. 7 drives were connected to each crate for a total of 42 drives, all running in parallel. Four drives writing at streaming speed were enough to saturate a SCSI bus, so each crate had two Magnetic Tape Controllers (MTCs), with four drives on one and three on the other. The controllers were Ciprico RF3513 VME to SCSI interfaces.

4.2.4 The VAX

The DA system was managed and monitored by a VAX-11/780 which was connected to the VME crates by a DR11-W on the VAX Unibus, a QBBC branch bus controller, and branch bus to VME interfaces (BVI) in each of the six VME crates [Ha87a]. Connected to the VAX were the control console, as well as several status displays. (See Figure 4.3.) The control console was used to start and stop runs as well as to monitor a run's progress. The Bosses polled the tape drives for remaining tape

capacities and sent that information to the VAX. When 20% of the drives were at least 95% full, the VAX ended the run. End of file marks were written on the tapes which were rewound and ejected, and the operator replaced the 42 tapes with blanks, closed the drive doors, and started the next run by issuing a start command on the control console. (Each run lasted roughly three hours, and contained about 90 gigabytes of data.)

The VAX allowed on-line monitoring of events in two ways. First, the event handlers always saved a few events which the VAX retrieved and placed in an online event pool, which could be accessed by other workstations in the cluster. More rapid monitoring of a part of the detector could be accomplished by requesting a histogram of detector elements, in order to identify noisy or dead channels. The user used a program to select a particular section of the detector. This program sent the request to the VAX using the DEC Mailbox facility. The VAX sent the request to the event handlers which accumulated all events for approximately one minute. The Bosses and then the VAX summed up the histogram contributions from each EH, and made the resulting histogram available in the online event pool. The program retrieved the histogram and allowed the user to view it in a variety of ways. This allowed the user to obtain very high statistics hit maps in a very short time.

At the end of each run, the VAX also created a disk file containing statistics about the run, such as number of events written to each tape, soft error counts, and other relevant data.

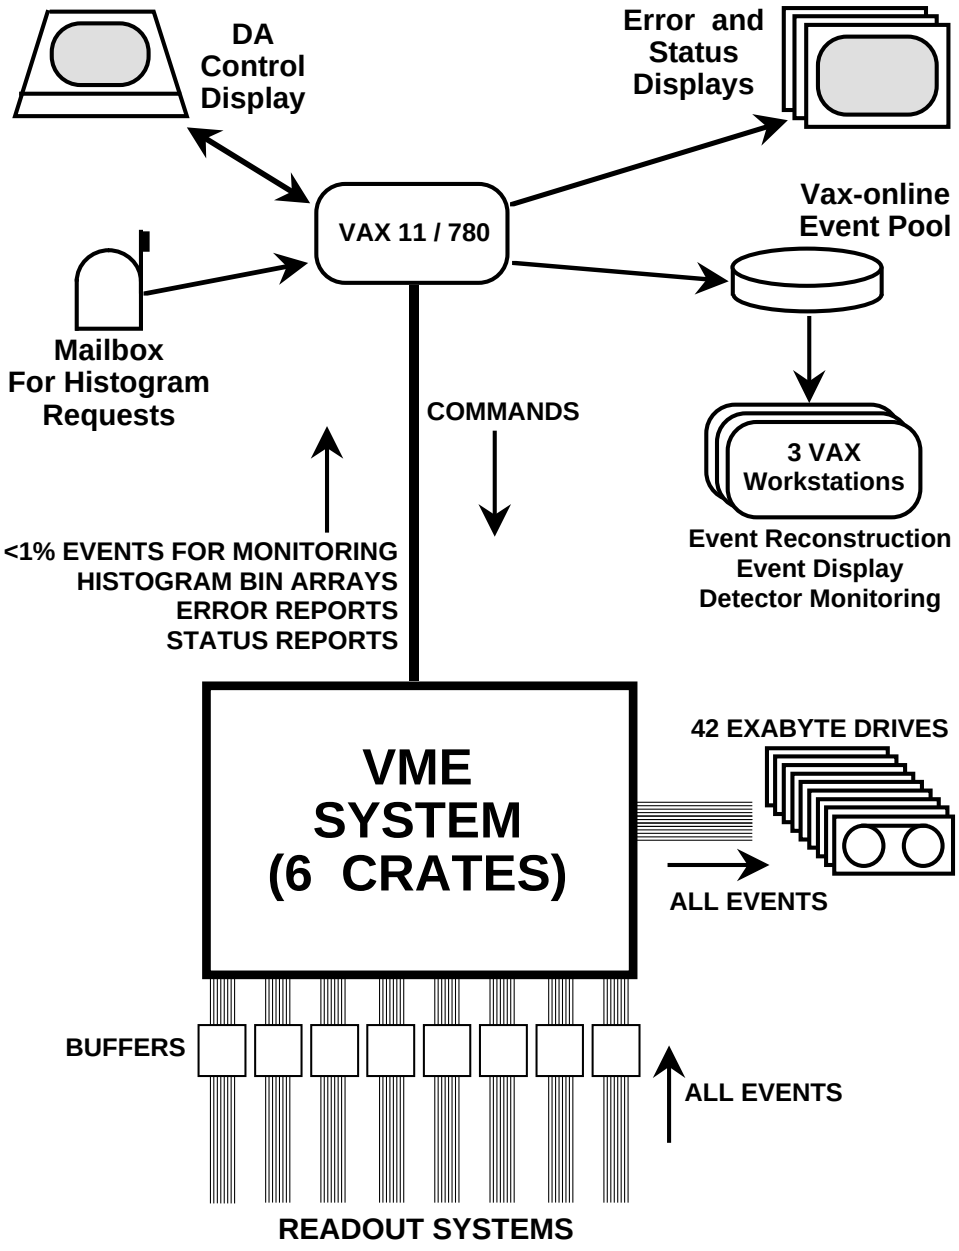

Figure 4.3: The role of the VAX in E791's DA system. (Source: [Am93])

4.2.5 Performance

The DA system performed very well. After problems were ironed out during initial testing, there were very few failures, with the exception of the Exabyte tape drives which often required head replacement after about 2000 hours of operation.

After a 5 month period of data taking, 20 billion raw events were recorded on a total of 24,000 8 mm tapes, comprising about 50 terabytes of raw data. E791 was the first high energy physics experiment to take so much raw data over such a short period of time.

Chapter 5

Event Reconstruction and Data Reduction

5.1 Introduction

As discussed in the previous chapter, E791's philosophy was to simply record as many interactions as possible and do the search for the relatively rare events in which charm particles were produced in software, after the data had been collected. As noted, this resulted in a huge (50 terabyte) sample of raw data, and thus a massive computing effort was required to process this data and reduce the dataset to a manageable size for physics analysis.

The first step involved event reconstruction and filtering. Reconstruction refers to the process by which raw detector hit data is used to create lists of tracks (hypothesized particle trajectories) and vertices (hypothesized points at which particles are created or decay). Filtering was the first stage in reducing the size of the event sample, rejecting those events in which it was unlikely a charm particle could be iden-

tified. The reconstruction and filtering were done together, as explained in subsequent sections.

The data used in this analysis underwent another reduction step called "stripping". Like filtering, stripping involved discarding events that weren't of interest according to some predefined selection criteria, which generally take the form of greater-than or less-than "cuts" on certain physics variables. The strip also involved multiple output streams optimized for different types of analyses.

Next, the sample was "substripped". This step was designed to reduce the number of events to an amount feasible for individuals to perform physics analyses, by retaining only those events for which it is very likely a charm candidate can be found.

5.2 Computing Farms

Of all the tasks outlined in the previous section, the first, reconstruction and filtering, was by far the most CPU intensive. In fact, although data was recorded over a period of only nine months, reconstruction and filtering took approximately two and half years to complete!

This enormous task was carried out at four institutions: the University of Mississippi, Kansas State University (moved from Ohio State University), Fermilab, and the Centro Brasileiro de Pesquisas Fisicas (CBPF) in Brazil. Although a large amount of data had to be processed, the fact that each event was processed independently of all other events meant that the job was highly parallelizable, with large clusters (or "farms") of CPUs working at essentially 100% capacity. The work at Fermilab was done on IBM RS6000 and SGI farms, totaling approximately 3000 MIPS (Million Instructions Per Second). The University of Mississippi and Kansas State University each used clusters of Unix based DecStations, 2900 MIPS total at U. Miss., and 2500

Table 5.1: Number of events processed at each institution.

Institution	Events processed ($\times 10^9$)
CBPF	1.8
Fermilab	4.7
Ohio/Kansas State Univ.	6.2
Univ. of Miss.	6.4

MIPS at KSU. CBPF used ACP-II computers from Fermilab totaling 1000 MIPS. The total number of events processed by each institution is listed in Table 5.1.

Each institution ran the same filtering and reconstruction code, but each used its own software to manage the parallel processing. What follows is a brief summary of the system used at Mississippi [Br96]. Although the systems at other institutions differed slightly, the basic principles were the same. A set of Unix workstations were connected via Ethernet, and used TCP/IP protocols to communicate. The farm consisted of a few server computers, and a much larger number of rack-mounted diskless client computers. Communication between the servers and clients occurred through both the NFS (Network File System) protocol by which clients could mount disks connected to the servers and read and write files to them as if they were connected locally, and directly through TCP/IP sockets where greater throughput was required than could be supplied by disks. The servers read events from the raw data tapes, and distributed events to the least busy client computer through TCP/IP. After processing, the clients wrote the reconstructed and filtered events to server disks using the NFS protocol. Servers were responsible for transferring these disk files to tape.

The reconstruction effort did go smoothly, and all 24 000 tapes were processed between November, 1992, and June, 1994.

5.3 Filtering and Reconstruction

The reconstruction and filtering process took place in three stages. First, tracks were reconstructed using information from proportional wire chambers (PWCs), silicon microstrip detectors (SMDs), and drift chambers (DCs). Vertices were then reconstructed using these tracks. The second stage was to apply the filter selection criteria and reject those events which did not pass the filter. Those events which passed were then fully reconstructed by using Čerenkov and calorimetry information to assign particle identification probabilities and to reconstruct the showers.

Some improvements were made to the reconstruction and filtering code while the data was being processed. Because of the massive computing effort required to reconstruct the data, it was decided not to re-reconstruct data that had been processed with the older code. The original code is referred to the "Release 5" code, and the newer code as "Release 7". The numbers "5" and "7" are simply code version numbers. Approximately 75% of the data was reconstructed with the Release 5 code, and the remaining 25% with the Release 7 code. Differences between the two versions are discussed when relevant.

5.3.1 Track reconstruction

Two different algorithms were used in SMD track reconstruction. Recall that the SMD planes measured the hit position in one of three "views": X, Y, and W upstream of the target, and X, Y, and V downstream. (See Section 3.4.)

For track reconstruction in the region downstream of the target, "1-view" tracks were first reconstructed in each of the three views. This was done by first forming "seed tracks", which were all possible lines between pairs of hits in two different planes in a given view. The line was projected through other planes in its view, and

hits from these other planes were added if the track passed sufficiently close to them. Tracks were kept if they had at least 4 hits in X or Y view, or 3 hits in V. Track parameters were adjusted to minimize their χ^2 to pass through their assigned hits. These 1-view tracks were then combined with tracks from other views to form the full three dimensional tracks. This algorithm was approximately 92% efficient at finding tracks.

For beam tracking (upstream of the target) three dimensional tracks were searched for directly, by first looping over all hits in all upstream SMD planes. Track candidates were selected based on a chi-square requirement which depended on the number of hits assigned to the track. A final three dimensional fit was done using PWC information as well. The beam tracking algorithm had an efficiency of approximately 95%.

Tracks found in the SMD system were then projected into the drift chamber system. Note that drift chamber planes occur in (X, U, V) triplets (see Table 3.4). The first step in tracking was to combine hit information from each of the three planes in a triplet to calculate hit positions. Silicon tracks were matched with drift chamber tracks by comparing their projected y positions since the magnets bent the tracks in the xz plane. Tracks were first projected into D3 (see Section 3.7) and upstream hits were added to the tracks. All tracks with an acceptable χ^2 per degree of freedom were kept. The track momenta were also calculated at this time. (See Section 3.6 for a discussion on momentum calculation.)

After silicon tracks were matched to drift chamber hits, remaining drift chamber hits were used to reconstruct "drift chamber only" tracks. These were used chiefly to search for K_s^0 and Λ particles which tended to decay downstream of the SMD system and were not used in the present analysis.

5.3.2 Vertex finding

Vertices were found geometrically, by considering the chi-square per degrees of freedom (χ^2/ν) for tracks to intersect. The *primary vertex* is the point at which the beam particle interacts in the target foil, possibly producing a charm particle. *Secondary vertices* are the points at which particles produced at the primary vertex decay, producing two or more "daughter" particles.

Primary vertices were found by finding the intersection point between a beam track and two or more good SMD tracks, while constraining the vertex to be in or very near a target foil. Other good tracks were added to the primary vertex as long as χ^2/ν remained less than 10. Primary vertex reconstruction was 95% efficient for events with a beam track, and less than 2% of reconstructed events had no beam track.

All good tracks with a momentum greater than 2.0 GeV were candidates for inclusion in secondary vertices, with the exception of those tracks belonging to the primary vertex with a $\chi^2/\nu < 3.5$ with respect to the primary. A secondary vertex was created by combining the two tracks with the smallest distance of closest approach downstream of the primary vertex, and then adding tracks which do not inflate the χ^2/ν beyond 15. More than one secondary vertex was allowed per event.

5.3.3 Filter

After the tracking and vertexing described above, events were filtered before they were fully reconstructed. To pass the filter, any one of the following had to be true:

1. A secondary vertex clearly separated from the primary vertex. This requirement is discussed in detail below.

2. A reconstructed $\Lambda \to p\pi$ or $K_s^0 \to \pi^-\pi^+$ candidate which decayed upstream of the first magnet (Release 5 & 7) or in the aperture of the first magnet (Release 7 only).

3. A reconstructed $\phi \to K^+K^-$ candidate (Release 5) or a $D_s \to \phi + X$ or Pentaquark $\to p\phi\pi, pK^*K$ (Release 7).

4. Events whose reconstructed tracks had a net negative charge, a total momentum greater than 350 GeV/c, and momentum transverse to the beam, p_T less than 2 GeV/c. (Release 7) (This requirement was used in diffractive jet analyses.)

The present analysis used events selected based on the first of the above criteria; the others are not relevant here. The specific requirements of item 1 are as follows. We first define the quantity $\Delta z = z_{\rm sec} - z_{\rm pri}$ to be the longitudinal separation (in z) between the primary and secondary vertices. Also, let σ_{z_p} and σ_{z_s} denote the measurement errors on the z coordinates of the primary and secondary vertices, respectively. Assuming these errors are uncorrelated (a reasonable assumption), the error on the separation becomes $\sigma_{\Delta z} = \sqrt{\sigma_{z_p}^2 + \sigma_{z_s}^2}$. We then define the quantity

$$\text{SDZ} = \frac{\Delta z}{\sigma_{\Delta z}}. \tag{5.1}$$

So SDZ is the number of standard deviations that the secondary vertex is separated from the primary vertex in the z direction. We also define an n-prong vertex to be a decay vertex with n reconstructed tracks associated with it. Requirement 1 can now be stated precisely as follows: at least one secondary vertex must be present, and further, we require that SDZ > 6 for 2-prong vertices, or SDZ > 4 for n-prong vertices with $n \geq 3$.

Events which passed one of the above filter requirements were then completely reconstructed by using the information from the Čerenkov detector to assign particle identification probabilities as described in Section 3.8, and calorimetry information to reconstruct the showers as described in Section 3.9.

Approximately 15% of events passed the filter. 55% of those 15% were selected based on the secondary vertex SDZ requirement. The average event size increased to 5.2 kilobytes with the addition of reconstructed information (track lists, vertex lists, etc.). The filtered dataset occupied roughly 8000 Exabyte tapes.

5.3.4 Release 5 and Release 7 reconstruction code differences

In addition to the changes to the filter requirements, there were several differences between the Release 5 and 7 reconstruction code, some of which have an effect on the acceptance shapes used in the present analysis. The important differences are outlined briefly here, and detailed comparisons of results are shown in the analysis chapters where appropriate.

In Release 5, a single bend point approximation for the magnets was used in projecting tracks through the magnets, as discussed in Section 3.6. In Release 7, the effect of the full field was used instead which more accurately modeled the trajectories of the particles through the magnets. Also, a misalignment between the coordinate systems in the SMD system and drift chambers which existed in Release 5 was found and corrected in Release 7. Both these changes improved the quality of the momentum measurements of tracks.

Drift chamber resolutions were updated more frequently in Release 7 than 5, and a position dependent resolution was also employed in Release 7 to account for some loss in resolution near the center of the drift chamber where the beam pions pass

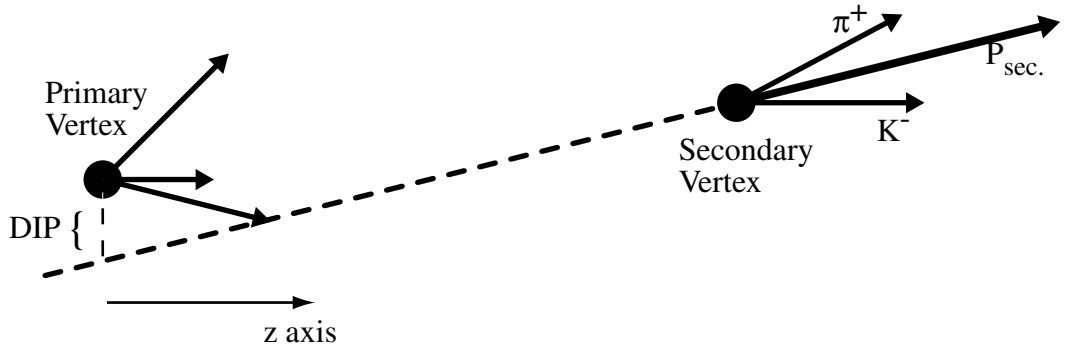

Figure 5.1: The DIP cut variable. The dashed line is the extension of the $\mathbf{P}_{\rm sec}$ vector backward toward the primary vertex.

through. These changes increased the track reconstruction efficiency. Some minor changes were made to the calculation of the χ^2/ν for tracks, and some bugs were corrected in the momentum error calculations. Also, some changes were made to the information written out to tape, but no changes which affect the present analysis.

5.4 Physics Cut Variables

Before discussing the further stages of data reduction, it will be helpful to introduce some of the physics variables used in the selection criteria. These selection criteria generally take the form of "cuts", which are typically upper or lower bounds placed on one of these variables. Events for which the variable in question does not satisfy the inequality are rejected. Most of the variables described in this section are categorized as upper bound or lower bound variables.

5.4.1 *D* Impact Parameter (DIP)

(Upper bound.) The D impact parameter (DIP) is defined as follows: First the vector momenta of the daughter particles associated with the secondary vertex are summed

to form a total momentum, $\mathbf{P}_{\text{sec}}$. A line passing through the secondary vertex with the same direction as $\mathbf{P}_{\text{sec}}$ is extended backward to the z position of the primary vertex. The DIP is defined as the distance from this line to the primary vertex at the z position of the primary vertex. (See Figure 5.1.)

The primary purpose of this cut is to ensure that the particle decaying at the secondary vertex indeed originated at the primary vertex. Also, if tracks were misassigned to the secondary vertex, or if tracks were missing from the secondary vertex (neutral particles, which don't leave a track, for instance), then the DIP would be larger, as well, so this cut also helps to ensure correct vertexing.

5.4.2 Invariant Mass (m)

(Lower and/or upper bound.) To calculate the invariant mass for a secondary vertex, it is first necessary to assign particle identification hypotheses to the daughter particles. Once this is done, their masses, obtained from the Particle Data Book [Ha02] are used, along with momentum information, to obtain the invariant mass according to the standard formula (in units with $c = 1$):

$$m = \sqrt{\left(\sum_{i=1}^{n} \sqrt{\mathbf{p}_i \cdot \mathbf{p}_i + m_i^2}\right)^2 - \left(\sum_{i=1}^{n} \mathbf{p}_i\right)^2} \tag{5.2}$$

where n is the number of daughter particles, $\mathbf{p}_i$ is the vector 3-momentum of particle i, and m_i is its mass hypothesis.

Bounds are placed on the invariant mass to exclude those events which are inconsistent with charm. The cuts are chosen to allow ample sidebands in mass histograms.

5.4.3 Proper Decay Time (τ)

(Upper bound.) Charm particles are fairly short-lived, so by making a cut on the decay time at several lifetimes, one can hope to reduce background without much loss of signal.

The proper decay time is defined as the time interval between the creation of the particle at the primary vertex, and the decay at the secondary vertex in the rest frame of the decaying particle. It is given by the formula:

$$\tau = \frac{\ell m}{|\sum_i \mathbf{p}_i|} \tag{5.3}$$

where ℓ is the distance from the primary vertex to the secondary, $\mathbf{p}_i$ is the measured momentum of the ith daughter track associated with the secondary vertex, and m is either a mass hypothesis for the decaying particle obtained from the Particle Data Book, or the reconstructed mass obtained via Equation (5.2).

5.4.4 P_T balance (PTBAL)

(Upper bound.) The transverse momentum (p_T) balance is defined as the component of $\mathbf{P}_{\text{sec}}$ (see the definition of DIP, above) transverse to the line joining the primary and secondary vertices. (See Figure 5.2.) Note that if all daughter particles are reconstructed and correctly assigned to the secondary vertex and have zero momentum errors, then this quantity should be zero.

The primary purpose of this cut is to ensure that all daughter particles tracks are reconstructed and assigned to the secondary vertex, and that no additional tracks are. In particular, this cut helps to reduce the number of secondary vertices which include neutral particles in their decays, which can't be reconstructed.

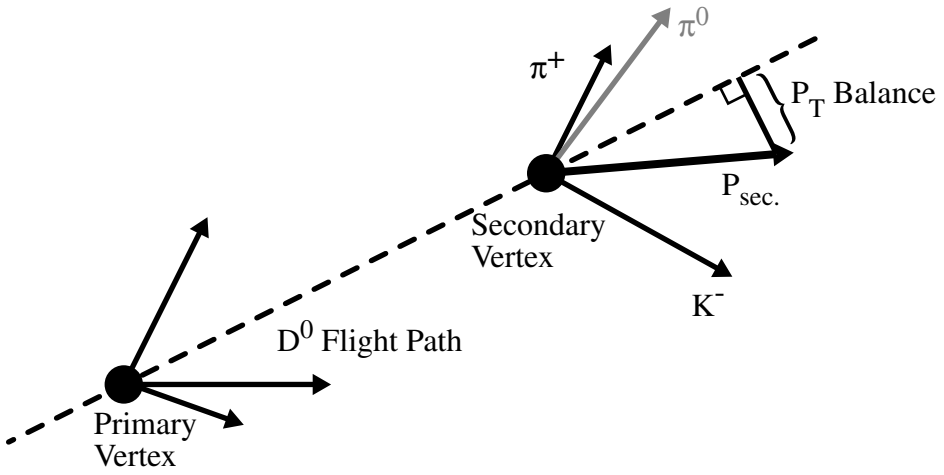

Figure 5.2: The P_T balance cut variable. Note that the π^0 momentum vector does not contribute to $\mathbf{P}_{\rm sec}$.

5.4.5 Ratio of impact parameters (RAT)

(Upper bound.) Most of the reconstructed tracks in a typical event correspond to particles produced in the primary interaction, therefore it is not unlikely that some tracks assigned to the secondary vertex are actually not daughter particles of the decaying particle, but mis-assigned primary tracks. The purpose of the ratio cut is to reduce the number of events where such a mis-assignment occurred.

The quantity is defined as follows. For any track, we define the *impact parameter* with respect to a vertex to be the distance from the vertex to the track at the z coordinate of the vertex. (This is a good approximation to the distance of closest approach and can be calculated more rapidly.) The quantity RAT is defined as the product over all secondary vertex daughter tracks of the ratio of the impact parameter with respect to the secondary vertex to the impact parameter with respect to the primary. In other words, if we let T_i refer to the ith daughter track assigned to the secondary vertex, and let V_P denote the primary vertex, and V_S denote the secondary,

and let $\mathrm{IP}(T, V)$ denote the impact parameter function defined above, then:

$$\mathrm{RAT} = \prod_i \frac{\mathrm{IP}(T_i, V_S)}{\mathrm{IP}(T_i, V_P)}. \tag{5.4}$$

5.4.6 Secondary vertex χ^2/ν

(Upper bound.) This is simply the chi-square per degree of freedom for the daughter tracks assigned to a secondary vertex to intersect. The measured quantities are the track geometry parameters (4 per track), and the unmeasured parameters are the three spatial coordinates of the secondary vertex. The constraint equations require that each track pass through the vertex (2 constraints per track). Thus there are a total of $4n - 2n - 3$ degrees of freedom for an n-prong vertex. A cut on this variable helps to reduce the number of fake secondary vertices.

5.4.7 Vertex Z coordinate

(Upper bound.) A cut is placed on this variable to ensure the primary and secondary vertices were upstream of the interaction counter. The reason has to do with triggering. If a vertex was downstream of the interaction counter, it didn't contribute to the trigger for the event, and so something else did, such as a second beam pion. Such events have a large number of tracks and contribute more to background than signal.

5.4.8 Significance of ΔZ (SDZ)

(Lower bound.) This quantity was defined in Equation (5.1) in Section 5.3.3. It is the significance of separation (in the z direction) between the primary and secondary vertices. Analysis selection criteria generally employ a tighter cut than that used in the filter. This is by far the most powerful cut at our disposal, and does a great deal

Table 5.2: Track category region definitions.

Physical region	Region number	JCATSG contribution ($= 2^{r-1}$)
Upstream of first magnet (SMDs, PWCs, D1)	1	1
Between magnets (D2)	2	2
Downstream of magnets (D3)	3	4
Upstream of calorimeters (D4)	4	8

to reduce the very high background levels present in fixed target hadroproduction experiments.

5.4.9 Target sigma (TSIG)

(Lower bound.) Some secondary vertices were found inside a target foil. Such vertices may have been the result of a secondary interaction rather than a decay. For this reason, it is desirable to eliminate those vertices. We define the quantity TSIG as:

$$\mathrm{TSIG} = \frac{\epsilon \left| z_{\mathrm{sec}} - z_{\mathrm{tgt}} \right|}{\sigma_{z_{\mathrm{sec}}}} \tag{5.5}$$

where z_{tgt} is the z coordinate of the target foil edge nearest the secondary vertex, z_{sec} is the z coordinate of the reconstructed secondary vertex, $\sigma_{z_{\mathrm{sec}}}$ is its error, and $\epsilon = 1(-1)$ if the secondary vertex is outside (inside) a target foil.

5.4.10 Track Category (JCATSG/NEWCATSG)

(Integer variable.) The variable JCATSG is a bitfield which indicates the regions of the detector which the track was observed to pass through. The regions were categorized according to Table 5.2, and the JCATSG contributions from each region where the track was observed were summed to form the variable JCATSG.

Thus, a particle which decays in the second magnet will be associated with a category 3 track, and a charged particle produced after the silicon system may be associated with a category 14 track. A real particle may have been missed by one of the regions during reconstruction if, for instance, it passed a region of low efficiency (center of the drift chamber), or if too many nearby hits created confusion that resulted in tracking inefficiency. Tracks that failed to be reconstructed downstream of the second magnet have a poorer momentum resolution than those observed in all regions, and tracks which fail to be seen beyond the first magnet have no momentum measurement associated with them at all. Drift chamber only tracks (those with JCATSG even) have poor geometrical resolution and can't be associated with a secondary vertex upstream of the interaction counter. So, for analyses which depended on momentum measurements and vertexing (such as the present one) tracks with category other than 3, 7 or 15, were generally rejected. Tracks with category 7 or 15 were generally accepted.

Tracks with category 3 were a more complicated matter. A large number of such tracks were "ghost" tracks, not corresponding to a real particle. (This can arise through combinatoric matching of hits which happen to fall in a straight line, but arise from different particles, for instance.) Many were good tracks with good measurements, however. (These facts were determined through Monte Carlo studies.) A neural net [Lo92] was trained to filter category three tracks based on chi-square, degrees of freedom, number of non-shared SMD hits, and distance from the center of the drift chamber, where efficiency is lower. To facilitate usage in older code, the routine was called NEWCATSG, and returned a value equal to JCATSG, except for tracks with JCATSG=3. For those tracks passing the neural net requirements, a value of 3 was returned, otherwise a value of 28 (indicating a "bad" track) was returned.

5.5 Stream A Strip

The 8000 tapes that resulted from the filtering and reconstruction process (see Sec. 5.3.3) were still too many for individual researchers to deal with when performing physics analyses. More data reduction steps were needed to bring the dataset to a reasonable size.

The next such step was known as "stripping", and included many physics tags, written to two separate output streams, labeled Stream A and Stream B. Stream A primarily included events with secondary vertices, making it useful for charm meson studies, as well as charm baryons which decay to charged particles, such as $\Lambda_c^+ \to pK^-\pi^+$. Stream B included tags for longer lived particles such as K_s and Λ which could be used for other charm baryon studies, as well as searches for B particles. Other tags were included as well. The Stream A strip was used in the present analysis.

Stripping took place at the same institutions as the reconstruction/filtering, with the exception of CBPF whose tapes were stripped at Fermilab. Because the data sample had already been reconstructed, stripping was much less computationally intensive than reconstruction/filtering. It was more "I/O-bound" rather than "CPU-bound", and thus stripping was done on individual workstations equipped with Exabyte EXB-10 tape stackers.

The tags used in the present analysis were the 2, 3 and 4 prong vertex tags which all required that:

- JCATSG = 3, 7, or 15 for all daughter tracks.
- SDZ > 5
- PTBAL < 1.0 GeV/c
- Total charge $= 0(\pm e)$ for even(odd) numbered prongs.

(See Section 5.4 for definitions of cut variables.) In addition, for two prong vertices, it was required that

- $m > 1.7$ GeV/c^2

where m is the invariant mass under the following hypotheses: $K\pi$, KK, $\pi\pi$ (charged particles only). If any of the hypotheses satisfied the inequality, the condition was satisfied.

Other tags saved events based on other criteria, such as two secondary vertices, dilepton events, and events with proton daughters, but they were not used in the present analysis. The stream A strip data occupied approximately 2400 Exabyte tapes.

5.6 Kansas State Substrip

Even 2400 tapes were quite a lot to deal with. Fortunately some of the physicists who did the first analyses performed another data reduction stage before applying their own final analysis cuts. This next stage of data reduction, referred to as a "substrip" was designed to be of general usefulness to many common types of meson analyses. One such substrip was begun at Ohio State University, and moved to Kansas State University along with their physics team and reconstruction farms. Another was begun at Princeton University and moved to the University of South Carolina. Because KSU was a reconstruction institution, they had earlier access to the strip data, and thus their substrip was completed first. Since it was just as inclusive as the other, it was decided not to complete the Princeton/South Carolina substrip. The "Kansas State Substrip" as it has become known, was the starting point for a substantial fraction of the E791 analyses, including the present one.

Table 5.3: Cuts used in the Kansas State Substrip.

Prongs	2	3	4
$Z_{\rm prim}$ (cm)	< -0.35	< -0.35	< -0.35
Total charge	0	$\pm e$	0
m (GeV/c^2)	> 1.7	> 1.7	> 1.7
PTBAL (GeV/c)	< 0.40	< 0.35	< 0.45
SDZ	> 8	> 8	> 7
τ (ps)	< 5.0	< 5.0	< 4.0
JCATSG	3, 7, 15	3, 7, 15	3, 7, 15
DIP (cm)	—	< 0.01	< 0.012

The KSU substrip was designed to keep events with a 2, 3, or 4 prong vertex consistent with a D^0, $D^{\pm}$, $D_s^{\pm}$, or $\Lambda_c^{\pm}$. The cuts are summarized in Table 5.3. The mass is computed for several hypotheses of daughter particles, the heaviest being K^-K^+, $K^-K^+\pi^+$, $K^-K^+\pi^+\pi^+$ for 2, 3, and 4 prongs respectively. Since the lower bound cut is the same for all hypotheses, only the heaviest daughter hypotheses are relevant as far as the cut is concerned. For the proper decay time cut, the mass was assumed to be that of the D^0, $D^{\pm}$, and D^0 for the 2, 3, and 4 prong vertices respectively.

Also, when writing the KSU substrip tapes, the raw data (detector hit positions, etc.) were removed from the event records, and only the reconstruction information was retained. The resulting data sample fit on 33 Exabyte tapes, copies of which were made for several participating institutions for physics analyses.

Chapter 6

$D^{*\pm}$ Sample and Monte Carlo

6.1 Decay channels

The first step in collecting a data sample for analysis is to decide which decay modes to use in the search for the particle of interest. Since accurate momentum measurements of all the daughter particles were needed to accurately determine the momentum distribution of the $D^{*\pm}$ particles, any decays which had neutral particles in their final states were excluded. This follows from the simple fact that momentum measurements are done by measuring the angle through which tracks are bent by the magnets, and neutral particles don't leave tracks. Also, it was pointless to search in modes which are heavily suppressed; only modes with substantial branching ratios were considered. These two considerations alone determined the decay modes used in this analysis.

The D^{*+} decays predominantly via the strong interaction into either $D^{+}\pi^{0}$ or $D^{0}\pi^{+}$ [Ha02, p. 623]. Since the π^{0} is a neutral particle, we considered only the decay:

$$D^{*+} \rightarrow D^{0}\pi^{+} \tag{6.1}$$

It is important to note at this point that unless otherwise indicated, charge conjugate modes are implied as being included in any decay written. Thus, (6.1) is understood to include $D^{*-} \to \overline{D}^0 \pi^-$ as well.

Of course, the D^0 is neutral, but its short lifetime generally resulted in decays upstream of the E791 tracking system. There are an almost overwhelming number of decay modes which have been observed for the D^0 [Ha02, p. 607]. They can be divided broadly into semi-leptonic modes, which contain leptons and hadrons in their final states, and hadronic modes which contain only hadrons in their final states. The semi-leptonic modes all include neutrinos in their final states, and could thus be eliminated. Again, hadronic modes with neutral particles such as π^0's or K^0's were eliminated. The only decays remaining, which are not suppressed by the Cabibbo rules [Pe87, p. 233], are those involving K^- and charged pions in their final states. This leaves only two decay modes:

$$D^0 \to K^- \pi^+ \tag{6.2}$$

and

$$D^0 \to K^- \pi^+ \pi^+ \pi^-. \tag{6.3}$$

Both modes were used in this analysis.

6.2 Decay Q-value

In any data sample containing decays one wishes to study, there are both "signal" events which correspond to the true particle, and "background" events which arise most often from random track combinations which happen to satisfy the selection criteria. To examine such a sample, one generally creates a *mass distribution*. That

is, for each candidate decay, the invariant mass is calculated by Equation (5.2), and the results are plotted in a histogram. The signal events will give rise to a peak near the particle's true mass[1], and the background events will generally exhibit a much broader shape.

In D^* studies, it is more common to deal with Q-value or mass difference distributions rather than straight mass distributions. First, we shall define Q, then the reasons for its use will be explained.

From the Particle Data Book [Ha02, p. 623], the difference in mass between the D^{*+} and the D^0 is $\Delta m = 145.42$ MeV/c^2, while the charged pion mass [Ha02, p. 440] is $m_{\pi^+} = 139.57$ MeV/c^2. Thus the phase space available for the decay shown in Equation (6.1) is a mere $\Delta m - m_{\pi^+} = 5.85$ MeV/c^2. This means the π^+ resulting from that decay will have a relatively small momentum. For this reason, it is referred to as the "slow pion". Contrast this with the decay (6.3), which has a phase space of 952 MeV/c^2, so each of the 4 final particles has an average momentum of 238 MeV/c. Thus the slow pion indeed generally has a much lower momentum than the other daughter particles.

We define the Q-value for the decay (6.1) by the following equation:

$$Q = m(D^{*+}) - m(D^0) - m(\pi^+) \tag{6.4}$$

where $m(D^{*+})$ and $m(D^0)$ refer to the reconstructed masses of the D^{*+} and D^0, respectively, and $m(\pi^+)$ is the PDB charged pion mass (139.57 MeV/c^2).

The advantage in using Q instead of the mass is the much greater experimental resolution in Q rather than m. The reason, put simply, is that the contributions of

[1]Measurement errors give rise to deviations from the true mass: statistical errors give the peak a finite width, and systematic errors can result in a shift in the mean from the true value as well as contributing to the width.

momentum measurement errors of the D^0 daughter particles tend to approximately cancel in the difference $m(D^{*+}) - m(D^0)$, and thus the error in Q is usually dominated by the momentum measurement errors (both magnitude and direction) in the slow pion.

To see how this approximate cancellation works, consider the general expression for invariant mass of a decaying particle (in units with $c = 1$):

$$m = \sqrt{E^2 - P^2} \tag{6.5}$$

here

$$E = \sum_i E_i = \sum_i \sqrt{\mathbf{p}_i \cdot \mathbf{p}_i + m_i^2} \tag{6.6}$$

where E_i is the energy of the ith daughter particle, $\mathbf{p}_i$ is its vector momentum, and m_i is its mass. Also,

$$P^2 = \left(\sum_i \mathbf{p}_i \right)^2 . \tag{6.7}$$

Let p_{ij} denote the jth component of $\mathbf{p}_i$, and let P_j denote the jth component of the decaying particle's momentum, i.e.

$$P_j = \sum_i p_{ij} \tag{6.8}$$

Now, consider the effect of a measurement error δp_{ij} on p_{ij}. Denote the resulting change in the calculated mass m by δm_{ij}, given by

$$\delta m_{ij} \approx \frac{\partial m}{\partial p_{ij}} \, \delta p_{ij} \tag{6.9}$$

Carrying out the partial derivative, we obtain

$$\delta m_{ij} \approx \frac{1}{m}\left(\frac{E}{E_i}p_{ij} - P_j\right)\delta p_{ij} \tag{6.10}$$

Equation (6.10) is valid for any decaying particle. In particular, let m refer to the D^{*+} mass. The daughter particles are the slow pion and the daughters of the D^0 which is created in the D^* decay. Let i refer to one of the D^0 daughters. Then Equation (6.10) becomes:

$$\delta m_{D^*ij} \approx \frac{1}{m_{D^*}}\left(\frac{E_{D^*}}{E_i}p_{ij} - P_{D^*j}\right)\delta p_{ij} \tag{6.11}$$

where we have added a subscript to identify the decaying particle in the total mass, momentum, and energy variables.

Now, let $\Delta m = m(D^{*+}) - m(D^0)$ be the difference in the reconstructed masses, and so,

$$\delta\Delta m_{ij} \approx \frac{\partial\Delta m}{\partial p_{ij}}\,\delta p_{ij} = \left[\frac{1}{E_i}\left(\frac{E_{D^*}}{m_{D^*}} - \frac{E_{D^0}}{m_{D^0}}\right)p_{ij} - \left(\frac{P_{D^*j}}{m_{D^*}} - \frac{P_{D^0j}}{m_{D^0}}\right)\right]\delta p_{ij} \tag{6.12}$$

Now, $m(D^*) - m(D^0) \ll m(D^*)$, and as the slow pion carries away only a relatively small amount of energy then $E_{D^*} - E_{D^0} \ll E_{D^*}$, and $P_{D^*j} - P_{D^0j} \ll P_{D^*j}$, as well. Thus the difference terms in parentheses in Equation (6.12) are small compared with the corresponding terms in Equation (6.11). So we would expect $\delta\Delta m_{ij} \ll \delta m_{D^*ij}$. The overall error on Δm, however, contains contributions from the slow pion momentum measurements, but it should still be much less than the overall error on the D^* mass. Since Q differs from Δm by a constant, then its error is the same as that on Δm.

This improvement in resolution resulting from the use of Q rather than m is

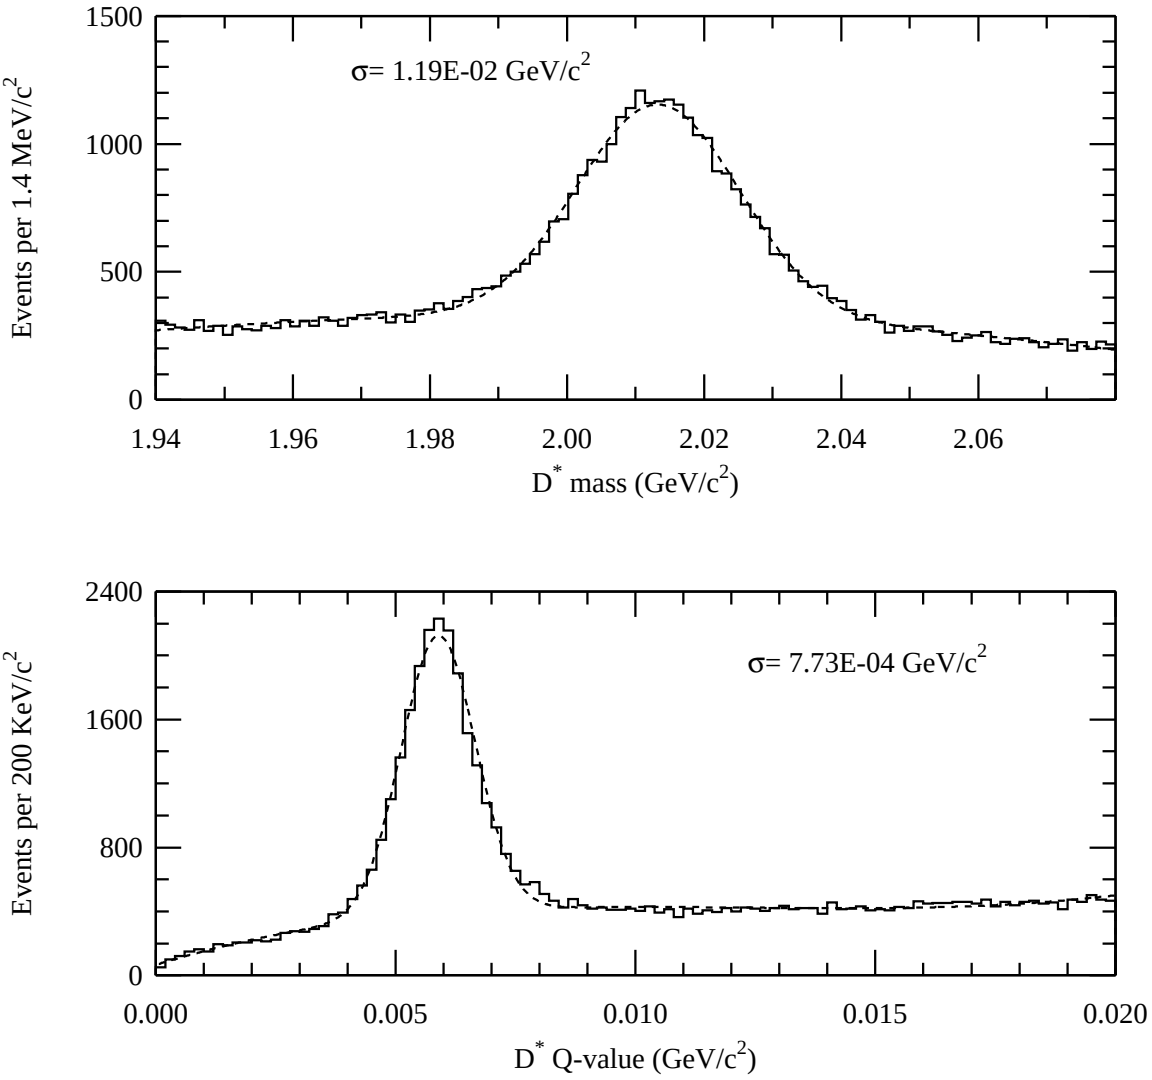

Figure 6.1: Comparison of widths of $D^{*\pm}$ mass and Q-value distributions. Note the difference in horizontal scale in the two histograms.

illustrated quite dramatically in Figure 6.1, which shows both mass and Q-value distributions for the D^* sample used in this analysis. (The selection criteria used to obtain this sample are described later in this chapter.) The fits shown are for illustrative purposes only, and assume a Gaussian shape for the signal, and a third order polynomial for the background. The signal widths displayed show more than an order of magnitude difference in resolution. Note that the large tick marks in the top plot are 20 MeV apart, the full width of the lower histogram!

Because the Q-value has a much narrower signal distribution than the mass, there are many fewer background events under the signal. This allows for a much more accurate estimation of the number of signal events in a sample. Since this analysis hinges on such estimations, Q-value histograms are used throughout.

6.3 Signal optimization

Once the decay modes are chosen for an analysis, it is then necessary to establish selection criteria which maximize the statistical significance of the signal.

For very small samples, Monte Carlo methods are typically used to determine selection criteria. This prevents bias due to tuning cut variables on statistical fluctuations, and artificially enhancing the size of the signal. For samples which are very large compared to the number of cut variables, however, it is more common to tune cuts on the data directly. This way, one is not biased by imperfections in the Monte Carlo model. The latter approach was used in this analysis.

To obtain a sample of D^* candidates a first-order filter was used on the 33 Kansas State Substrip tapes. (See Section 5.6.) The filter worked as follows. Firstly, D^0 candidates were formed by examining two and four prong vertices from the vertex

list.[2] The track with the greatest Čerenkov probability to be a kaon was assumed to be such, and the other tracks were assumed to be pions. The invariant mass was calculated, and required to fall within the range $1.77 \leq m \leq 1.97$, measured in GeV/c^2. Next, tracks not used in the D^0 vertex were examined as slow pion candidates. The charge was required to be opposite to that of the kaon candidate (which must be true in Cabibbo-favored D^0 decays). The invariant mass and Q-value for the D^* candidate was computed and the Q-value was required to be less than 21 MeV/c^2. All events containing at least one D^* candidate were saved to disk. The resulting file was about 500 Mb long and contained 2.3×10^5 events. (This file included the 3-prong candidates discussed in the previous footnote which were later discarded.)

The statistical significance of the signal is defined as the measured number of signal events divided by the error on that measurement. Let S denote the number of measured signal events. If there were no background at all, the error on S would simply be $\sqrt{S}$. In a realistic situation, we must subtract the number of background events B in the signal region from a total count of events in the signal region to obtain S. Adding the errors due to these two components in quadrature, we obtain an expression $\sqrt{S+B}$ for the error on S. So, the task of signal optimization is one of maximizing the quantity

$$\frac{S}{\sqrt{S+B}}.$$

A rapid method of computing $S/\sqrt{S+B}$ was needed. It was assumed for this purpose that the shape of the background in the D^* Q-value distribution was independent

[2]Three prong vertices were also examined both as $K\pi$, and $K3\pi$ candidates, by discarding one track or adding another track from the event, respectively, and revertexing. Although a number of good candidates were found this way, the background levels were sufficiently higher as to not enhance the overall statistical significance of the signal, even with further optimization. As a result, only two and four prong vertices from the vertex list were used in this analysis.

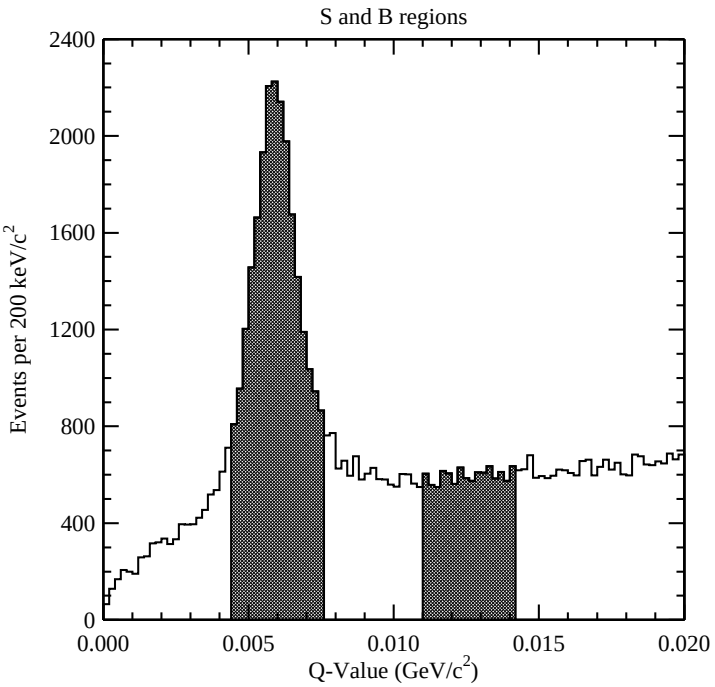

Figure 6.2: Signal and background regions used in optimization.

of the choice of cuts. This was not strictly true, but a good enough approximation for optimization purposes. A signal region was defined which included most of the signal peak (4.4 MeV/$c < Q <$ 7.6 MeV/c), and a background region was chosen to the right of the signal. The width of the background region was chosen to contain the same number of events as a fitted background shape integrated under the signal region. Thus for any D^* sample under consideration, B was simply the number of events in this background region, and $S + B$ was the number of events in the signal region. These two quantities were used to compute $S/\sqrt{S+B}$. The signal and background regions chosen are shown in Figure 6.2.

The strongly decaying $D^{*\pm}$ has a lifetime far too short to produce a secondary vertex which could be distinguished from the primary, so most of the cuts examined were based on the D^0 vertex. Plots were made of $S/\sqrt{S+B}$ as functions of cut values of various cut variables. These are shown the modes $D^0 \to K^-\pi^+$ and $D^0 \to$

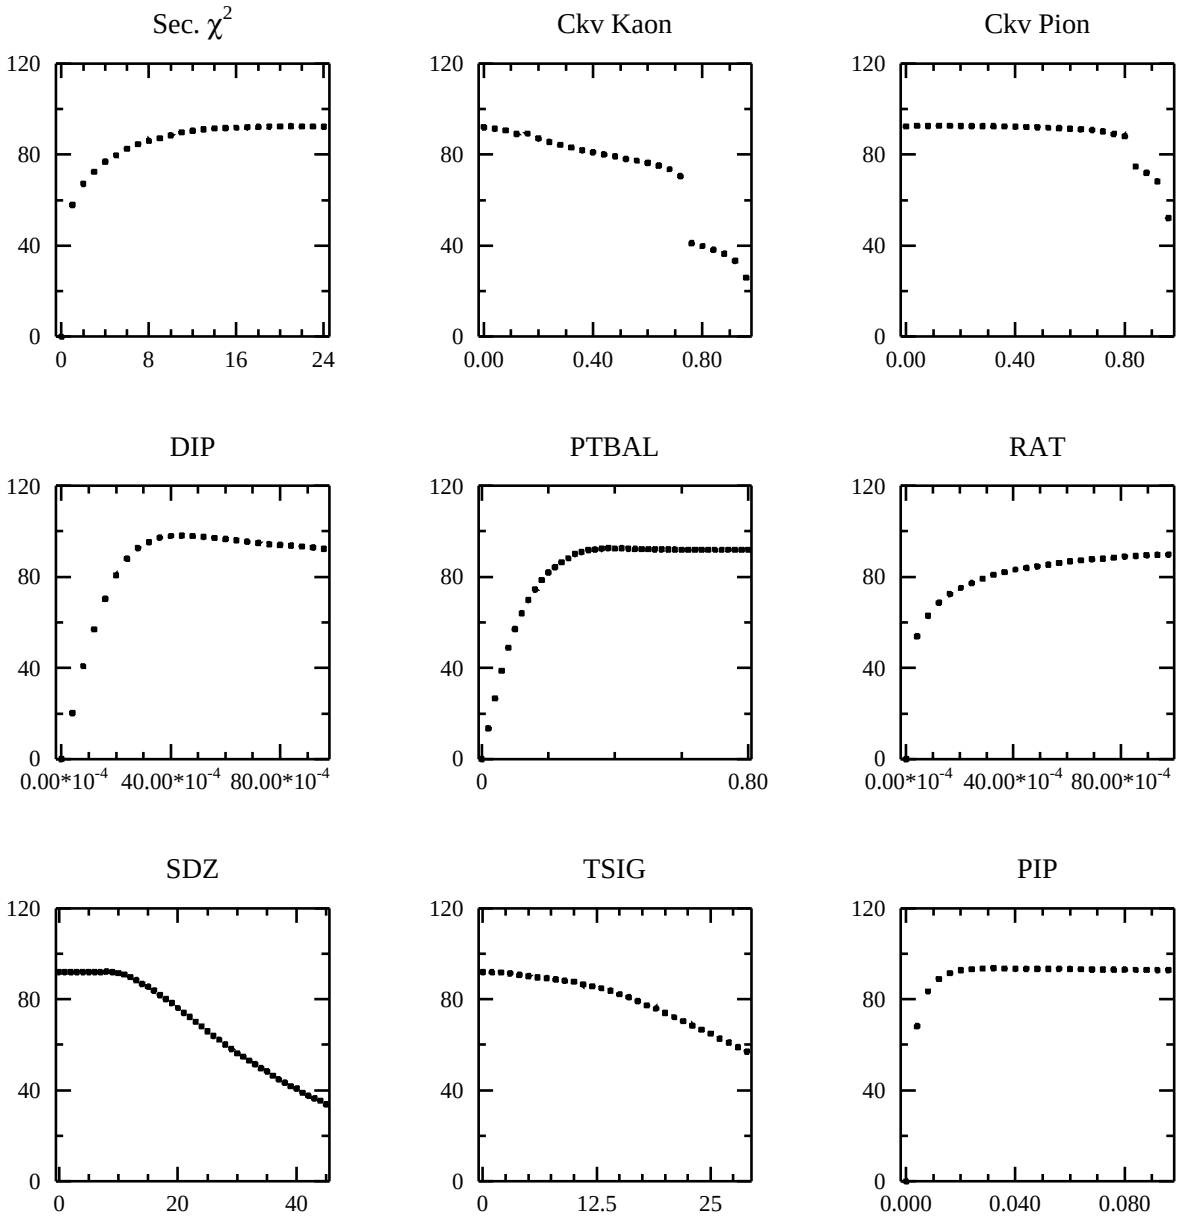

Figure 6.3: $S/\sqrt{S+B}$ of $D^{*\pm}$ vs. cut value for cut variables for $(K\pi)\pi$ sample.

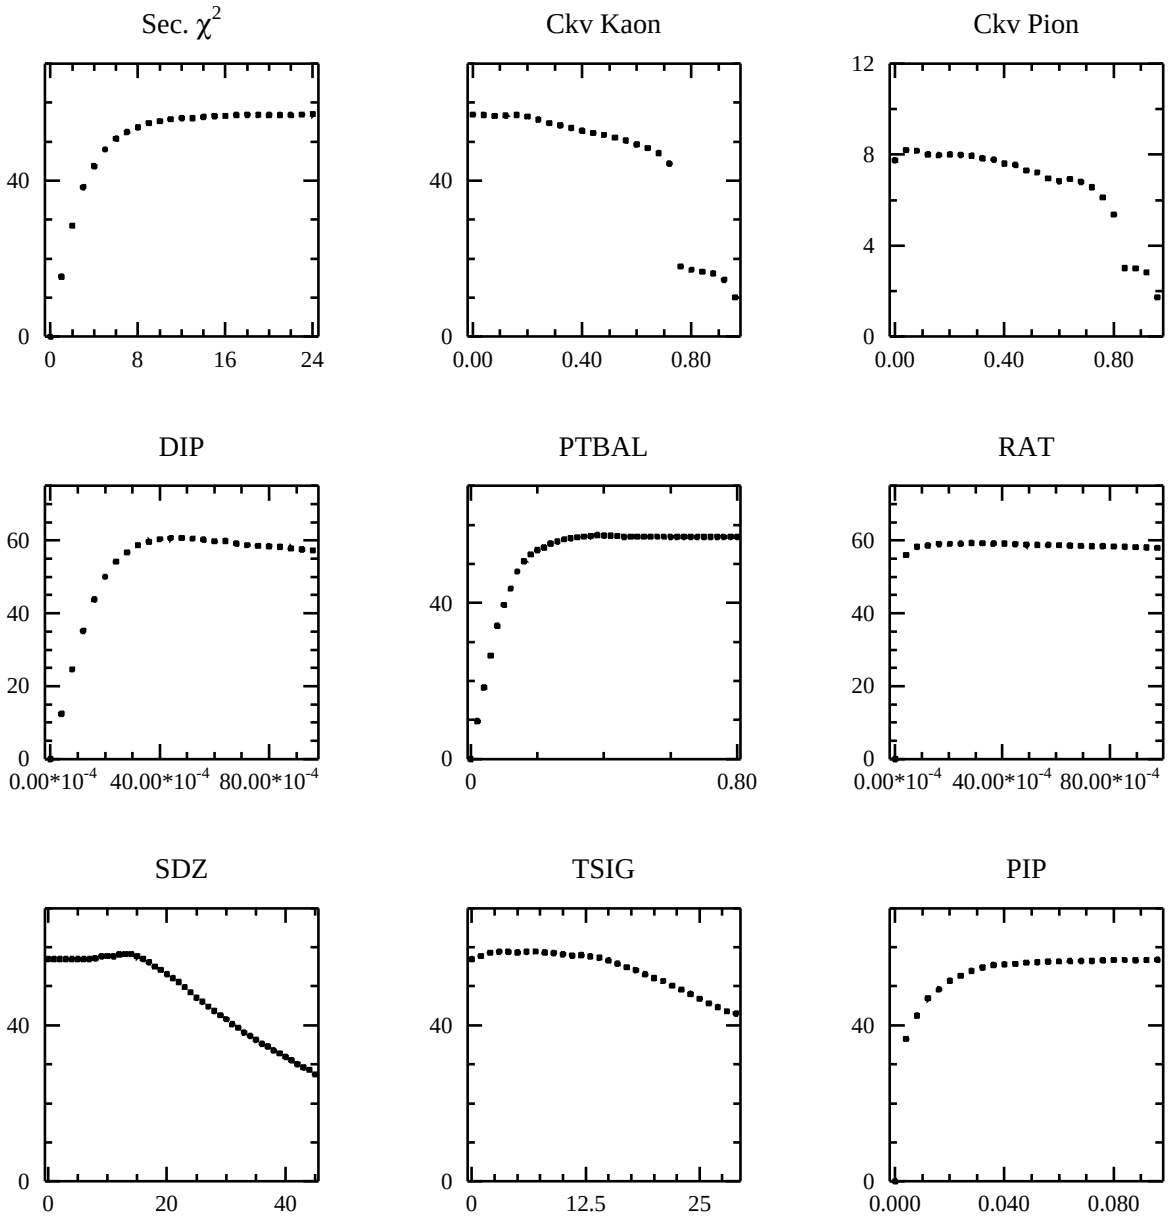

Figure 6.4: $S/\sqrt{S+B}$ of $D^{*\pm}$ vs. cut value for cut variables for $(K3\pi)\pi$ sample.

$K^-\pi^+\pi^+\pi^-$ in Figures 6.3 and 6.4, respectively. The first eight variables shown: secondary vertex χ^2/ν, Čerenkov kaon probability, Čerenkov pion probability, DIP, P_T balance, RAT, SDZ, and TSIG are based on the candidate D^0 vertex (these cuts are defined in Section 5.4, and the last, PIP (pion impact parameter), is defined as the distance between the candidate slow pion track and the primary vertex at the z coordinate of the primary vertex. (It is an upper bound cut variable.) Despite all this, the quantity $S/\sqrt{S+B}$ was computed, for all cuts, from the $D^{*\pm}$ Q-value histograms, since that was the signal to be optimized.

Examination of the plots in Figures 6.3 and 6.4 reveals no clear peaks. This suggests that signal significance cannot be significantly enhanced by tightening the cuts on any of these variables. This is due chiefly to the fact that the $D^{*\pm}$ peak sits right at the edge of phase space and already has very little background under the signal. It also should be noted that all the $S/\sqrt{S+B}$ curves flatten out as they reach the point where cuts had already been applied (at the Kansas State substrip level). This is a very good thing, as it demonstrates that previously applied cuts were not already too tight. Very loose cuts if any needed to be applied to create a final analysis sample.

Tighter cuts on the invariant mass of the D^0 would, of course, increase the significance of the $D^{*\pm}$ signal as the above mentioned samples included ample sidebands in the D^0 mass histogram. The resolution of the D^0 is strongly dependent on the momenta of its daughter particles, and hence the width of the signal is highly correlated with x_F, one of the production variables under investigation. For this reason a variable cut was placed on the D^0 mass at the time of analysis. This is discussed in detail in Section 7.3.

To summarize, the final cuts used to create the analysis sample are listed in Tables 6.1 and 6.2. The resulting signals are shown for the two modes separately

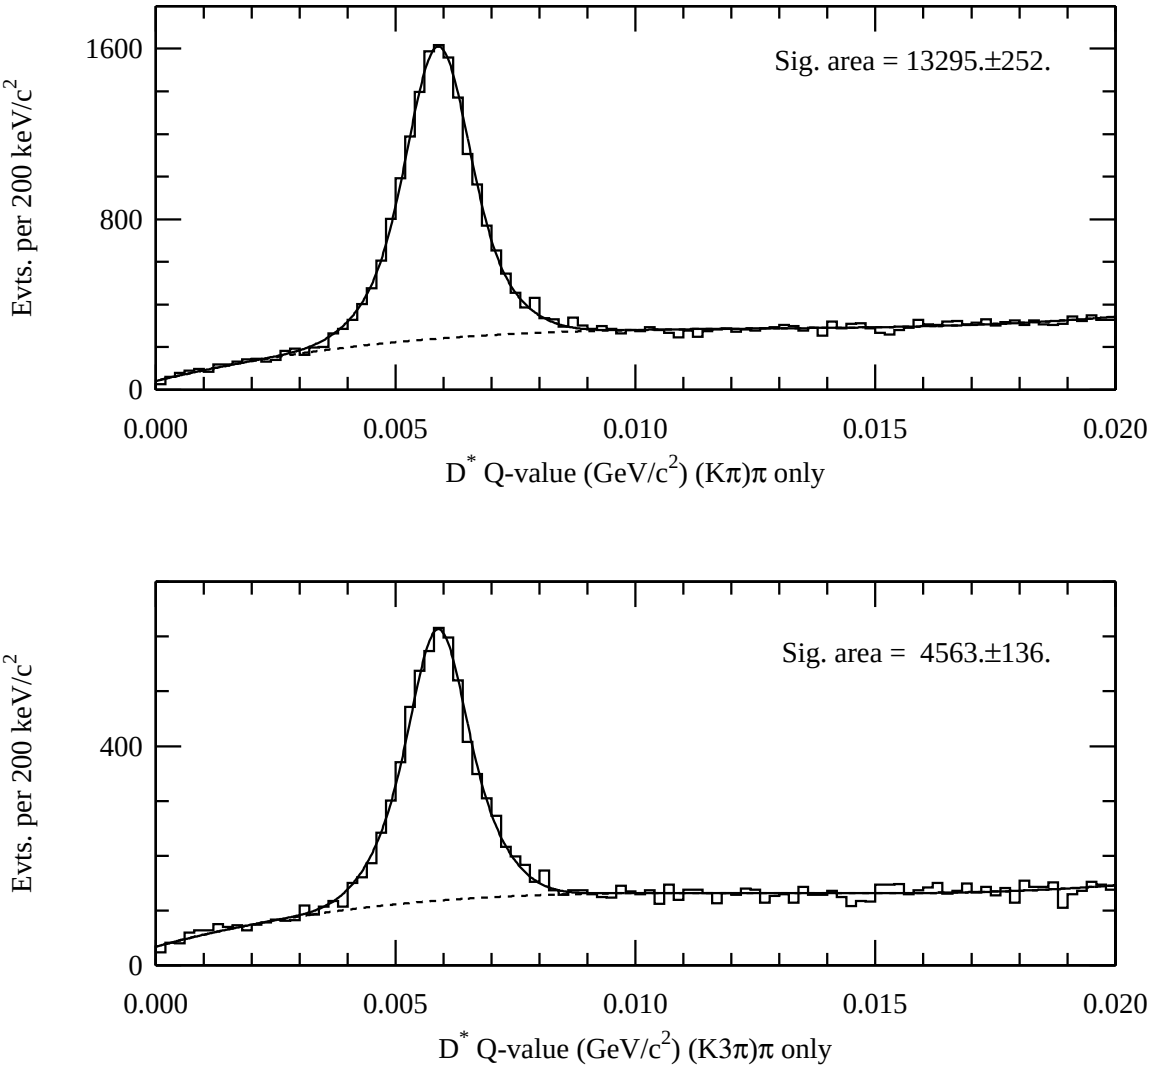

Figure 6.5: $D^{*\pm}$ samples passing final analysis cuts obtained from the two decay modes of the D^0.

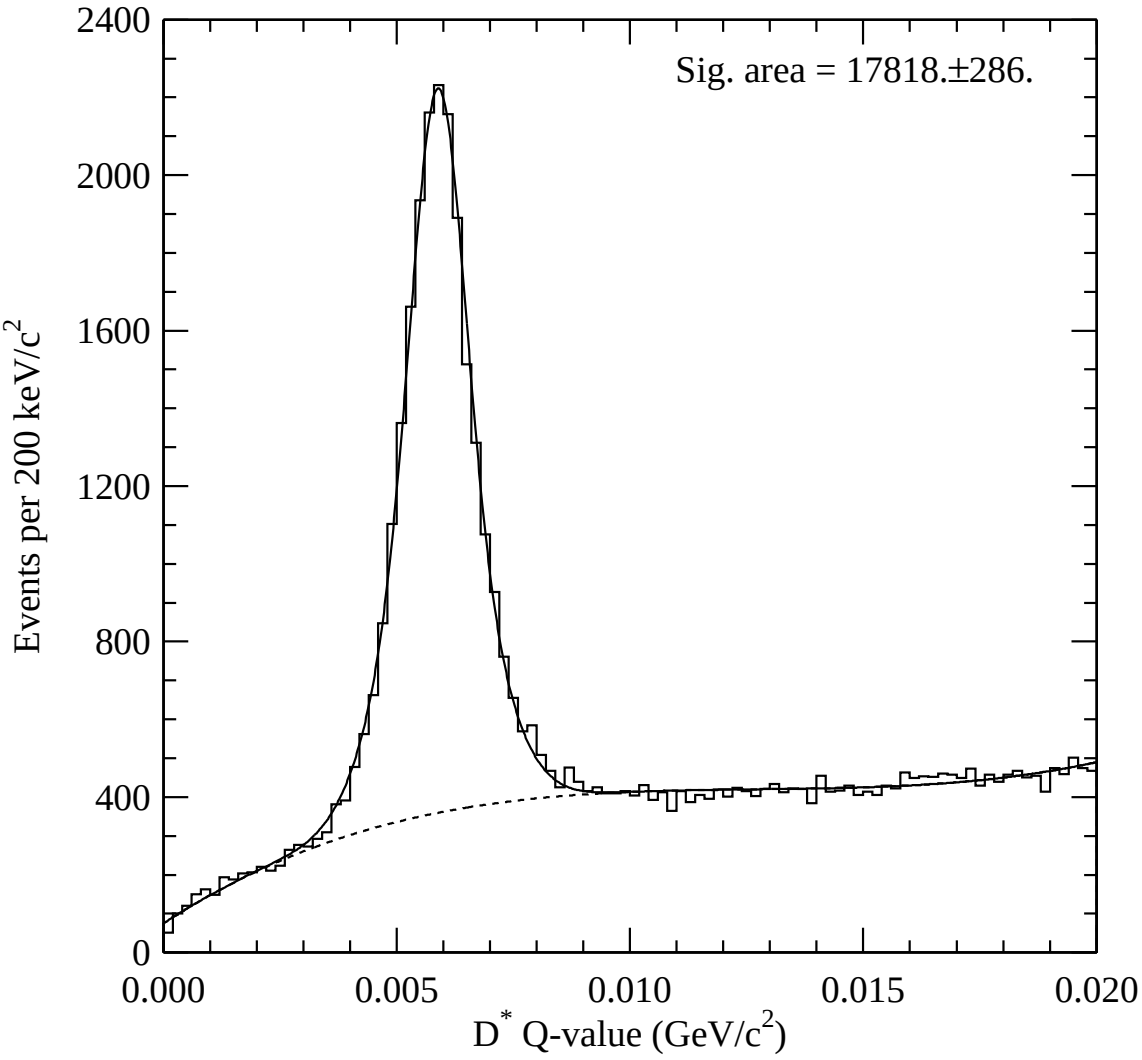

Figure 6.6: Combined $D^{*\pm}$ samples passing final analysis cuts.

Table 6.1: Final analysis cuts on D^0 candidates.

Cut Var.	$D^0 \to K^-\pi^+$	$D^0 \to K^-\pi^+\pi^+\pi^-$
$z_{\rm sec}$ (cm)	< -0.35	< -0.35
$\chi^2_{\rm sec}/\nu$	< 100	—
SDZ	> 8	> 7
P_T bal. (GeV/c)	< 0.4	< 0.45
TSIG	> 0	> 0
DIP (cm)	< 0.01	< 0.01
RAT	< 1	< 1
NEWCATSG(K, π)	$\in \{3, 7, 15\}$	$\in \{3, 7, 15\}$
$m(D^0)$ (GeV/c^2)	1.77–1.97	1.77–1.97

Table 6.2: Final analysis cuts on $D^{*\pm}$ candidates, in addition to D^0 cuts.

Cut Var.	$D^{*+} \to D^0\pi^+$
NEWCATSG(π)	$\in \{3, 7, 15\}$
$Q(D^*)$ (MeV/c^2)	≤ 21
Charge(D^*)	$\neq$ Charge(K)

in Figure 6.5 and combined in Figure 6.6. To fit the distributions, the signal was modeled as the sum of two Gaussian distributions with different areas and widths, but constrained to have the same mean. The background was modeled by a third order polynomial. The reported signal area is the sum of the areas of the two Gaussian distributions. A binned maximum likelihood fit was employed. The numbers of signal events found were $13\,295 \pm 252$ and 4563 ± 136 for the $(K\pi)\pi$ and $(K3\pi)\pi$ modes, respectively. The fit to the combined sample yields $17\,818 \pm 286$ signal events. The errors on the numbers of signal events are larger than one would naively expect due to uncertainty in the background shape under the signal. This is discussed more thoroughly in the next chapter where a more powerful fitting technique is employed which virtually eliminates this uncertainty and hence reduces these error bars.

The D^0 signal corresponding to those events with with a $D^{*\pm}$ candidate is shown in Figure 6.7.

6.4 Monte Carlo

6.4.1 Introduction

Not every $D^{*\pm}$ which is produced and decays in one of the examined modes shows up as an event in the final data sample. In fact, the vast majority do not. There are many reasons for this: one of the daughter particles might have a trajectory which takes it outside the instrumented region of the detector, detector inefficiencies might result in too few hits for a track to be reconstructed, tracks left by more than one particle might be too close to each other to be resolved, selection criteria might reject good events that happen to fall outside the required window, etc.

If these inefficiencies were uniform in the variables examined in this study, there

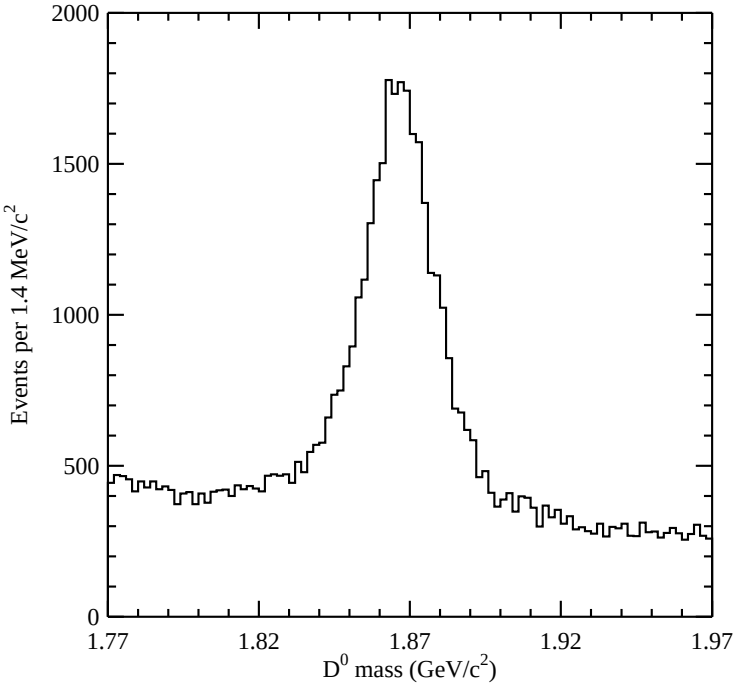

Figure 6.7: D^0 signal from final analysis sample.

would be no cause for concern, but unfortunately that was not even approximately the case. Events with a very high Feynman-x (x_F), for instance, tended to produce tracks which passed through the central regions of the the drift chambers where efficiencies were much poorer. Events of very low or negative x_F frequently escaped the detector entirely. Transverse momentum (p_T^2) efficiencies were somewhat flatter, but not flat enough to be neglected.

In a physics analysis, we are interested in obtaining true distributions arising from physics processes, not detector and reconstruction effects. We therefore model these effects using *acceptance* functions (also known as *efficiency* functions). An acceptance function is proportional to the probability that an event will be "accepted" (i.e. reach the final data sample) as a function of some variable. So, if we wish to measure a distribution of some physics variable x, we divide the observed distribution by the acceptance function of x to obtain an estimate of the true distribution.

In high energy physics, by far the most common technique for calculating acceptance functions is the *Monte Carlo* method. It involves creating a software simulation of the experiment in which large numbers of simulated events are generated according to physics models and then reconstructed. The resulting reconstructed distributions are divided by the generated distributions to obtain acceptance functions.

Three stages are involved in the creation of a Monte Carlo sample to be used in modeling acceptance: *generation*, *digitization*, and *reconstruction*.

Generation refers to the process by which the interactions of the beam particles and target are modeled in order to produce further particles. Also, the subsequent decay of such particles, as well as further interactions (elastic or inelastic) in downstream target foils or detector matter must be simulated as well. Pre-existing software libraries are often employed. E791 used the PYTHIA/JETSET package [Sj94] to model these interactions. The output of the generation phase is a list of particles which includes the particle identity, momentum components, birth and death coordinates, etc. This list is referred to as a *truth table*, since it lists the "true" quantities which do not yet have measurement errors or inefficiencies associated with them.

The next step is digitization. This is where the response of the detectors to particles is simulated. Truth table information is used to determine a particle's trajectory through the detector. Where a charged particle passes through an SMD plane, for instance, a hit is recorded on the strip through which it passed. It is at this stage where detector performance characteristics must be taken into account. These characteristics were frequently measured using real data. To measure the efficiency of an SMD plane, for instance, tracks were reconstructed using all but the plane under consideration. Good tracks were projected through that plane and it was observed whether or not a hit was recorded on the appropriate strip. This was done for numerous tracks to determine efficiencies. Also, noisy and "hot" strips were observed.

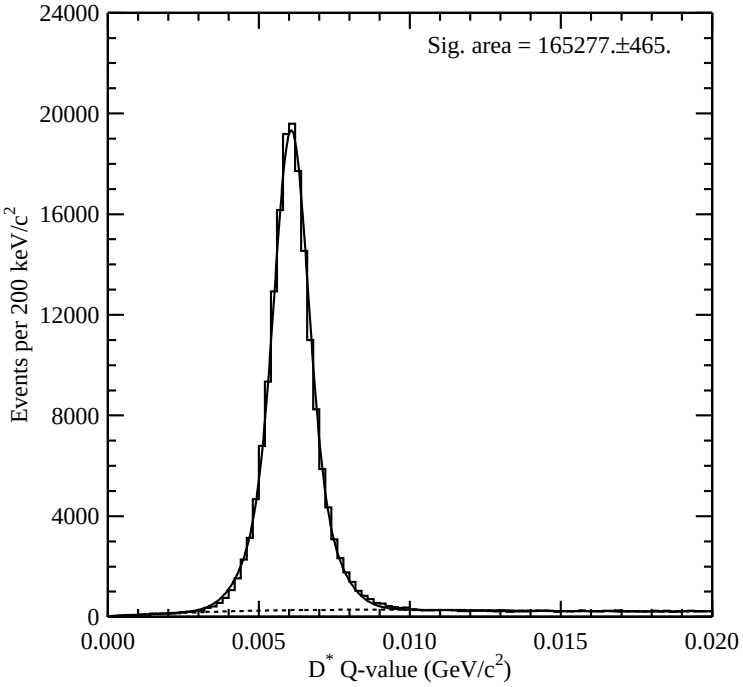

Figure 6.8: $D^{*\pm}$ signal from Monte Carlo sample.

Special muon runs were used to calibrate calorimetry systems. All this performance information is used during digitization to simulate the way the detector responds to an event. The output of the digitization phase is an event record in exactly the same format as that used by the DA system when recording real data.

Reconstruction takes place exactly as for real data. In fact, in E791, the very same code is used, ensuring that reconstruction effects are modeled correctly by the Monte Carlo.

6.4.2 Monte Carlo Sample

To generate acceptance functions, it was necessary to create a Monte Carlo sample of $D^{*\pm}$ particles decaying through the channels used in the analysis. Because statistical errors associated with acceptance functions should not dominate those in data, a $D^{*\pm}$

sample with approximately ten times the statistics of the signal observed in real data was generated. The reconstructed Monte Carlo was then passed through all the same cuts as the real data: Filter, Stream A strip, Kansas State Substrip, and final analysis strip.

This was a fairly CPU intensive task. A DEC Alpha and dual processor Pentium II computer at the University of South Carolina spent more than a month on the task. Both the release 5 and release 7 reconstruction code (see Section 5.3) were employed in the same ratio as that used in reconstructing the data. The $D^{*\pm}$ signal obtained from this Monte Carlo sample is shown in Figure 6.8. The actual acceptance functions are shown in the relevant analysis chapters.

6.4.3 Correcting Monte Carlo Distributions

The accuracy of the acceptance functions depends on how closely the Monte Carlo system models reality. The reconstruction phase is identical for both data and Monte Carlo, and the digitization depends on measurements of detector characteristics made using real data. The generator, however, depends on physics models and parameters which are determined, in part, by analyses such as the present one. But this analysis, of course, depends on accurate acceptance functions. This apparent "chicken and egg" situation is partly resolved by observing that a measurement of a differential cross section in one production variable, say x_F, is independent of the Monte Carlo distribution of that variable. The reason is simply that the acceptance function is found by dividing a reconstructed Monte Carlo distribution by a truth table distribution, so the shape of the generated distribution cancels in the ratio.

It is important to note, however, that when computing an acceptance function of one variable, we are effectively integrating over all other variables, and discrepancies

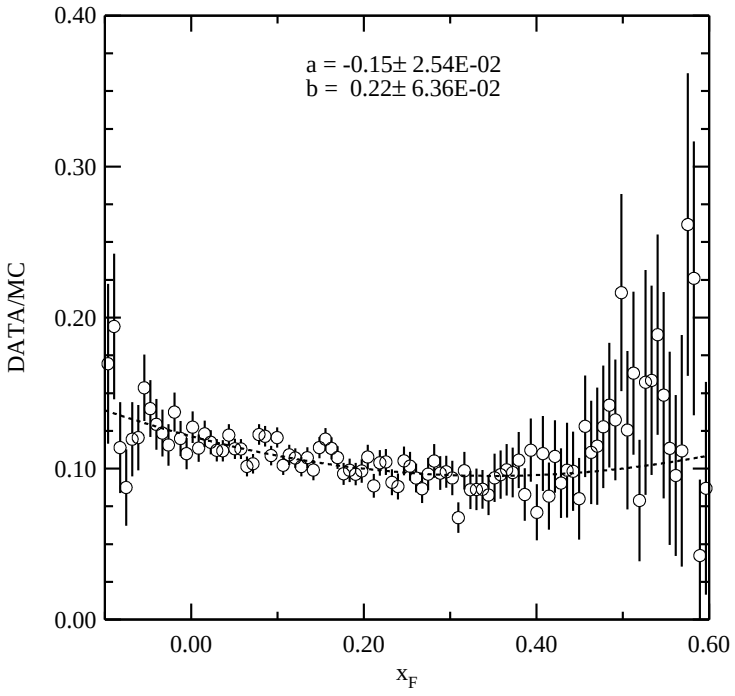

Figure 6.9: Ratio of data x_F distribution to reconstructed Monte Carlo with fit of form $ax + bx^2 + c$.

between real and simulated distributions in some of these variables may lead to inaccuracies in the acceptance function. These variables fall into two broad categories. In the first category are characteristics of the event as a whole, such as track multiplicity. If there are more high momentum tracks coming from the primary vertex in real data than are simulated in Monte Carlo, then the acceptance at high x_F will be overestimated in the Monte Carlo due to an underestimation of the amount of track "confusion" in the central regions of the detector through which high momentum tracks tend to pass, for example. The second category includes production characteristics and lifetimes of the particles of interest. There are correlations between x_F and transverse momentum squared p_T^2 in the generation, so errors in the distribution of one variable can give rise to inaccuracies in the acceptance functions of the other.

As to the first category described above, the E791 Monte Carlo program had been tuned so that relevant quantities such as track multiplicity were correct. The production characteristics of individual particles had some discrepancies, however. Of particular importance are the momentum distributions of the $D^{*\pm}$ particles studied, as reflected in their x_F and p_T^2 distributions.

As will be shown below, substantial differences exist between the data and Monte Carlo distributions, particularly in p_T^2. Since this analysis involves one dimensional acceptance corrections in both x_F and p_T^2, it is necessary to correct these distributions in Monte Carlo. There are two possible ways to do this. The first would be to tune some parameters in the generator in order to attempt to make the distributions agree. The second would be to apply a weighting factor to events in the Monte Carlo sample after the fact, to correct for the differences. The latter approach was used in this analysis for two reasons. Firstly, tuning the generator would involve regenerating the entire Monte Carlo sample, possibly several times, and as explained earlier, it was a very CPU intensive task. Secondly, changing generation parameters could conceivably

disturb other event characteristics which had already been tuned to agree with data. For this reason, a weighting scheme was used on the existing sample.

To examine the quality of agreement between real and Monte Carlo distributions, it was assumed that the ratio of reconstructed data to reconstructed Monte Carlo is a good approximation to the ratio of the true distributions to the generated distributions. It was in fact the best approximation available; acceptance-correcting the data and comparing to the generated distribution would lead to the exact same ratio. The only complication was that our samples contain both signal and background events, and only signal events should be used in the comparison. To estimate the signal-only distributions, the technique of *background subtraction* was employed. A background-only distribution was created using only events which fall within a Q-value window well away from the signal peak. A signal plus background distribution was created by making tight cuts around the signal peak and using only events falling inside that Q-value window. The former distribution was subtracted from the latter, after normalizing it to have the same number of events as the estimated number of background events in the signal plus background distribution. This estimation was obtained by fitting the signal and background of the Q-value histogram, and integrating the background function over the signal region chosen.

The ratio of the data x_F distribution to reconstructed Monte Carlo is shown in Figure 6.9. As can be seen, there is a bit more data at very low and high x_F than is modeled by the Monte Carlo. The discrepancy is not too bad. A quadratic polynomial was fit to the distribution and the linear and quadratic coefficients (a and b respectively) are shown on the plot. Both are several standard deviations from zero, so some correction was needed. For this reason, a "weighting curve" of

$$w = 9.426(0.12 - 0.15x_F + 0.215x_F^2) \tag{6.13}$$

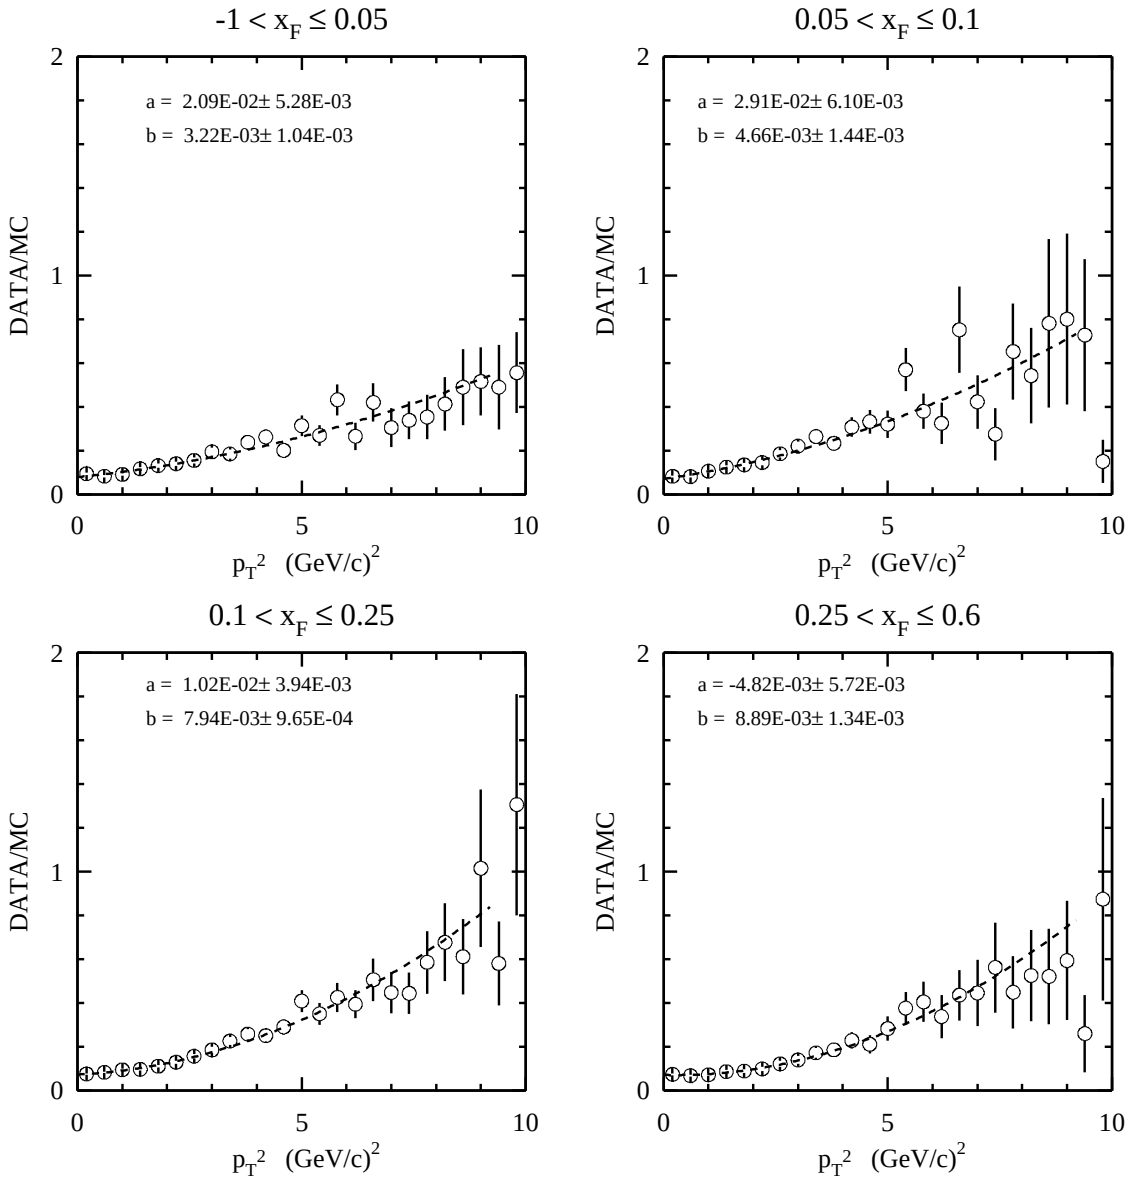

Figure 6.10: Ratios of data p_T^2 distributions to reconstructed Monte Carlo in four x_F regions with fits of form $ax + bx^2 + c$.

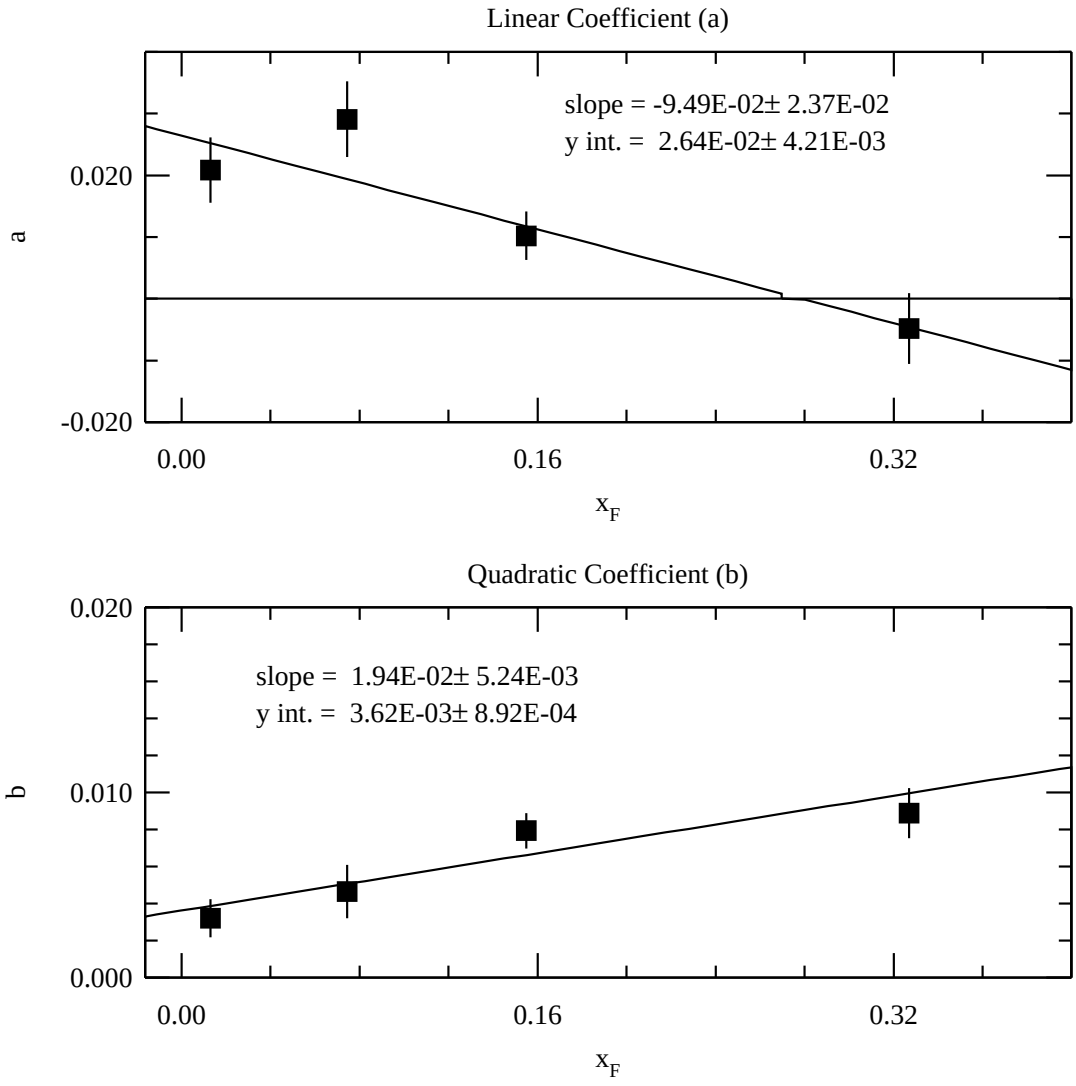

Figure 6.11: Straight line fits of linear and quadratic coefficients of p_T^2 correction curves as functions of x_F.

was employed. All Monte Carlo events were weighted according to the above equation, except where generating acceptance functions in x_F.

The x_F discrepancy was found to be fairly small, but the situation is quite different for p_T^2. Ratios of data to reconstructed MC are plotted in Figure 6.10. Clearly the data shows a much stiffer p_T^2 distribution than the Monte Carlo. Because it was found to be much more significant than the x_F ratio, the p_T^2 ratio was computed in four different regions of x_F to investigate a second-order dependence of the correction curve on x_F. As can be clearly seen, there is such a dependence. To obtain a weighting curve, the parameters themselves had to be fit as functions of x_F. Straight lines were used because a higher order curve would have been inappropriate for only four points. Smaller bins in x_F would have yielded more points, but the error bars would have become too large to be useful. Although a line doesn't fit really well, it is a step better than assuming no dependence on x_F at all. These fits are shown in Figure 6.11. The fits yielded a weighting curve of:

$$\begin{aligned} w &= 3.12[7.2 \times 10^{-2} + \\ &\quad (2.64 \times 10^{-}2 - 9.494 \times 10^{-}2x_F)p_T^2 + \\ &\quad (3.624 \times 10^{-}3 + 1.9385 \times 10^{-}2x_F)(p_T^2)^2]. \end{aligned} \tag{6.14}$$

Chapter 7

Differential Production Cross Sections

7.1 Kinematic Variables

As discussed in Chapter 1, one of our goals is to measure the dependence of $D^{*\pm}$ production cross sections on the kinematic variables Feynman-x (x_F) and square transverse momentum (p_T^2). The transverse momentum p_T is simply the component of the D^* momentum vector transverse to the beam particle direction. Since the beam tracking system measures this direction for us, and the D^* momentum is reconstructed from its decay tracks, calculation of p_T^2 is trivial.

x_F is slightly more complicated. Recall its definition (Eq. 1.3),

$$x_F = \frac{p^*_{\shortparallel}}{p^*_{\shortparallel\,\mathrm{max}}} \tag{7.1}$$

where $p_{\shortparallel}$ is the component of the D^* momentum vector parallel to the beam particle direction, and the asterisks indicate the quantities are evaluated in the production

center of mass frame. It is thus necessary to define precisely what is meant by the "production center of mass frame" and how is $p^*_{\parallel\,\max}$ computed.

To define the center of mass frame, it is necessary to know precisely which particles are interacting to produce the D^*. Clearly the beam pion is one such particle. A nucleon in the target material is clearly also involved. One might argue that because the nucleon is bound in an atomic nucleus that we should consider the entire nucleus as the target particle, but at the energy scales involved here the nuclear binding forces are not significant, and it makes sense to define the target as a free nucleon. One might then argue that at our energy scales it is more correct to define a parton as the target, and not the whole nucleon, but since one of the purposes of a differential cross section study is to probe the parton structure of the nucleon, the entire nucleon must be considered as the target particle. The target particle might, of course, be either a proton or neutron, but fortunately, their masses are close enough that it doesn't matter; we simply use the proton mass (938 MeV/c^2) as the target mass. The target nucleon is assumed to be at rest in the lab frame, and the beam pion is assumed to have a momentum of 500 GeV/c. This, together with a measurement of the beam particle's direction, is sufficient to precisely define the center of mass frame.

It is also possible to calculate the square total energy in the COM frame (in units with $c = 1$):

$$s = (\tilde{p}_B + \tilde{p}_T)^2 = m_p^2 + m_\pi^2 + 2m_p E_B \approx 939 \text{ GeV}^2 \tag{7.2}$$

where $\tilde{p}_B$ and $\tilde{p}_T$ are the 4-momenta of the beam and target particles respectively. The D^* will have maximum forward momentum when all other products move backward in the COM frame with total momentum equal and opposite to that of the D^*. Assuming both are relativistic, the energy is split evenly between the two, so $E_{D^*} = \sqrt{s}/2$. Thus,

$p_{D^*} = \sqrt{s/4 - m_{D^*}}$. We thus arrive at the form used to compute x_F in this analysis:

$$x_F = \frac{p^*_{\shortparallel}}{\sqrt{\frac{s}{4} - m_{D^*}}} \tag{7.3}$$

where m_{D^*} is the reconstructed mass of the D^*.

7.2 Fitting Procedure

In Section 6.4.3, background subtraction was used to create x_F and p_T^2 distributions in order to compare data and Monte Carlo. That technique was sufficient for that purpose, but for the analysis proper a more accurate approach is required. In doing a simple background subtraction, it is necessary to assume that the x_F and p_T^2 distributions in the background region are the same as those of the background events in the signal region. This is equivalent to assuming that the background shape does not depend on either x_F or p_T^2. As we shall see, that assumption is not really valid.

A more precise technique which does not depend on that assumption is to first divide the data sample into bins of x_F (or p_T^2) and create Q-value distributions for each of those bins. Then, one can fit the Q-value distributions to determine the numbers of signal events. These numbers are then plotted as functions of x_F (or p_T^2) to give the data distributions. Thus, a reliable and efficient method for fitting D^* Q-value distributions is needed.

In Section 6.3, the final data sample was fit with a double Gaussian model for the signal and a third order polynomial for the background (see Figure 6.6). Although the fit looks good to the eye, the error bar on the signal area reveals it is, in fact, not a good fit at all. We argued in Section 6.3 that the error on this number should be approximately $\sqrt{S+B}$. Integrating the background curve under the signal ($3.5 \leq$

$Q \leq 8.5$ MeV/c^2), yields roughly 8890 background events under the signal. The fit reports 17 818 signal events, so $\sqrt{S+B} \approx 163$. The reported error of 286 is about 75% larger than it should be.

The explanation for this discrepancy lies in the difficulty in ascertaining the background shape under the signal. Because the D^* Q-value signal sits right at the edge of phase space, the background shape changes rapidly under the signal. To the right it is almost described by a horizontal line; to the left, by a small shoulder much lower than the level on the right. Contrast this with a D^0 mass plot (see Figure 6.7), where a straight line fit to the background on one side could be projected under the signal and closely match the background on the other side, thus leaving very little uncertainty about the background shape under the signal. With the D^*, however, there are many reasonable ways the right and left background regions could be joined, and thus the error on S is not purely statistical, but has a large systematic component due to this uncertainty.

A background model which does not suffer from this uncertainty can be obtained directly from the data by using the technique of *event mixing*. Since the background in a D^* sample is due to combining slow pion candidates with D^0 candidates which don't originate from a common D^*, background events are modeled by combining D^0 candidates with slow pion candidates from different events. The exact same selection code is used, only the slow pion candidates and D^0 candidates must be from different events. Since it is possible to combine a D^0 candidate with slow pion candidates from several other events, the event-mixed sample can be many times larger than the actual data sample, so that errors from fits are dominated by the statistics of the data sample, not the background model. Using this technique, the only free parameter associated with the background portion of the fit is an overall normalization, which

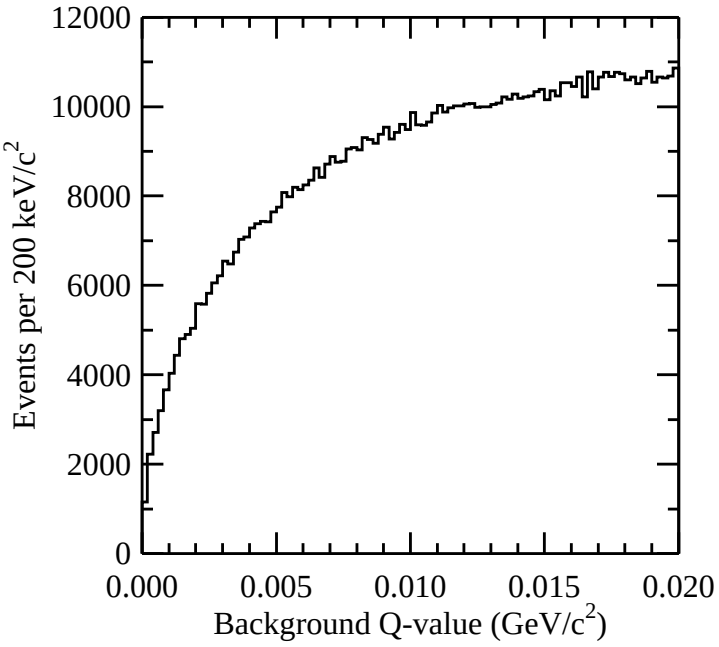

Figure 7.1: Background model obtained by event mixing.

can be well determined by the ample sideband to the right of the signal; there is no uncertainty in shape under the signal.

Figure 7.1 shows the background shape model obtained by event mixing from the entire analysis sample. Figure 7.2 shows a fit to the D^* signal using a double Gaussian for the signal (as in Figure 6.6) and event mixing to model the background. The error reported for S is 187, much closer to the expected value. Also, a visual examination of the fit suggests that the shape predicted by event mixing does indeed fit the real background very well on either side of the signal. It is therefore quite reasonable to assume it is correct under the signal, as well.

Although a double Gaussian models the signal shape for the complete data sample very well, problems were encountered during attempts to fit the sample when divided into bins of p_T^2 and x_F. In some bins of x_F a single Gaussian provided a better fit, and the double Gaussian fit was unstable. Other bins required a double Gaussian,

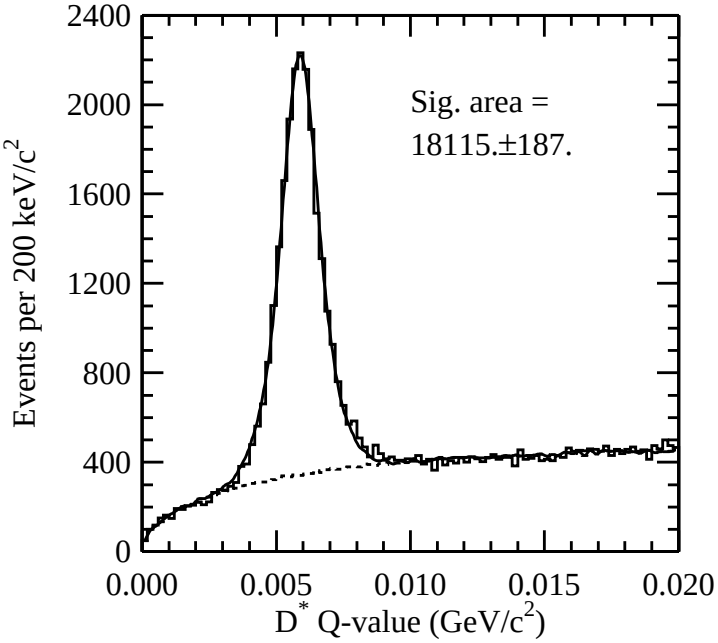

Figure 7.2: $D^{*\pm}$ analysis sample with fit. A double Gaussian was used to model the signal, and event mixing to model the background.

and some bins were not fit well by either. This resulted in some fits not converging properly, therefore having one model for the signal shape is not sufficient.

The background was fit with a shape obtained from event mixing. This is a physically meaningful hypothesis, unlike the double Gaussian fit for the signal, which is not motivated by physics, but merely by the fact that it fits the full data sample quite well. In fitting the event-mixed histogram to the background, the only parameter is a scale factor. Since we are interested in obtaining only the number of signal events, and not widths or means, we really don't need to fit the signal at all if we have a good model for the background. All that is necessary is to determine the correct scale factor for the background, and once that is done, the number of signal events can be obtained by subtracting the integral of the scaled background function over the signal region of the histogram from the total number of events in that region. To obtain the correct scale factor, we fit only in a region well away from the signal.

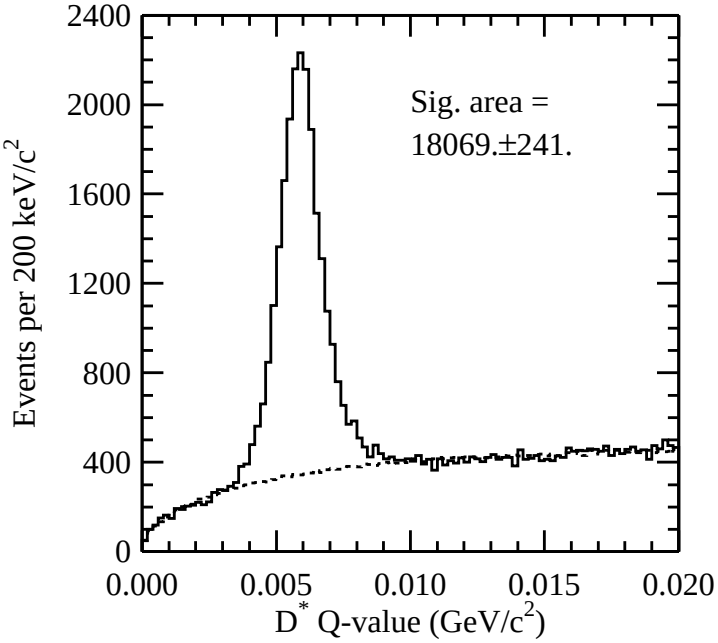

Figure 7.3: $D^{*\pm}$ analysis sample with background fit only.

As a consistency check, the analysis sample was fit using this technique, as shown in Figure 7.3. The signal region was chosen as $0 \leq Q < 12$ MeV/c^2 and the background region was chosen as $12 \leq Q < 20$ MeV/c^2. The reported yields differ by only 0.2σ, which suggests the fitting technique is sound. The error on the yield is higher (241) owing to the statistical errors obtained when fitting the background model to the histogram, but it is still substantially less than that shown in Figure 6.6.

7.3 Variable D^0 Mass Cut

In Section 6.3, it was indicated that the width of the D^0 signal is highly correlated with x_F. This is illustrated in Figure 7.4 which shows the D^0 signal in two different ranges of x_F. What would be a reasonable cut on the D^0 mass for the left-hand plot would cut into the signal on the right-hand plot. The reason for this can be found

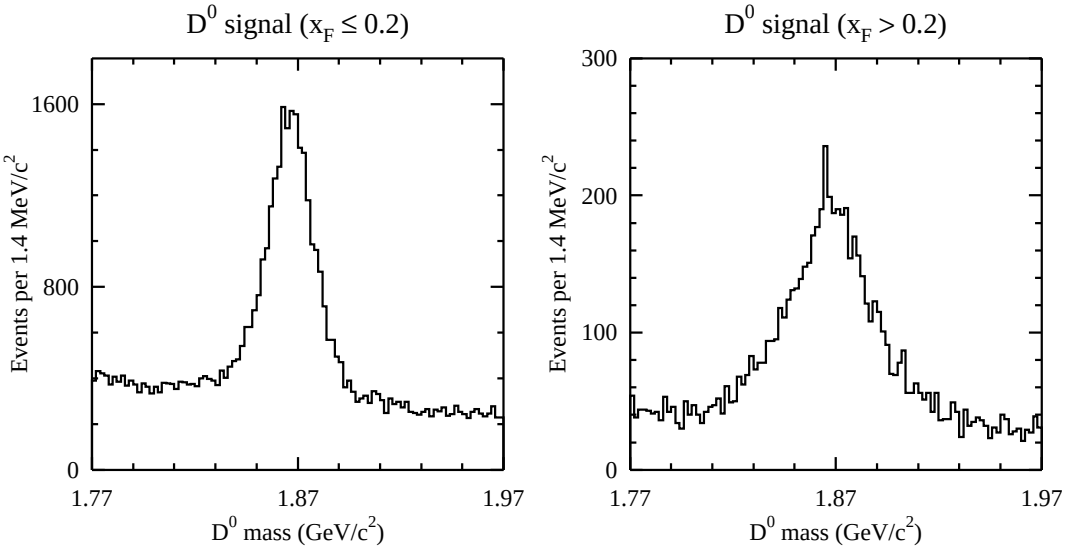

Figure 7.4: D^0 signal in regions of low and high x_F illustrating the dependence of resolution on x_F.

in Equation (3.6): the error on momentum measurements goes up with momentum. High x_F D^0s have a greater momentum than low x_F ones, and thus their daughter tracks tend to have a high momentum as well. This leads to a greater measurement error on the reconstructed mass, and a greater signal width.

For this reason, different cuts on the D^0 mass had to be employed in the different bins of x_F or p_T^2 under consideration. In each such bin, a D^0 mass histogram was created and fit with a single Gaussian and a linear background. The mean and width were extracted and only those D^0 candidates with a mass within a window of 2.5σ on either side of the mean were retained.

7.4 Differential Cross Sections

The data sample was divided into twelve bins each of x_F and p_T^2. Owing to the much higher statistics at lower x_F and p_T^2 narrower bins were used in those regions.

In the range $-0.1 < x_F < 0.4$, the bins have a width of 0.05, and in the range $0.4 < x_F < 0.6$ the bins have a width of 0.1. In the range $0 < p_T^2 < 2$ $(\mathrm{GeV}/c)^2$ bins have a width of 0.5 $(\mathrm{GeV}/c)^2$, and in the range $2 < p_T^2 < 10$ $(\mathrm{GeV}/c)^2$ bins have a width of 1 $(\mathrm{GeV}/c)^2$. When using non-uniform bins, it is necessary to divide the yields by the corresponding bin widths before plotting them, in order that the points on the plot model a probability density function. This procedure is followed throughout. For the sake of consistency, the p_T^2 analysis used only data with an x_F in the range studied ($-0.1 < x_F < 0.6$) and the x_F analysis used only data in the range $0 < p_T^2 < 10$ $(\mathrm{GeV}/c)^2$.

In each bin of x_F or p_T^2, $D^{*\pm}$ Q-value histograms were created. The event-mixed sample was divided into the same bins in order that the differences in background shape in these different bins were modeled correctly. The histograms were fit as described in Section 7.2, using the same signal and background regions as before.

In Figure 7.5, the Q-value histograms, along with their fits are plotted for each of the x_F bins. Figure 7.6 shows the same for the p_T^2 bins.

The procedure described above for the data was applied to the Monte Carlo sample (see Section 6.4.2) and each of the resulting yields was divided by the number of events generated (as obtained from truth tables) in the corresponding bin to obtain binned relative acceptance functions. The term "relative" here means that the absolute vertical scale is arbitrary. For a differential cross section study the absolute acceptance is not required. These acceptance functions are plotted in Figure 7.7.

The shapes of the acceptance functions are consistent with expectations. The detector is designed only to detect forward going particles, so there is very little acceptance at negative x_F. Because the center of mass frame is boosted forward in the lab frame, there is some acceptance at negative x_F, but nothing really below -0.1. At high x_F tracks were more likely to pass through the centers of the drift chambers

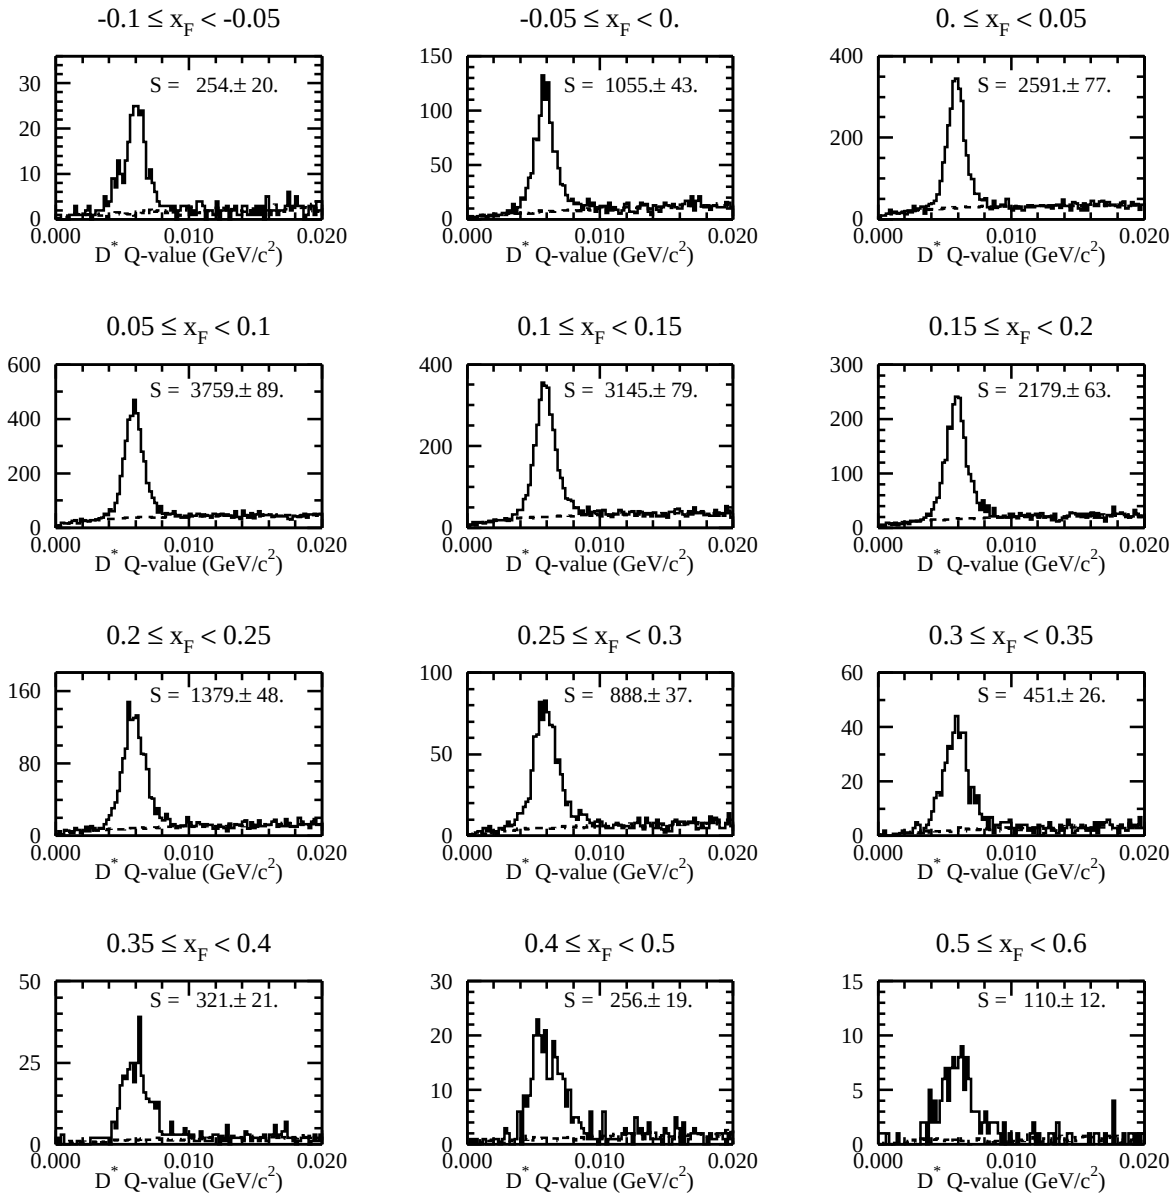

Figure 7.5: Fits of $D^{*\pm}$ Q-value histograms in bins of x_F to determine the numbers of signal events, using event mixing to model the backgrounds.

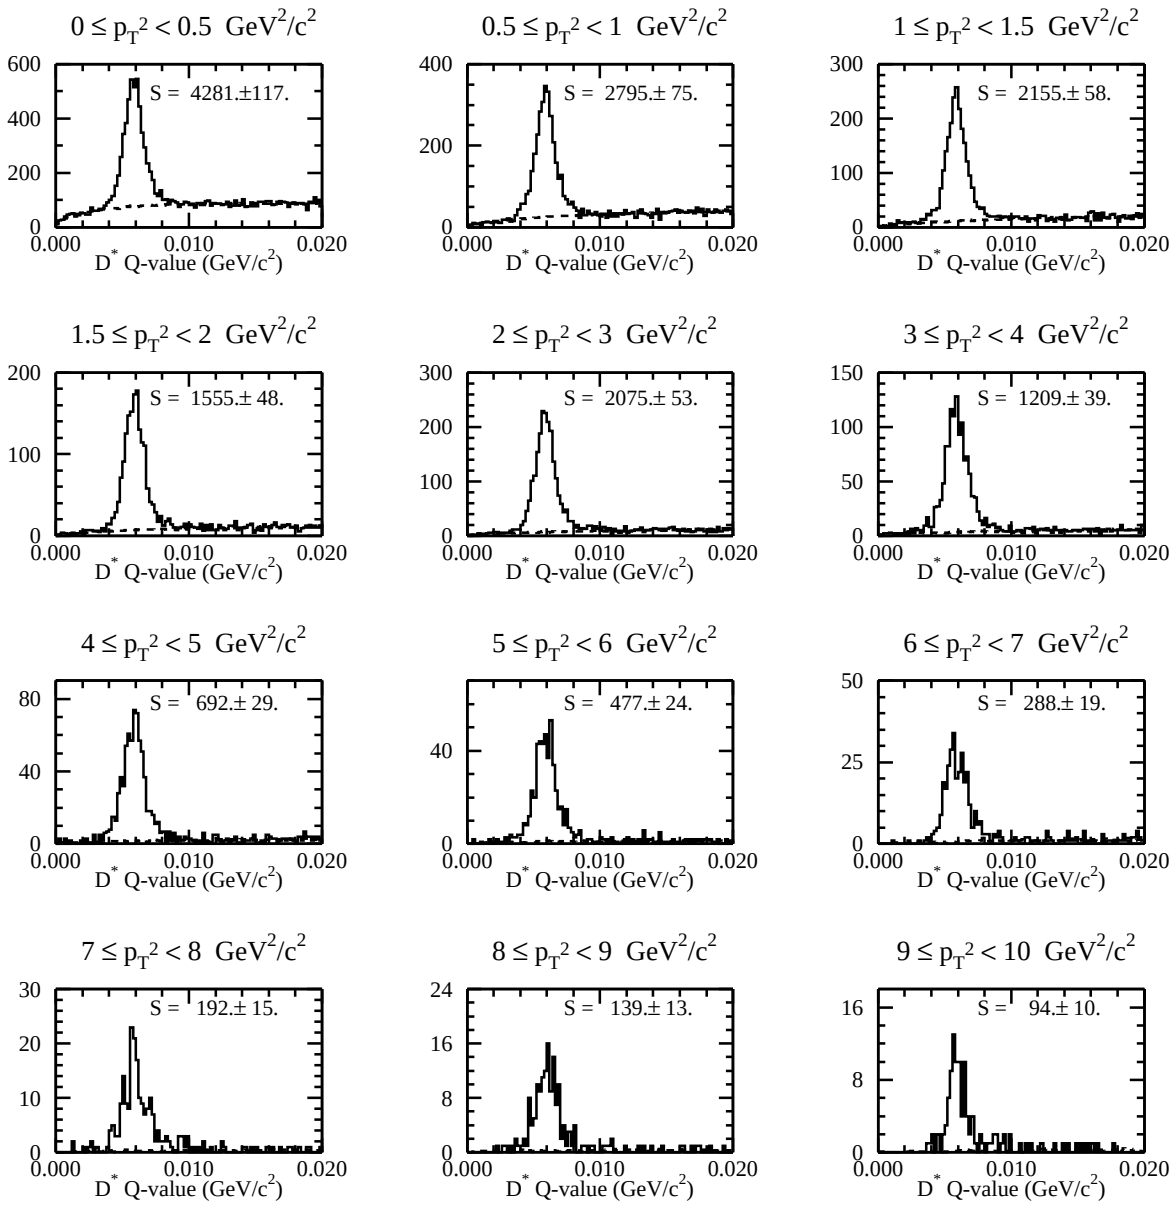

Figure 7.6: Fits of $D^{*\pm}$ Q-value histograms in bins of p_T^2 to determine the numbers of signal events, using event mixing to model the backgrounds.

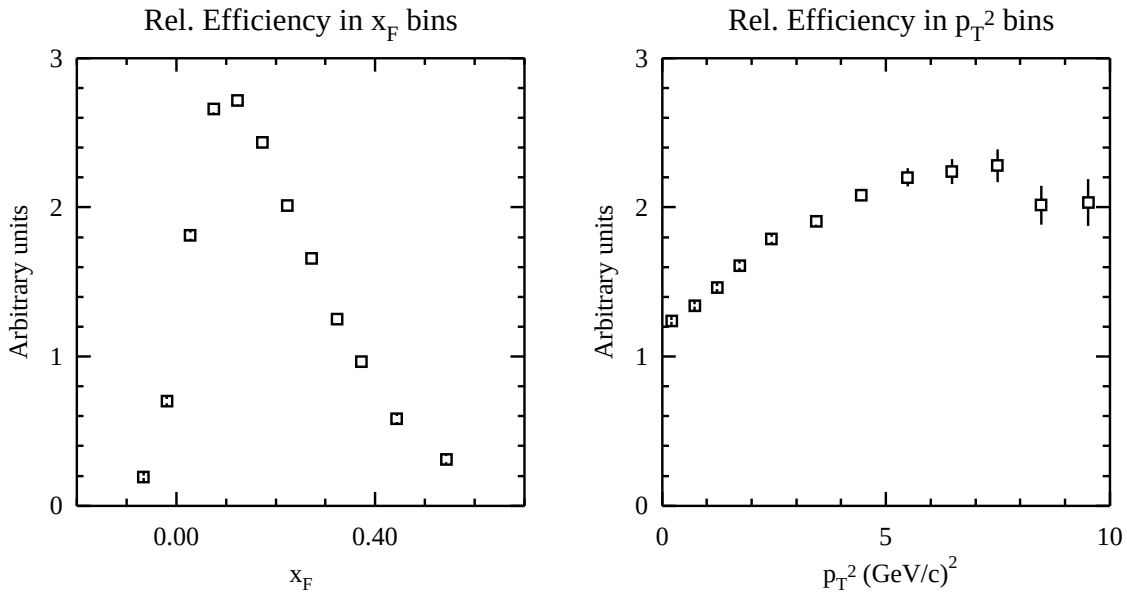

Figure 7.7: Binned relative efficiency functions of x_F and p_T^2 as determined from Monte Carlo.

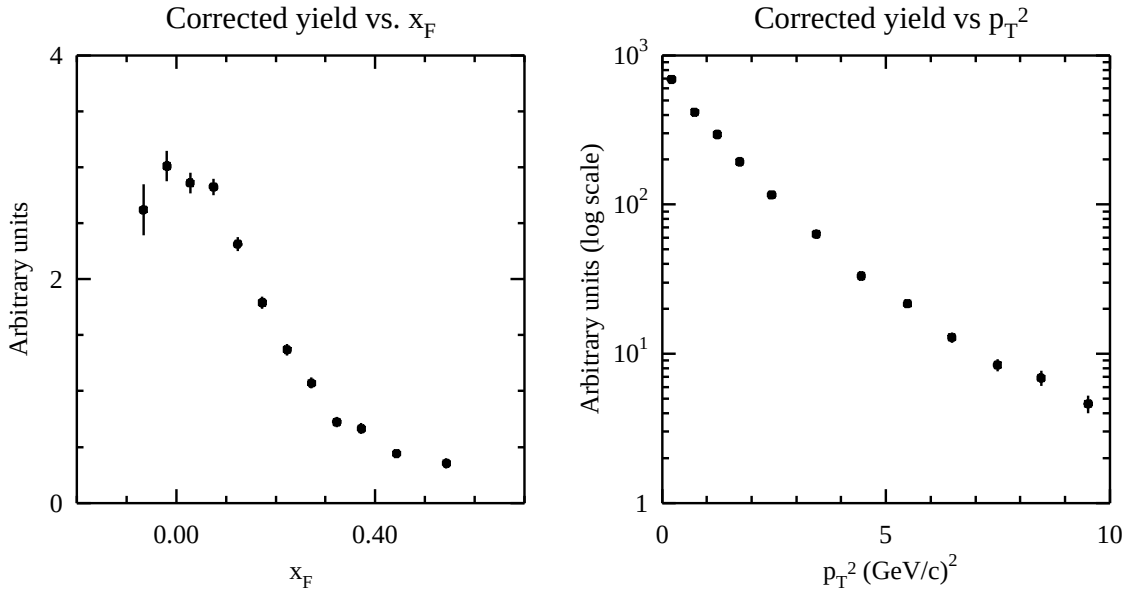

Figure 7.8: Acceptance corrected yields (differential cross sections) as functions of x_F and p_T^2.

where efficiency was lower. Also, at very high x_F, daughter tracks had less separation and a greater probability of approximately overlapping. It was difficult to resolve more than one track when they passed through the same channels of tracking system detectors. Also, at very high x_F, secondary vertices may have fallen outside the volume at which they were accepted. (Beyond the interaction counter, for example.) This accounts for the sharp drop in acceptance at high x_F. The p_T^2 acceptance is much flatter. The dip at low p_T^2 is probably due to the fact that such particles' daughters may be more likely to pass through the central regions of the drift chambers.

Each of the yields measured from the data sample was divided by the corresponding point in the acceptance function to obtain the points in the differential cross section plots shown in Figure 7.8. Because p_T^2 distributions are expected to decrease exponentially with p_T^2, a logarithmic vertical scale was used in that plot.

7.5 Consistency Checks

Although a detailed, quantitative study of systematic errors will be carried out in Chapter 10, some basic consistency checks on the result can be done at this point. Recall that our $D^{*\pm}$ sample was obtained through two different modes of the D^0 decay: $K\pi$ and $K\pi\pi\pi$. The momentum distributions of the D^*s depend on their production mechanism only, and can't be affected by the subsequent decay channels of their daughter D^0's. Thus any discrepancy in the measured D^* x_F and p_T^2 distributions between the two D^0 decay modes must be due to errors either in modeling acceptance or fitting our Q-value distributions. In fact, since the $K3\pi$ mode has two more daughter tracks than the $K\pi$ mode, such a comparison would be an excellent test of our modeling of the tracking efficiency. If the Monte Carlo overestimated our tracking efficiency at high x_F, for example, then the corrected x_F distribution for the $K\pi$ mode

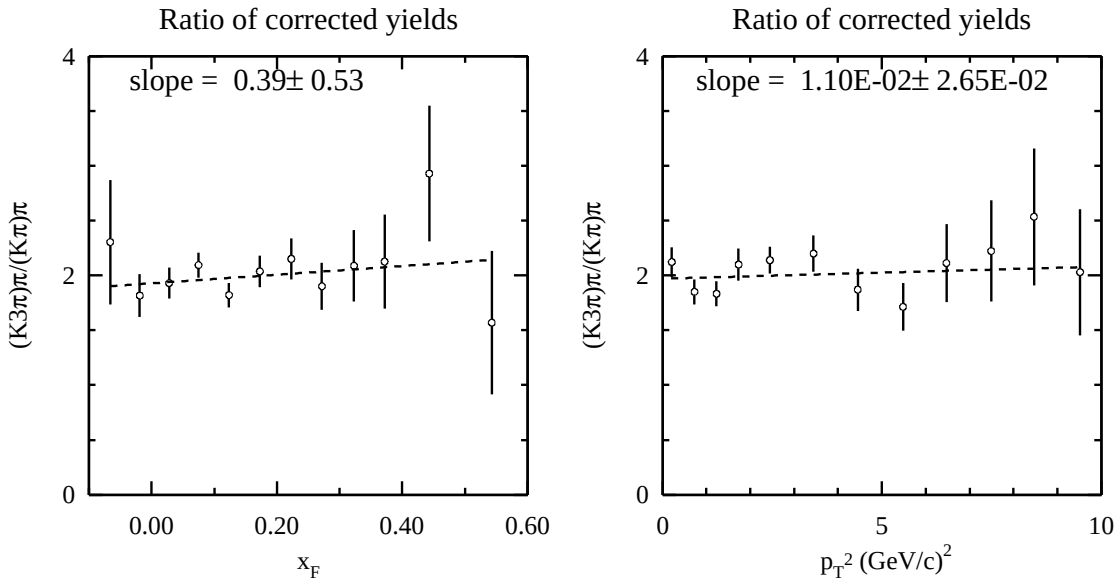

Figure 7.9: Ratios of corrected $D^{*\pm}$ x_F and p_T^2 distributions from channel $D^0 \to K3\pi$ to $D^0 \to K\pi$ with straight line fit.

should not fall as quickly as the $K3\pi$ mode. Since the acceptance is dominated by tracking efficiency, such a comparison provides an excellent consistency check.

To perform this comparison, the analysis described in Section 7.4 was carried out on the two subsamples of our data sample corresponding to the two decay modes of the D^0 considered. This involved computing new background models and acceptance functions for the two modes separately. Ratios of the resulting distributions are shown in Figure 7.9, along with straight line fits, which should be horizontal if the results are consistent. For the x_F ratio, the slope is 0.39 ± 0.53 and for p_T^2 it is $(1.1 \pm 2.6) \times 10^{-2}$. Both are consistent with zero, suggesting our acceptance model is good.

Recall from Section 5.3 that two different versions of reconstruction code were used on the dataset, named "Release 5" and "Release 7". It is important to check that the two subsamples reconstructed with the two different code versions agree. Ratio plots similar to those comparing the D^0 decay modes are shown in Figure 7.10, also with linear fits. The p_T^2 fit shows a slope of $(-0.44 \pm 1.59) \times 10^{-2}$, consistent

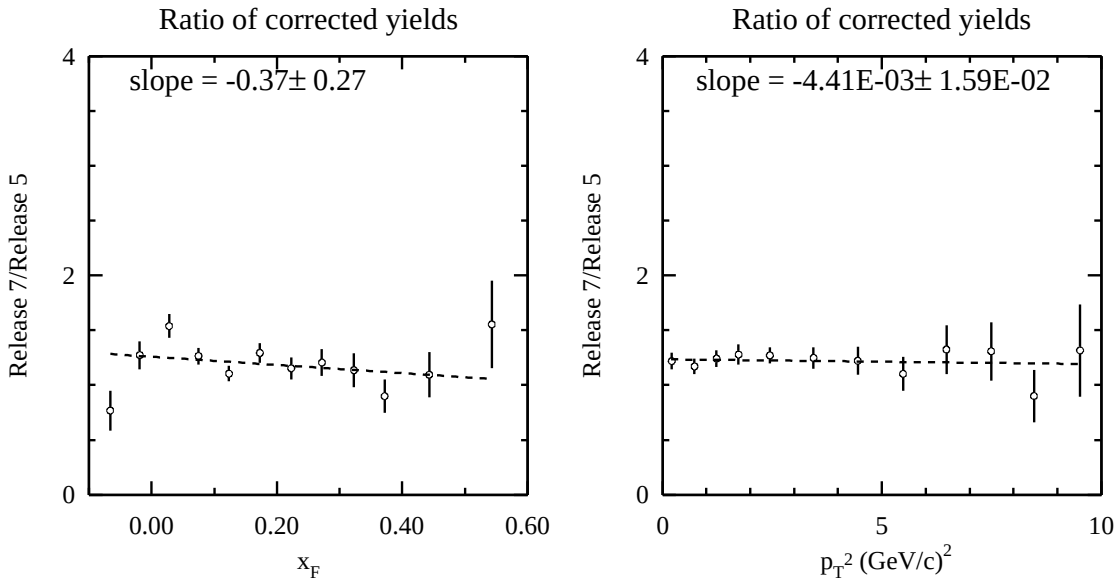

Figure 7.10: Ratios of corrected $D^{*\pm}$ x_F and p_T^2 distributions from release 7 data to release 5 with straight line fit.

with zero. The x_F fit has a slope of -0.37 ± 0.27. This is a 1.4σ effect, larger than one would like, but not entirely inconsistent with zero. In fact, visual inspection of the plot shows fluctuations of the points exhibiting no clear trend, so it is reasonable to conclude the results are consistent.

One might be tempted to compare subsamples of D^{*+} and D^{*-} this way, but differences there might be due to the physics of production, rather than acceptance modeling problems. In fact, such physics is the subject of the next chapter.

7.6 Fits and Comparisons

7.6.1 Traditional Forms

The x_F and p_T^2 distributions have traditionally been parameterized by the following functional forms [Al94b]:

$$\frac{d\sigma}{dx_F} \propto (1 - x_F)^n \tag{7.4}$$

and

$$\frac{d\sigma}{dp_T^2} \propto \exp(-Bp_T^2). \tag{7.5}$$

These have been found, in the past, to agree with data only in restricted ranges: at high x_F for Equation (7.4), and at low p_T^2 for Equation (7.5). For this reason, and for ease of comparison with other experiments, the x_F distribution was fit only in the range $0.1 < x_F < 0.6$ and the p_T^2 distribution was fit only in the range $0 < p_T^2 < 4$ $(\mathrm{GeV}/c)^2$. The results are shown in Figures 7.11 and 7.12. The error bars on these plots include systematic as well as statistical errors. Systematic errors will be dealt with in Chapter 10, but they are included here as they might affect the error bars on the fit results.

The fits both exhibit rather poor chi-square per degree of freedom (χ^2/ν) even in the restricted ranges used. For the x_F fit, $\chi^2/\nu = 2.4$ and for p_T^2, $\chi^2/\nu = 3.4$. This suggests the functional forms, especially for p_T^2, do not describe our data well at all.

7.6.2 Alternate Forms

Another functional form for the p_T^2 distribution has been suggested because it fits distributions computed in next-to-leading-order perturbative QCD extremely well

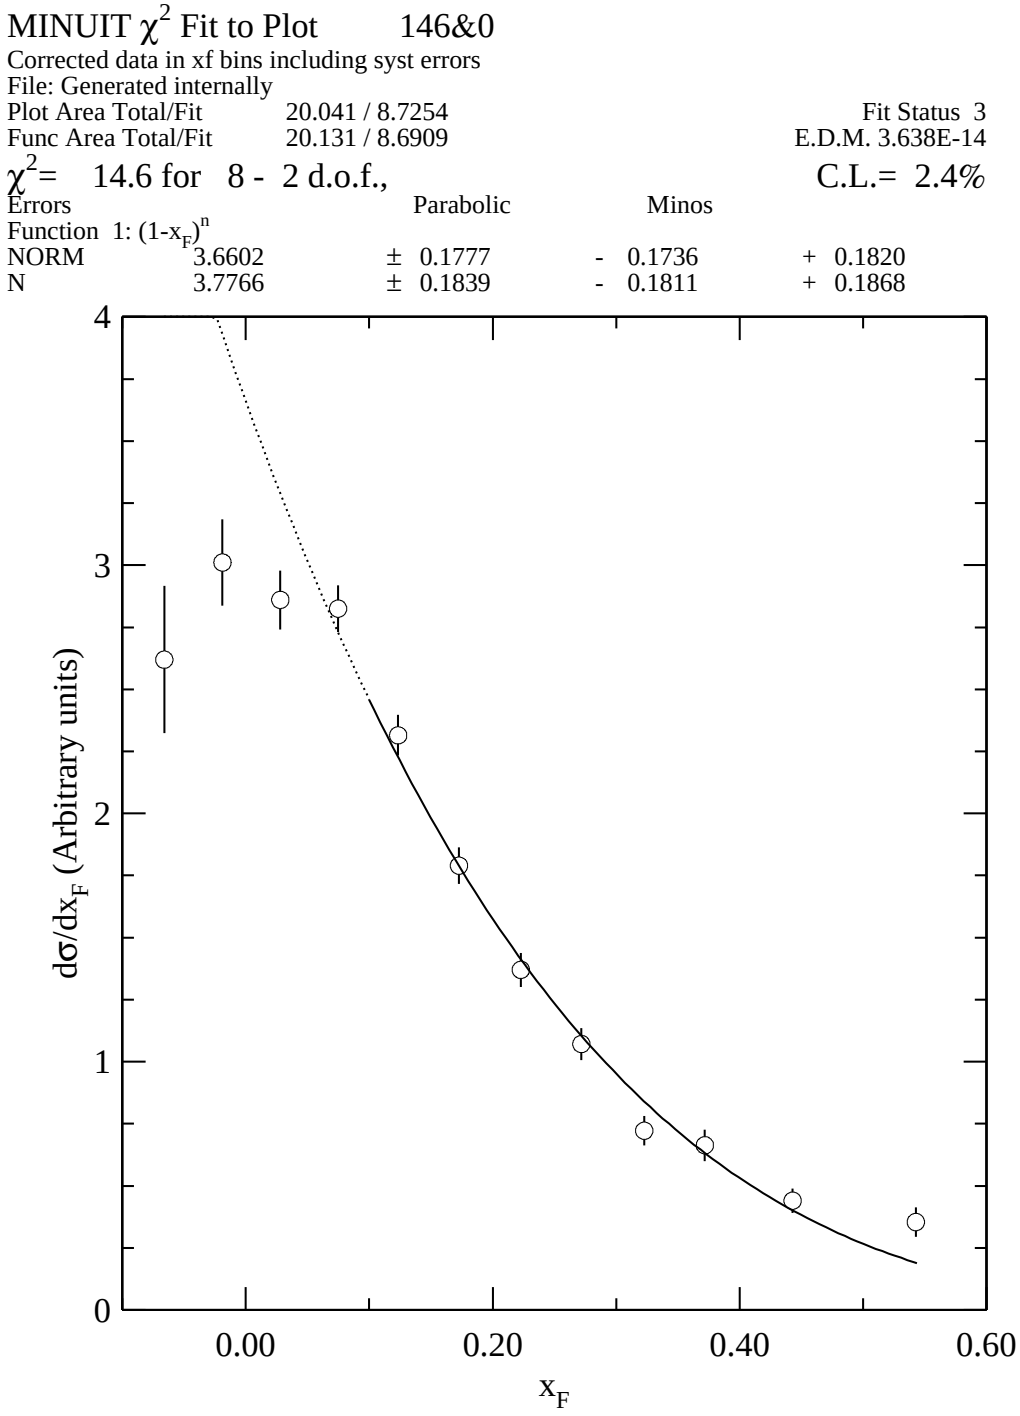

Figure 7.11: Acceptance corrected $D^{*\pm}$ x_F distribution fit with functional form $(1 - x_F)^n$ in the range $0.1 < x_F < 0.6$.

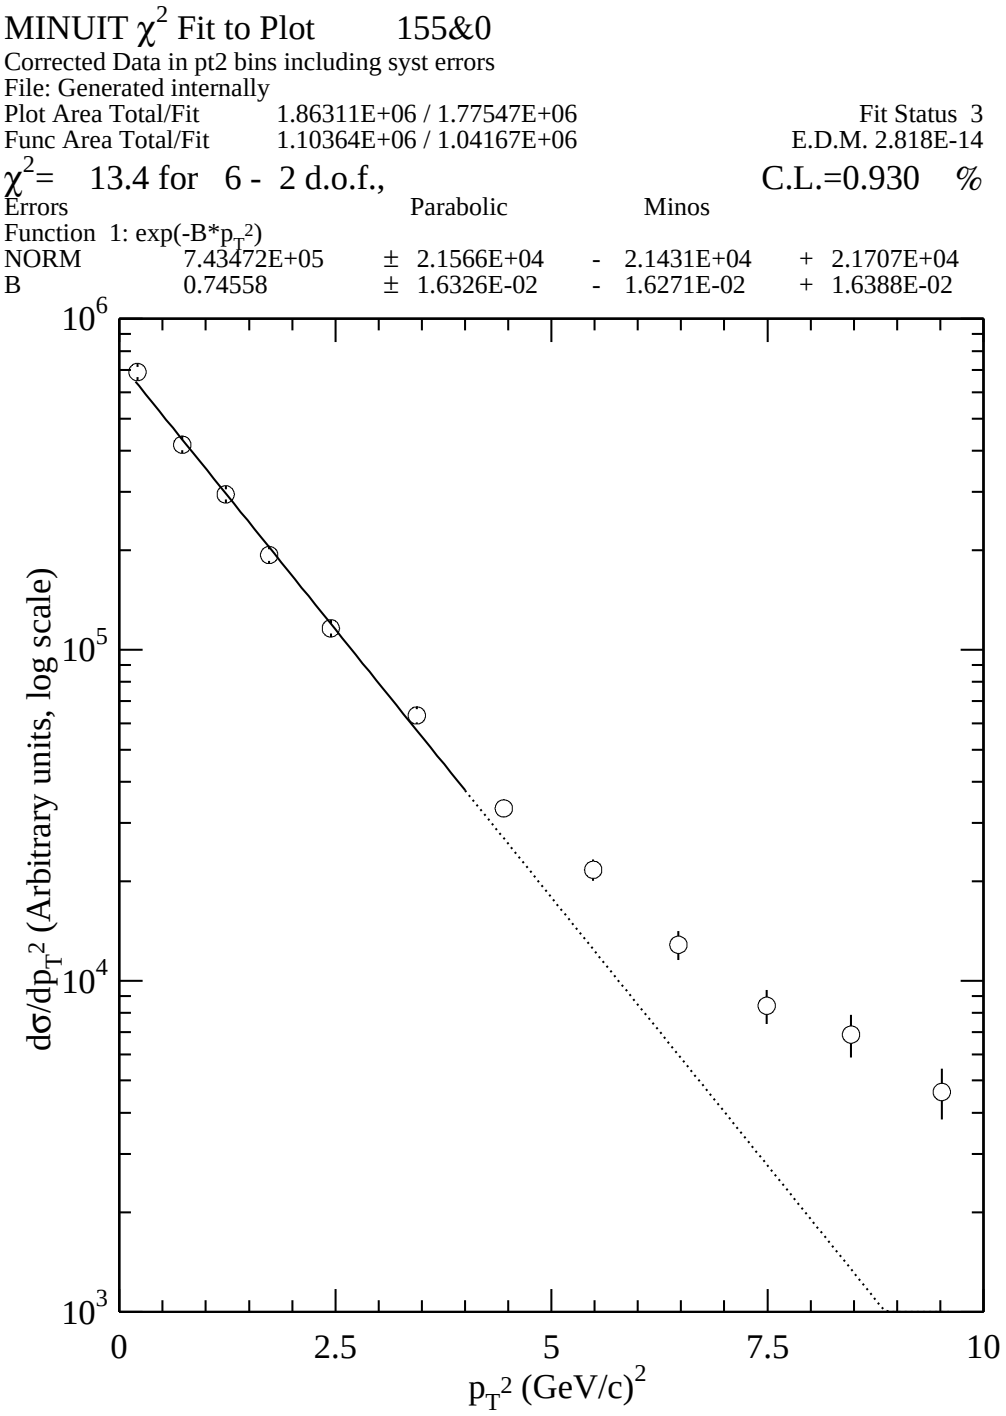

Figure 7.12: Acceptance corrected $D^{*\pm}$ p_T^2 distribution fit with functional form $\exp(-Bp_T^2)$ in the range $0 < p_T^2 < 4$ (GeV/c)^2.

over a broader p_T^2 range [Fr94]:

$$\frac{d\sigma}{dp_T^2} = \left(\frac{C}{bm_c^2 + p_T^2}\right)^{\beta} \tag{7.6}$$

The quantity m_c here is meant to be the charm quark mass, but we choose bm_c^2 as a free parameter as there is some uncertainty as to its exact value. A fit of the p_T^2 distribution in our full range of p_T^2 ($0 < p_T^2 < 10$ (GeV/c)^2) is shown in Figure 7.13.

Clearly, this functional form provides a much better fit than Equation (7.5). If the χ^2/ν of 0.57 seems impossibly low, remember that the error bars in the plot include a systematic component as well as statistical. If we perform a similar fit to a plot with statistical errors only, $\chi^2/\nu = 0.96$.

Because the x_F distribution is not peaked at exactly zero, a modified version of the functional form was tried:

$$\frac{d\sigma}{dx_F} \propto (1 - |x_F - x_0|)^{n'} \tag{7.7}$$

where x_0 is a parameter which is supposed to represent the peak of the distribution.

The function was fit in both in the range $x_F > 0.1$ and the full range of x_F. The results of these fits are shown in Figure 7.14, along with a fit of the standard function for comparison. Clearly in the restricted x_F range the x_0 parameter assumes the rather unphysical value of 0.12, the approximate location of the leftmost data point in the fit region. Also, the χ^2/ν is 2.5, greater than the value of 2.4 obtained when using the standard function. Putting bounds on the x_0 parameter didn't help; the fit kept forcing it to the upper limit.

When used on the full range of x_F, the x_0 is a more reasonable 9.86×10^{-3}, but χ^2/ν is 2.8, again worse than before. Moreover, the curve completely misses some

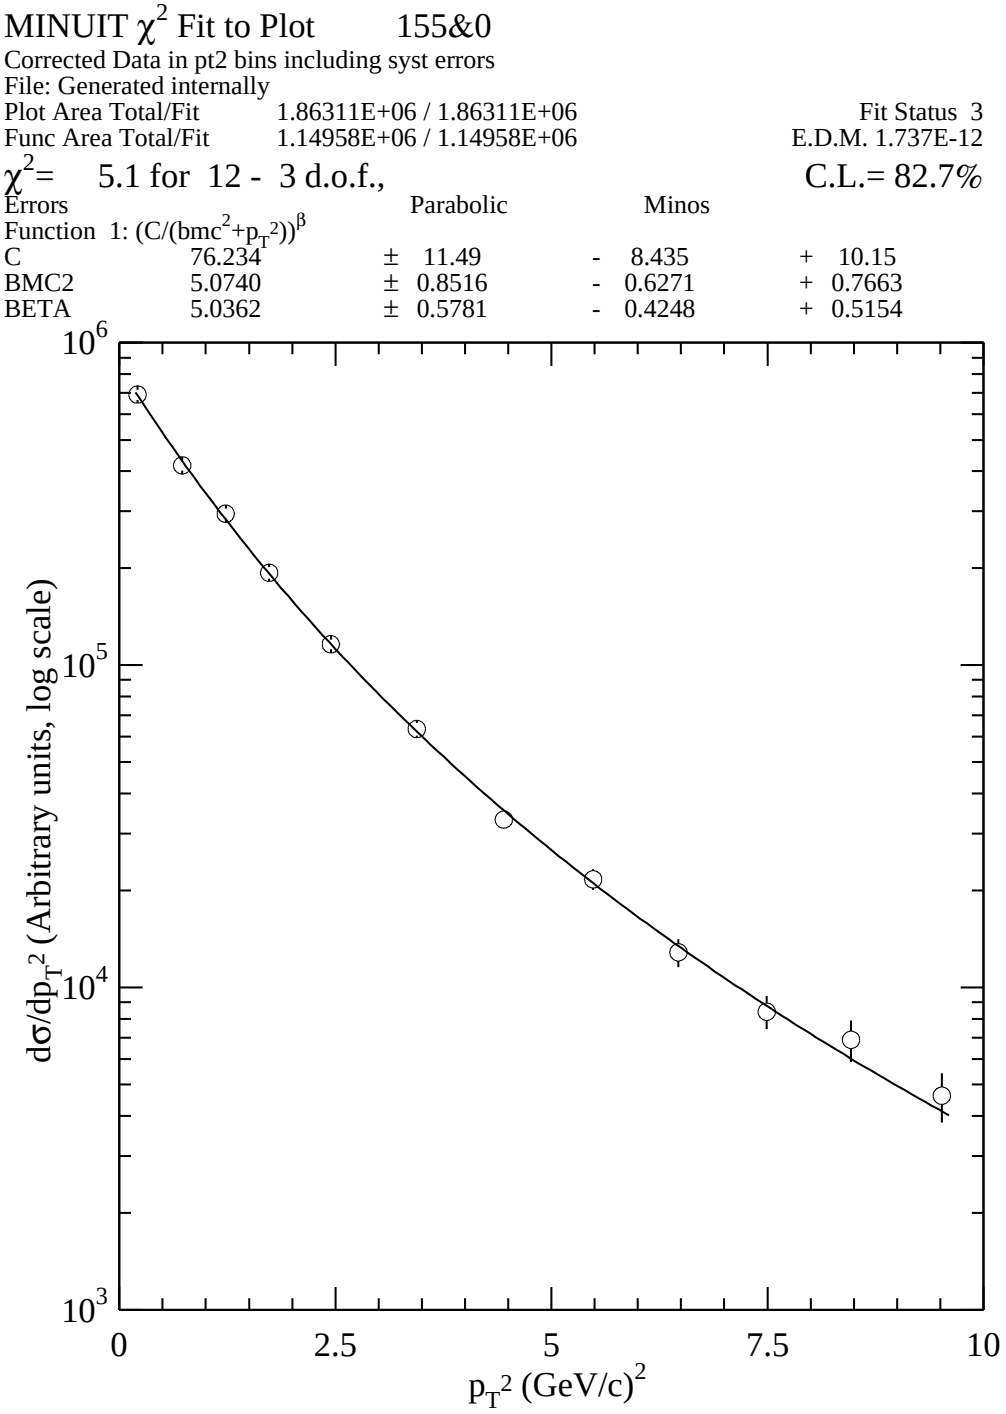

Figure 7.13: Acceptance corrected $D^{*\pm}$ p_T^2 distribution fit with functional form $C^\beta/(bm_c^2 + p_T^2)^\beta$ in the range $0 < p_T^2 < 10$ $(\mathrm{GeV}/c)^2$.

points at low x_F. This is not surprising as this functional form was never meant to model data at low x_F, and the peak shape of the curve (a sharp cusp, arising from the use of the absolute value function) is nothing like that of either data or Monte Carlo models.

Since the modified functional form neither models our data better (as evidenced by the χ^2/ν's) nor admits easy comparison to previous experiments, it is not used in this analysis.

Another form, used successfully in a D^0 analysis [Ai99], was the modified form above, joined to a Gaussian at lower x_F to model the peak area of the distribution. Ideally the Gaussian central value would converge on the peak of the distribution, enabling one to determine the "turnover" point.

$$\frac{d\sigma}{dx_F} = \begin{cases} A_1(1-|x_F-x_c|)^{n'}, & |x_F-x_c| > x_b \\ A_2 \exp[-\frac{1}{2}\left(\frac{x_F-x_c}{\sigma}\right)^2], & |x_F-x_c| < x_b \end{cases} \tag{7.8}$$

There are six parameters in the above equation: A_1, A_2, n', x_c, x_b, and σ. By requiring the function and its first derivative to be continuous at $x_F = x_b$, we eliminate two parameters (A_2 and σ) and allow the other four to float.

The result of the fit is shown in the lower right plot of Figure 7.14. Although the parameters converge to reasonable values ($x_c = 0.0122 \pm 0.0127$), the fit quality is not much improved (χ^2/ν=2.1), so it is not clear if it is meaningful to interpret x_c as the peak of the data distribution. Moreover, even if we are willing to accept that interpretation, the statistical significance is not good: we see the peak is less than 1σ from zero.

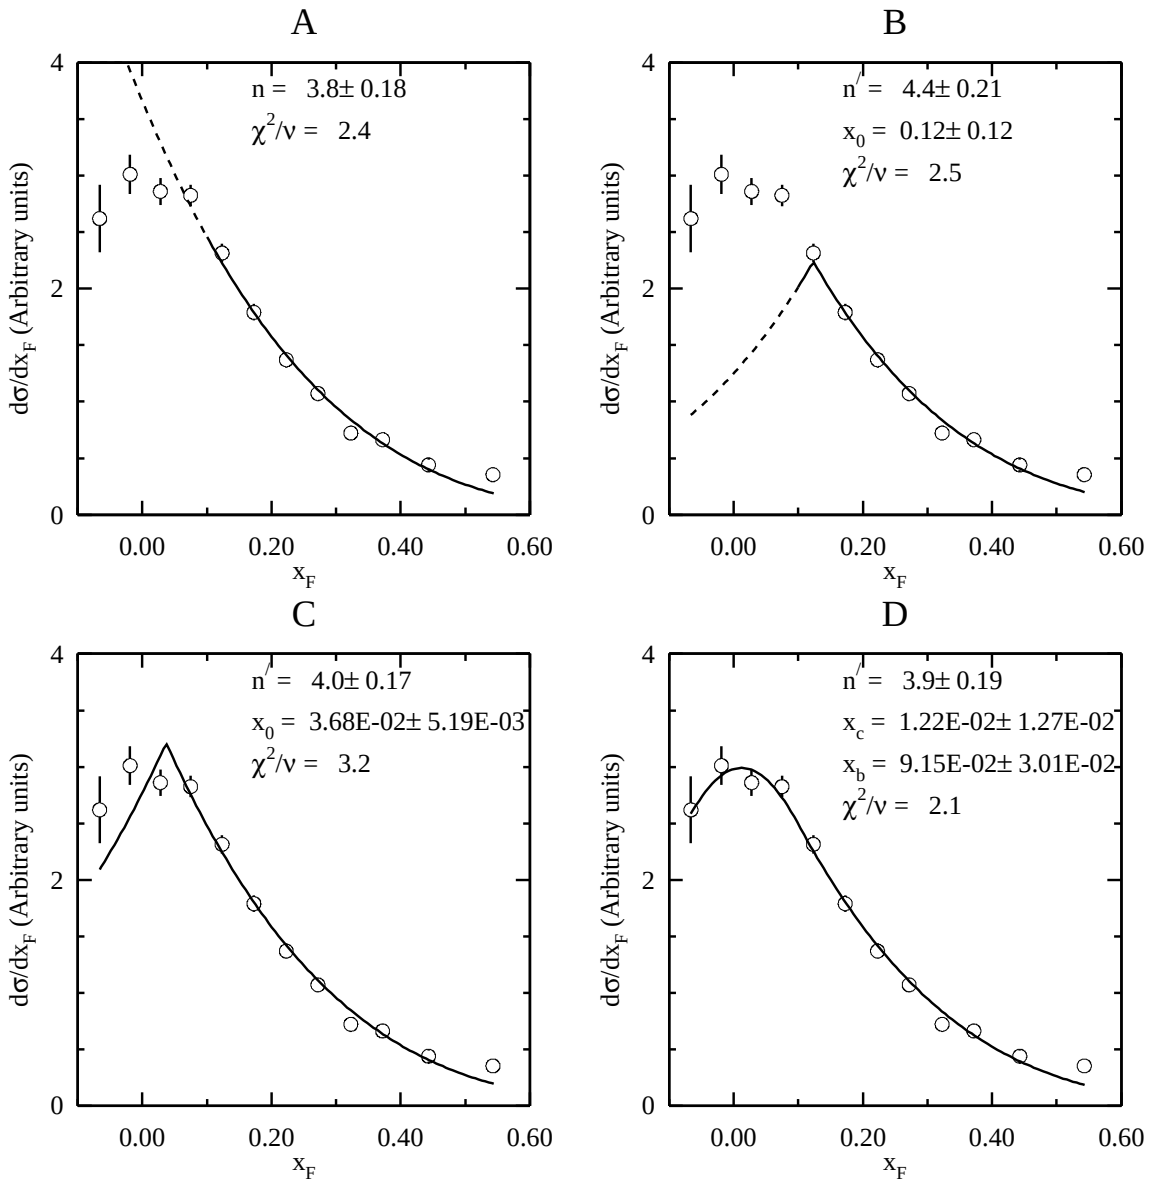

Figure 7.14: Acceptance corrected x_F distributions, fitted with (A) $(1 - x_F)^n$, as before over the range $0.1 < x_F < 0.6$, (B) $(1 - |x_F - x_0|)^{n'}$ in the same range, (C) same as (B) but in range $-0.1 < x_F < 0.6$, (D) same as (C) but with a Gaussian for $|x_F - x_C| < x_B$.

Table 7.1: Parameters from fits to $\pi^{\pm}N$ differential $D^{*\pm}$ cross sections vs. x_F and p_T^2, including comparisons to other experiments ([Ad99, Al94b, Ba91, Ba88, Ag86]). Column labels refer to fit parameters of the same name defined by Equations (7.4)–(7.6). The labels "l" and "nl" refer to leading and non-leading particles, respectively. Asterisks (*) indicate extended range of p_T^2 used in fit of Equation (7.6). Errors shown on E791 results include both statistical and systematic components.

Expt.	x_F fit range	n	p_T^2 fit rng. $(\mathrm{GeV/c})^2$	B $(\mathrm{GeV}/c)^{-2}$	bm_c^2 * $(\mathrm{GeV}/c)^2$	β *
E791	0.1 to 0.6	3.8 ± 0.2	0 to 4	0.75 ± 0.02	5.1 ± 0.8	5.0 ± 0.6
l. (D^{*-})		3.6 ± 0.2	or	0.73 ± 0.02	6.1 ± 1.5	5.6 ± 1.0
nl. (D^{*+})		4.1 ± 0.3	0 to 10*	0.77 ± 0.03	4.0 ± 1.1	4.4 ± 0.8
WA92[Ad99]	0 to 0.6	4.3 ± 0.4	0 to 4	0.84 ± 0.05	3.6 ± 0.8	4.7 ± 0.7
l.		4.9 ± 0.5	or			
nl.		3.9 ± 0.4	0 to 14*			
E769[Al94b]	0.1 to 0.6	3.5 ± 0.3	0 to 4	0.70 ± 0.07		
l.		2.9 ± 0.4		0.58 ± 0.09		
nl.		4.1 ± 0.5		0.79 ± 0.09		
NA32[Ba91]	0 to 0.8	$3.14^{+0.40}_{-0.39}$	0 to 10	0.79 ± 0.07		
l.		$2.62^{+0.53}_{-0.49}$		$0.71^{+0.09}_{-0.08}$		
nl.		$3.83^{+0.66}_{-0.62}$		0.90 ± 0.11		
NA32[Ba88]	0 to 0.7	$2.8^{+1.1}_{-0.9}$	0 to 5	$0.9^{+0.3}_{-0.2}$		
l.		$4.7^{+1.9}_{-1.6}$				
nl.		$1.7^{+1.4}_{-1.0}$				
NA27[Ag86]	0 to 0.5	$4.3^{+1.8}_{-1.5}$	0 to 3	0.9 ± 0.4		

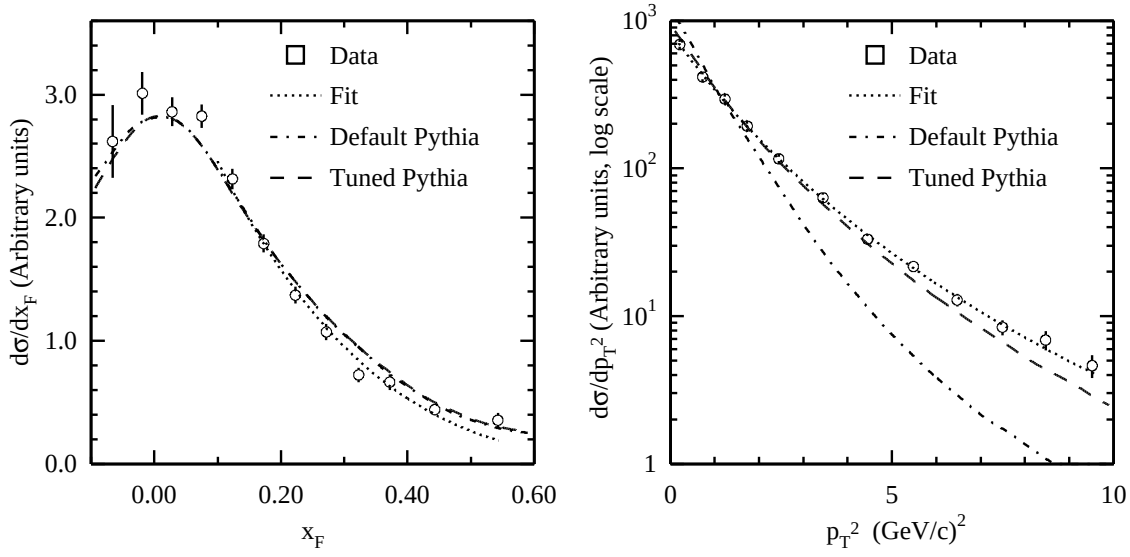

Figure 7.15: Acceptance corrected D^* x_F and p_T^2 distributions showing fits (Eqs. (7.4) and (7.6)), and both default and tuned Pythia models.

7.6.3 Comparisons to Previous Experiments

The D^{*-} and D^{*+} subsamples were also analyzed and fit separately, as their distributions might be different. The resulting fit parameters, along with comparisons to other experiments, are shown in Table 7.1. As can be seen, results are largely consistent with previous measurements.

7.6.4 Comparison to Pythia Models

The leading model for describing x_F and p_T^2 distributions is the PYTHIA model [Sj94], which employs the Lund string fragmentation model. This model has many tunable parameters. In a previous E791 analysis on D^+/D^- production asymmetry [Ai96], several parameters were tuned to improve the agreement between the PYTHIA model and experimental data. Specifically, PMAS(4,1), the charm quark mass in GeV/c^2 was changed from 1.35 to 1.7, PARP(91), the average k_T^2 in (GeV/c)^2 was increased

from 0.44 to 1.0, PARP(93), the maximum allowable k_T in GeV/c was increased from 2.0 to 5.0, and MSTP(92) the remnant quark-diquark energy splitting function, was changed from 4 to 3. These variables are all described in detail in reference [Sj94]. The effect of most of these changes is to increase the cross sections at higher p_T^2.

The x_F and p_T^2 distributions are shown in Figure 7.15, along with the fits shown previously, and comparisons to both the default and tuned PYTHIA models. In x_F, there is very little difference between the tuned and untuned models, and they fit the data reasonably well although they seem to predict a slightly stiffer x_F distribution. The tuned and untuned models are quite different in the p_T^2 distribution, the tuned model yielding a much stiffer distribution, which agrees with the data fairly well. Equation (7.6), which is based on fits to NLO QCD distributions, describes the data better, however, our data having a slightly stiffer distribution than the tuned PYTHIA model.

Chapter 8

Production Asymmetries between D^{*+} and D^{*-}

8.1 Definitions

Recall that a leading particle (for $x_F > 0$) is a final state particle which has a quark in common with the beam particle while its antiparticle does not. In E791, a π^- ($\overline{u}\,d$) beam was used. The D^{*+} ($c\,\overline{d}$) has no quark in common with the beam particle, whereas the D^{*-} ($\overline{c}\,d$) has a d quark in common with the beam particle. Therefore, the D^{*-} is a leading particle and the D^{*+} is a non-leading particle.

The asymmetry parameter A is thus defined as (cf. Eq. 1.4):

$$A = \frac{\sigma(D^{*-}) - \sigma(D^{*+})}{\sigma(D^{*-}) + \sigma(D^{*+})} \tag{8.1}$$

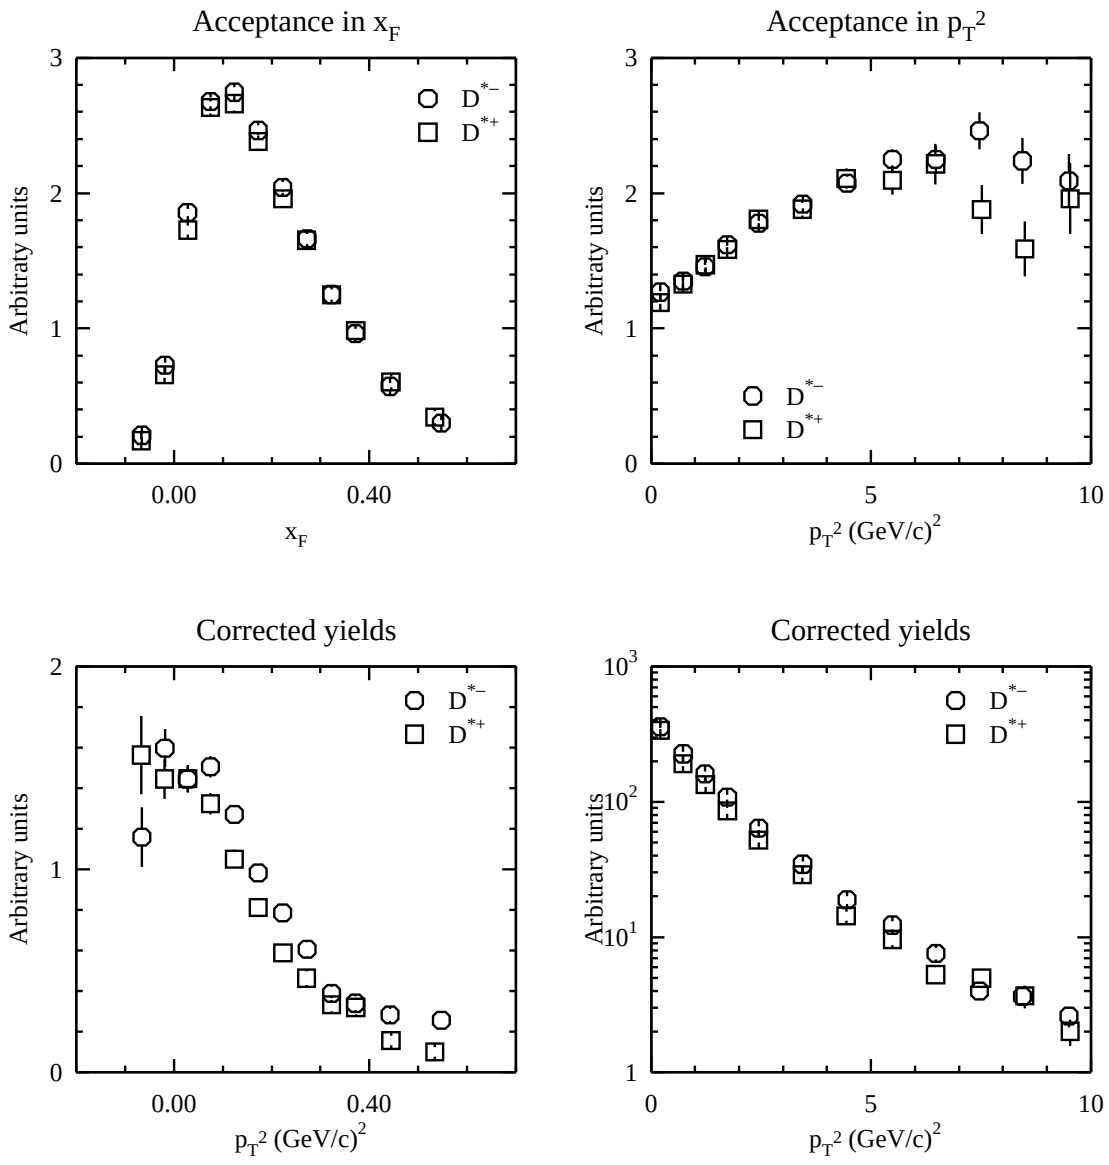

Figure 8.1: Binned relative acceptance functions and corrected distributions as functions of x_F and p_T^2 for D^{*+} and D^{*-}.

8.2 Asymmetry Measurement

The analysis technique is virtually identical to that described in the previous chapter on differential cross sections. The data sample was divided into bins of x_F and p_T^2 using the same binning as before. This time the D^{*+} and D^{*-} samples were analyzed separately. The fitting technique is the same as used previously, and separate acceptance functions were computed for D^{*+} and D^{*-} using the Monte Carlo sample. The acceptances and corrected distributions are shown in Figure 8.1.

The acceptances are roughly the same for the two charges. This is not surprising, as the $D^{*\pm}$ decays into a D^0 which has equal numbers of positively and negatively charged daughter tracks, and a charged pion with a very small momentum, making it unlikely to be bent into the drift chamber "holes" through which the beam particles passed, even at high x_F.

With these distributions, the asymmetry parameter can be computed in each bin using the formula

$$A = \frac{N(D^{*-}) - N(D^{*+})}{N(D^{*-}) + N(D^{*+})} \tag{8.2}$$

where N denotes the acceptance corrected number of signal events in each bin of the production variable.

8.3 Results and Comparisons

The asymmetry parameter is plotted as a function of x_F and p_T^2 in Figures 8.2 and 8.3 along with those measured by E791 for D^+/D^- asymmetry [Ai96] and the D_s^+/D_s^- asymmetry [Ai97] for comparison. The $D^{*\pm}$ results are consistent with the $D^{\pm}$ results. Note that neither the D_s^+ $(c\,\overline{s})$ nor the D_s^- $(\overline{c}\,s)$ are leading particles. It is therefore not surprising that the asymmetry is lower (in fact, consistent with zero) and does

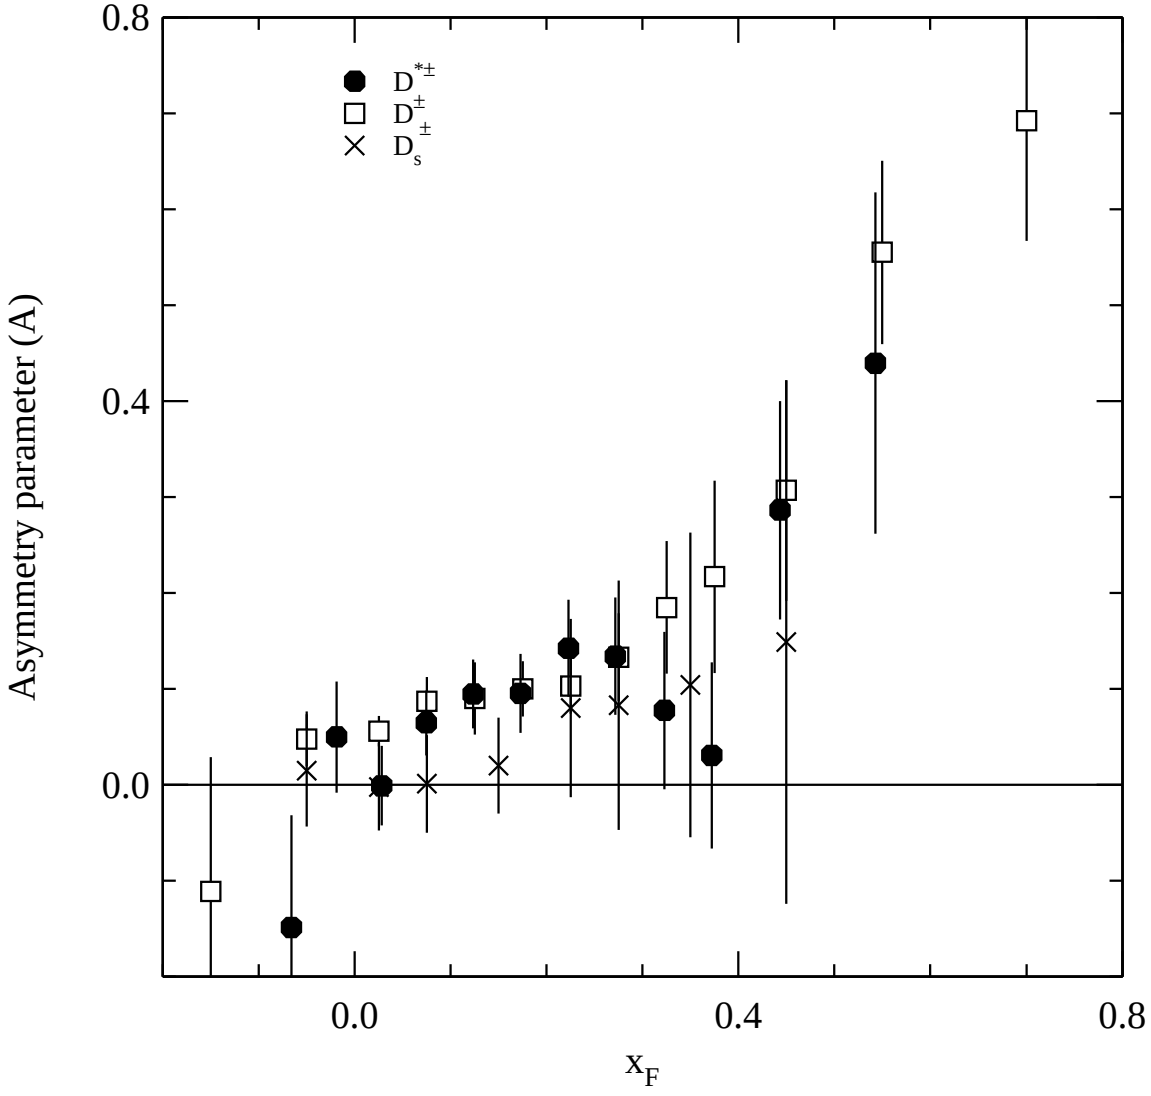

Figure 8.2: Asymmetry parameter for $D^{*\pm}$ as a function of x_F. Also shown for comparison are E791 results for $D^{\pm}$ [Ai96] and $D_s^{\pm}$ [Ai97].

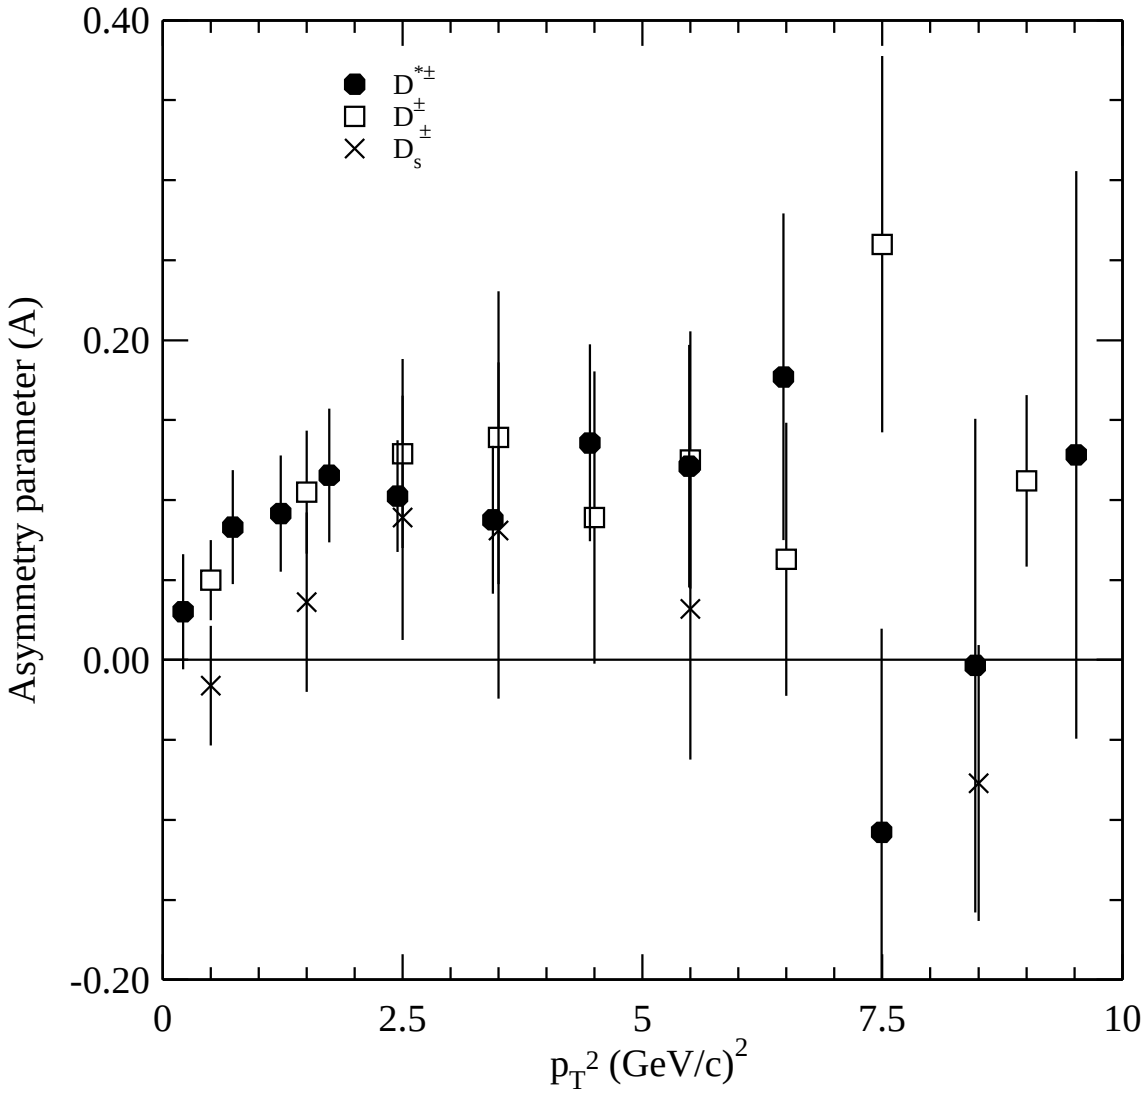

Figure 8.3: Asymmetry parameter for $D^{*\pm}$ as a function of p_T^2. Also shown for comparison are E791 results for $D^{\pm}$ [Ai96] and $D_s^{\pm}$ [Ai97].

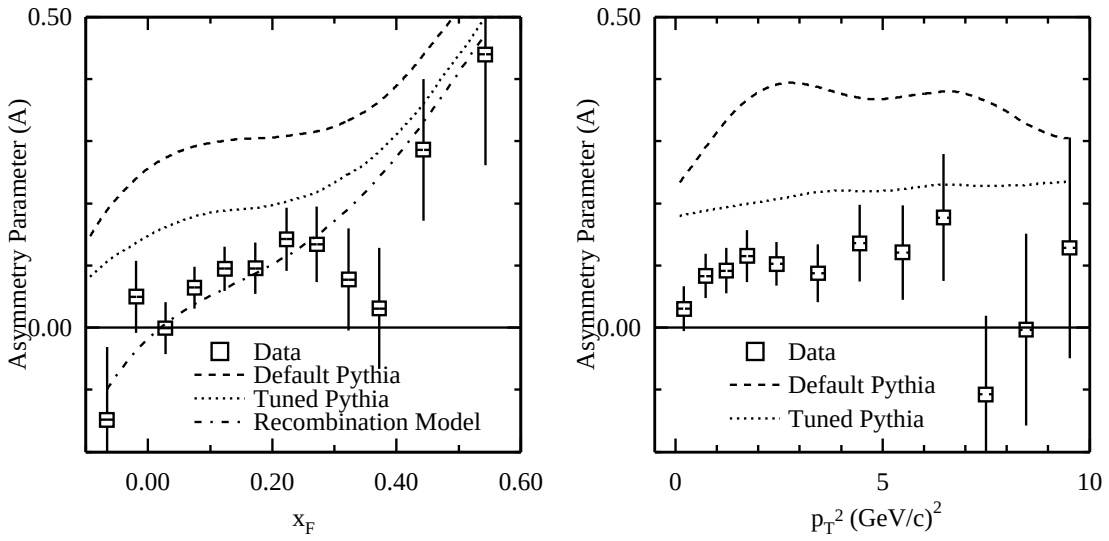

Figure 8.4: Asymmetry parameter for $D^{*\pm}$ compared to default and tuned PYTHIA models, as well as the recombination model (x_F plot only)[Av03].

not rise with x_F. As in the $D^{\pm}$ system, we observe a strong enhancement in A at high x_F, the so-called "leading particle effect", which suggests that NLO QCD calculations by themselves are not sufficient to account for the asymmetries, and therefore fragmentation plays a role. The p_T^2 asymmetry is consistent with being flat. The overall asymmetry observed is probably due to the fact that we use data only in the range $-0.1 < x_F < 0.6$ (essentially positive x_F) so the leading particle effect would thus give rise to an overall asymmetry in the p_T^2 curve.

The results are also plotted in Figure 8.4, along with comparisons to the PYTHIA model. Both the default PYTHIA model and a PYTHIA model with parameters tuned as discussed in Section 7.6.4 are shown. Recall that the parameters were originally tuned specifically to bring about better agreement with $D^{\pm}$ asymmetry results. The agreement with the observed $D^{*\pm}$ asymmetry is not as good. Both the tuned and untuned models predict roughly the correct shape but they predict a greater overall asymmetry than observed. The tuned model is closer to the measured results, how-

ever. It is interesting to note that although the $D^{*\pm}$ results agree with the measured $D^{\pm}$ results, the PYTHIA model predicts a greater asymmetry in the $D^{*\pm}$ system than it does in the $D^{\pm}$ system. Overall, it can be concluded that the Lund String model can account for the shape, but with default parameters, it predicts far too large an asymmetry.

Also shown in Figure 8.4 is a fit[Av03] which includes contributions from both a fragmentation model (the Peterson function) and a recombination model (Das and Hwa[Da77]). It provides a much better fit to the data than either the untuned or tuned PYTHIA models.

Chapter 9

$D^{*\pm}$ Polarization

9.1 Introduction

The $D^{*\pm}(2010)$ is a spin-1 particle, and therefore it can exist in one of three possible spin angular momentum states, with $m_s = -1$, 0, or 1 along any chosen quantization axis. In a high energy physics experiment, many such particles are observed, and so a means of characterizing the distributions of these spin states is needed.

A quantum mechanical system which consists of a statistical mixture of identical subsystems, each of which can be described by a state vector, can in general be characterized by a *density operator* [Bl96]. A density operator is a linear combination of projection operators onto the possible states weighted by their statistical probabilities:

$$\rho = \sum_i w_i |\psi_i\rangle\langle\psi_i| \tag{9.1}$$

where the $\{|\psi_i\rangle\}$ is a set of state vectors describing the subsystems, and w_i is the probability that any subsystem is in a state described by $|\psi_i\rangle$.

With respect to a given orthonormal basis $|u_i\rangle$, we can define the *density matrix*

as:

$$\rho_{ij} = \langle u_i|\rho|u_j\rangle \tag{9.2}$$

From the definition (9.1), it immediately follows that ρ_{ij} is hermitian, and the fact that $\sum w_i = 1$ implies that

$$\mathrm{Tr}(\rho_{ij}) = 1 \tag{9.3}$$

The angular momentum distribution of a sample of $D^{*\pm}$ particles can thus be described by a *spin density matrix*. In any given basis, there are three possible states, so ρ is a 3×3 matrix. A general complex 3×3 matrix has 18 real independent components, but hermiticity and the fact that $\mathrm{Tr}(\rho_{ij}) = 1$ reduces the number to 8. The D^*'s are produced in a strong interaction, and parity conservation can be shown to further reduce this number to 4. (See Appendix A). For relativistic particles, it is often most convenient to work in a *helicity basis* where the spin quantization axis is chosen along the particle's momentum vector. In such a basis, the spin density matrix takes the form given by Equation (A.38):

$$\rho = \begin{pmatrix} \rho_{11} & \rho_{10} & \rho_{1-1} \\ \rho_{10}^* & 1-2\rho_{11} & -\rho_{10}^* \\ \rho_{1-1} & -\rho_{10} & \rho_{11} \end{pmatrix} \tag{9.4}$$

with ρ_{1-1} and of course ρ_{11} real.

The angular distributions of the decay products of the D^* depend on the elements of the spin density matrix, and so measurements of these angular distributions can be used to determine the spin density matrix elements of the D^*. It is necessary to define a coordinate system and reference frame in which to define the angles that describe the directions of the decay products. The E791 coordinate system defined

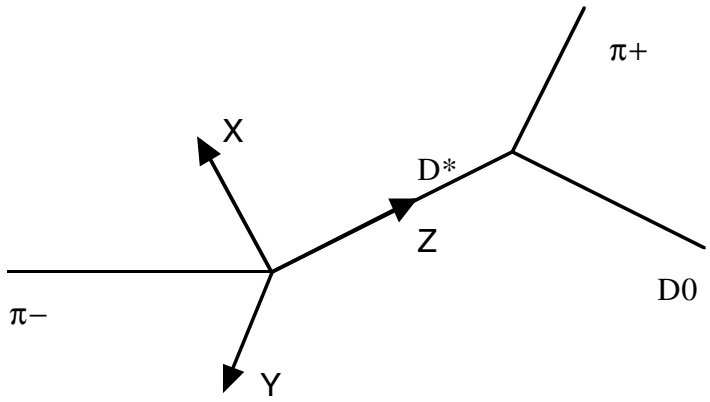

Figure 9.1: Coordinate system used to define angular distribution of D^* decay products.

in Section 3.3 is not convenient, as it is related to detector elements rather than the particles under consideration. A helicity basis was chosen, so that the z-axis points along the D^* boost vector in the lab frame. The y-axis is defined as pointing normal to the production plane. That is, the positive y-axis points in the direction $\mathbf{d} \times \mathbf{b}$ where $\mathbf{d}$ is the D^* boost vector and $\mathbf{b}$ is a vector pointing in the direction of the beam particle's momentum. The choice of a right-handed coordinate system then uniquely determines the direction of the x-axis. (See Figure 9.1).

The decay product angles are evaluated in the rest frame of the D^*. One might object to the fact that the definition of the z-axis (along the D^* boost vector) makes no sense in the D^*'s rest frame, but if the x and y axes are defined in the lab frame, they will be invariant under a boost to the D^* rest frame as they are perpendicular to the boost direction, and they can be used to define the z-axis in the D^* rest frame. Spherical polar coordinates are used to measure the direction of one of the decay products (the slow pion). θ is the angle between the slow pion momentum vector and

the z-axis. ϕ is the angle between the projection of the slow pion momentum vector in the xy-plane and the x-axis.

With these definitions, the angular dependence of the cross section is given by Equation (A.52):

$$\frac{d\sigma}{d\Omega} \propto \frac{1}{4\pi}[1 + (1 - 3\cos^2\theta)(3\rho_{11} - 1) - (3\sqrt{2}\sin 2\theta\cos\phi)\mathrm{Re}\,\rho_{10} - (3\sin^2\theta\cos 2\phi)\rho_{1-1}]. \tag{9.5}$$

Note that this distribution depends on only three of the four independent components of the spin density matrix, namely: ρ_{11}, $\mathrm{Re}\,\rho_{10}$, and ρ_{1-1}. Thus $\mathrm{Im}\,\rho_{10}$ cannot be determined from an angular distribution study.

Recall that the slow pion was somewhat arbitrarily chosen as the decay particle whose angles were measured. The D^* has two decay products: the slow pion and the D^0, which emerge back to back in the D^* rest frame. Thus choosing the D^0 rather than the slow pion is equivalent to replacing θ by $\pi - \theta$ and ϕ by $\phi + \pi$. A quick check shows that Equation (9.5) is invariant under such a replacement, so the results don't depend on that arbitrary choice.

Equation (9.5) is a probability defined per unit solid angle:

$$d\Omega = \sin\theta\, d\theta\, d\phi = -d(\cos\theta)\, d\phi \tag{9.6}$$

and so in order for an angular distribution histogram to approximate a probability density function, $\cos\theta$ and ϕ must be chosen as independent variables, rather than θ and ϕ. (It would also be possible to multiply both sides of Equation (9.5) by $\sin\theta$ and use θ and ϕ as independent variables, but the former technique is simpler as the unpolarized case leads to a flat distribution in $\cos\theta$ and ϕ.)

Note that if the D^*s are unpolarized, (i.e. all three helicity states occur with equal

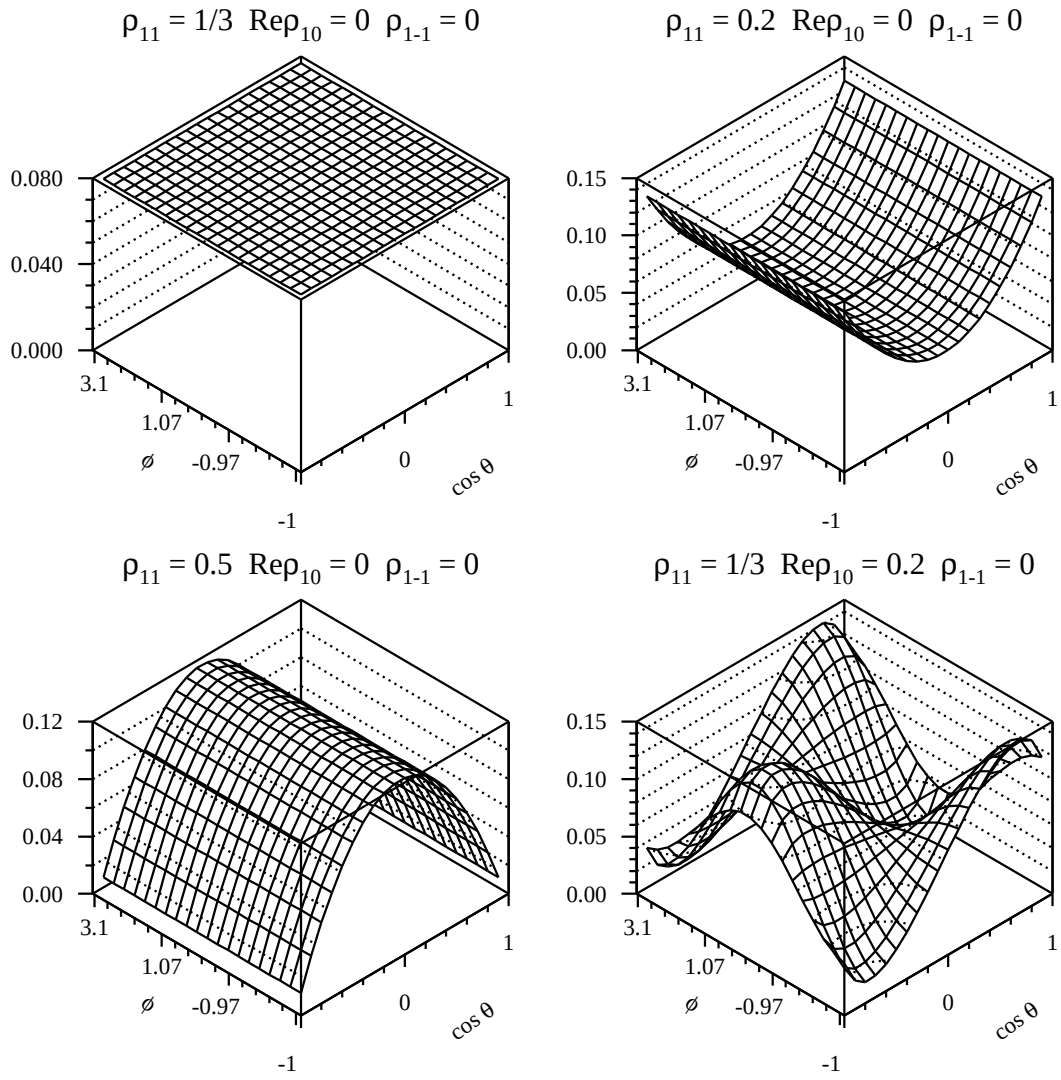

Figure 9.2: Angular dependence of $d\sigma/d\Omega$ for various values of spin density matrix elements. (Plot 1 of 2)

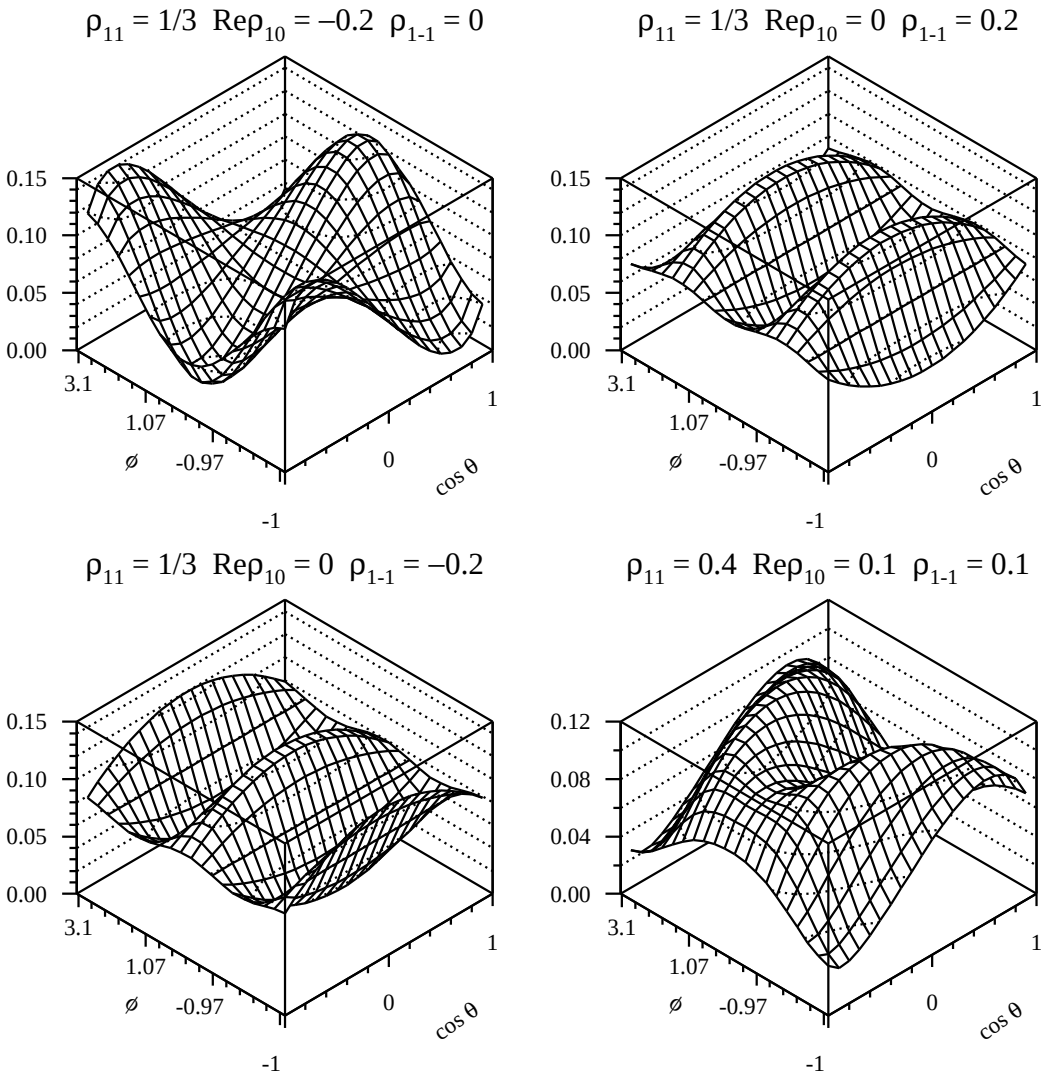

Figure 9.3: Angular dependence of $d\sigma/d\Omega$ for various values of spin density matrix elements. (Plot 2 of 2)

probability) then $\rho_{ij} = (1/3)\delta_{ij}$, so that the three parameters are $\rho_{11} = 1/3$, $\mathrm{Re}\,\rho_{10} = 0$, and $\rho_{1-1} = 0$. The functional form in Equation (9.5) is illustrated in Figures 9.2 and 9.3 where it is plotted for various values of the parameters. For the unpolarized case discussed above, the distribution is flat, as expected. The remaining plots show the effect of one parameter being different from the unpolarized value, and the last plot shows a more general case where all three parameters differ from their unpolarized values.

9.2 Measurement of Spin Density Matrix Elements

Measuring the spin density matrix elements involved two steps. First, angular distribution histograms were created, and second, these histograms were fit with Equation (9.5) to obtain the three measurable spin density matrix elements.

The angular distribution histograms are two dimensional histograms in which the numbers of events falling in bins of $\cos\theta$ and ϕ are plotted versus $\cos\theta$ and ϕ.

As before, the data sample was divided into bins of x_F and p_T^2. Because the subsamples needed to be further divided into bins of $\cos\theta$ and ϕ, fewer x_F and p_T^2 bins were used than in Chapters 7 and 8 so as not to reduce the statistics in the angular distribution histograms to a level where bin populations no longer approximate the values of a probability density function. For this reason as well, non-uniform bins in x_F and p_T^2 were chosen such that each bin contained roughly the same number of raw data events. The bins are shown in Tables 9.1 and 9.2.

For each of these bins, an 8×8 two dimensional histogram was created by uniformly dividing both the regions $-1 \leq \cos\theta \leq 1$ and $-\pi \leq \phi \leq \pi$ each into 8 equal parts. To efficiently calculate the bin populations, a bin by bin background subtraction technique was employed. In each of these bins, Q-value histograms were created

Table 9.1: x_F bins used in angular distribution study.

Bin	Lower limit	Upper limit
1	-0.1	0.05
2	0.05	0.09
3	0.09	0.13
4	0.13	0.19
5	0.19	0.6

Table 9.2: p_T^2 bins used in angular distribution study.

Bin	Lower limit $(\mathrm{GeV}/c)^2$	Upper limit $(\mathrm{GeV}/c)^2$
1	0.	0.025
2	0.025	0.0.175
3	0.175	0.75
4	0.75	3.5
5	3.5	10.

from both the real data and event-mixed data samples (see Section 7.2). Events were counted both above and below a Q-value of 10 MeV/c^2 in both histograms. A comparison of the number of events in the data and mixed data sample in the background region ($Q > 10$ MeV/c^2) yielded a scale factor which was used in subtracting the number of background events from the signal region ($Q < 10$ MeV/c^2). In other words, if S_1 denotes the number of events counted in the data sample with a Q-value less than 10 MeV/c^2, S_2 denotes the number of events from the same sample with a Q-value greater than 10 MeV/c^2, and B_1 and B_2 denote the corresponding quantities obtained from the mixed data sample, then the number of signal events could be computed as

$$S = S_1 - B_1 \frac{S_2}{B_2}. \tag{9.7}$$

These S values, and their errors were then used to fill the 2-d histograms. (Errors

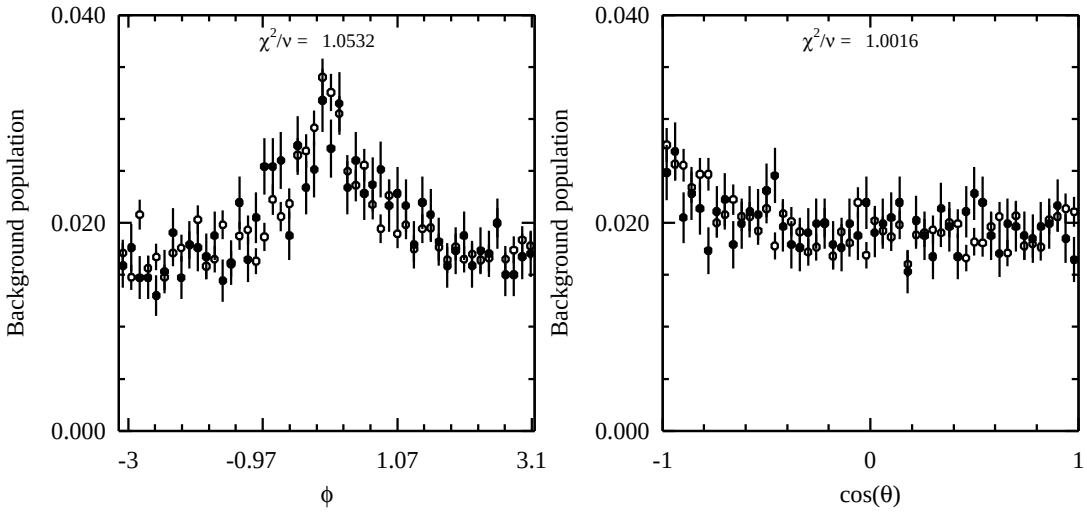

Figure 9.4: Comparison of $\cos\theta$ and ϕ distributions for background regions to the left (solid circles) and right (hollow circles) of D^* signal. The chi-squares per degrees of freedom shown are for the left and right regions of those one-dimensional projections to match. The chi-square per degrees of freedom for the full 2-dimensional histograms (not shown) to match is 1.02.

were computed assuming a $\sqrt{n}$ error on each of the counted values S_1, B_1, S_2, and B_2, and assuming these errors to be uncorrelated.)

This technique assumes that the background shape is modeled accurately by event mixing in each bin of $\cos\theta$ and ϕ. This can be checked by comparing the $\cos\theta$ and ϕ distributions of the sidebands to the left and right of the signal. If they agree, it is likely that the background events under the signal have the same shape. This would imply that variations in the background shape among the $\cos\theta$ and ϕ bins are modeled correctly by event mixing. The left sideband was defined as $Q < 3.5$ MeV/c^2, and the right sideband by $10 < Q < 15$ MeV/c^2. Two dimensional histograms were created in these two sidebands and compared. The chi-square per degree of freedom for them to agree was 1.02, suggesting a good match. In Figure 9.4, projections of

these histograms on the $\cos\theta$ and ϕ axes are shown, with the two sideband regions overlaid. The agreement appears to be good, consistent with the good chi-square.

The procedure used to create the data histograms was also applied to the Monte Carlo sample, and those resulting histograms were divided by the generated distributions to obtain binned acceptance functions in $\cos\theta$ and ϕ. The data histograms were divided by the acceptance histograms to obtain acceptance corrected histograms. (See Figure 9.5.)

These two dimensional histograms were fit with Equation (9.5) to obtain values for the spin density matrix elements on which it depends. It is difficult to show two dimensional functions overlaid on two dimensional histograms, so one dimensional projections of the histograms and fits are shown in Figure 9.6. The χ^2/ν for the full two dimensional fits range from 0.76 to 2.73, with an average value of 1.46.

The resulting spin density matrix elements are plotted in Figure 9.7. The results are clearly consistent with no polarization. (The dashed line in the upper plots is at $\rho_{11} = 1/3$, the unpolarized value.) There are fluctuations, of course, from bin to bin, but no clear trends as functions of x_F or p_T^2.

As a consistency check, the results are shown separately for the two decay modes of the D^0 in Figure 9.8. The results are consistent, with no trends in either case.

Finally, there may be physical reasons for D^{*+} and D^{*-} to be individually polarized when total sample shows no average polarization. The results from the two D^* charges are shown in Figure 9.9. Again, no polarization or trends were observed.

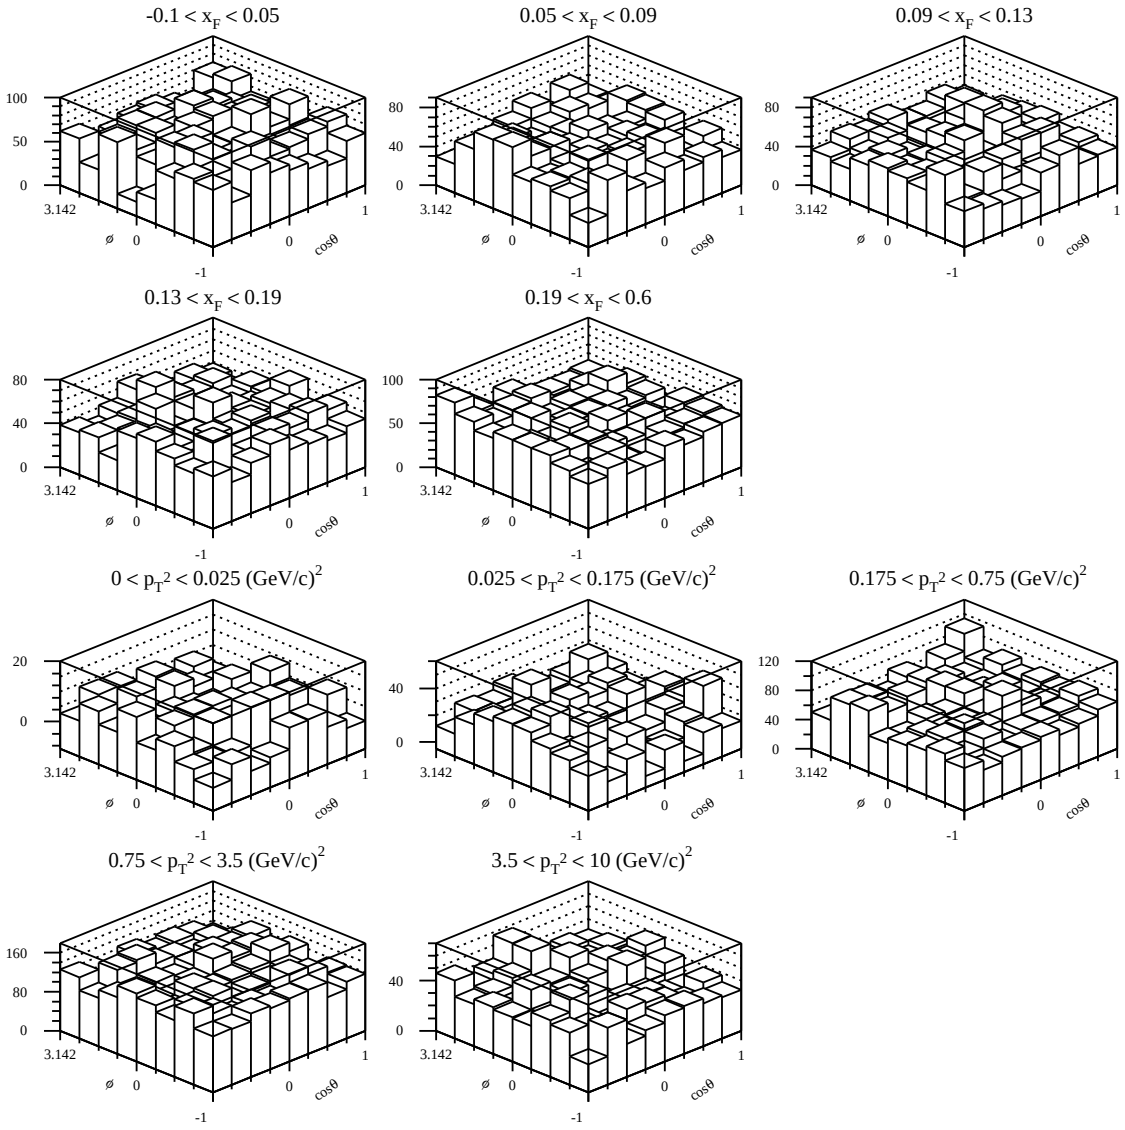

Figure 9.5: Acceptance corrected angular distribution histograms.

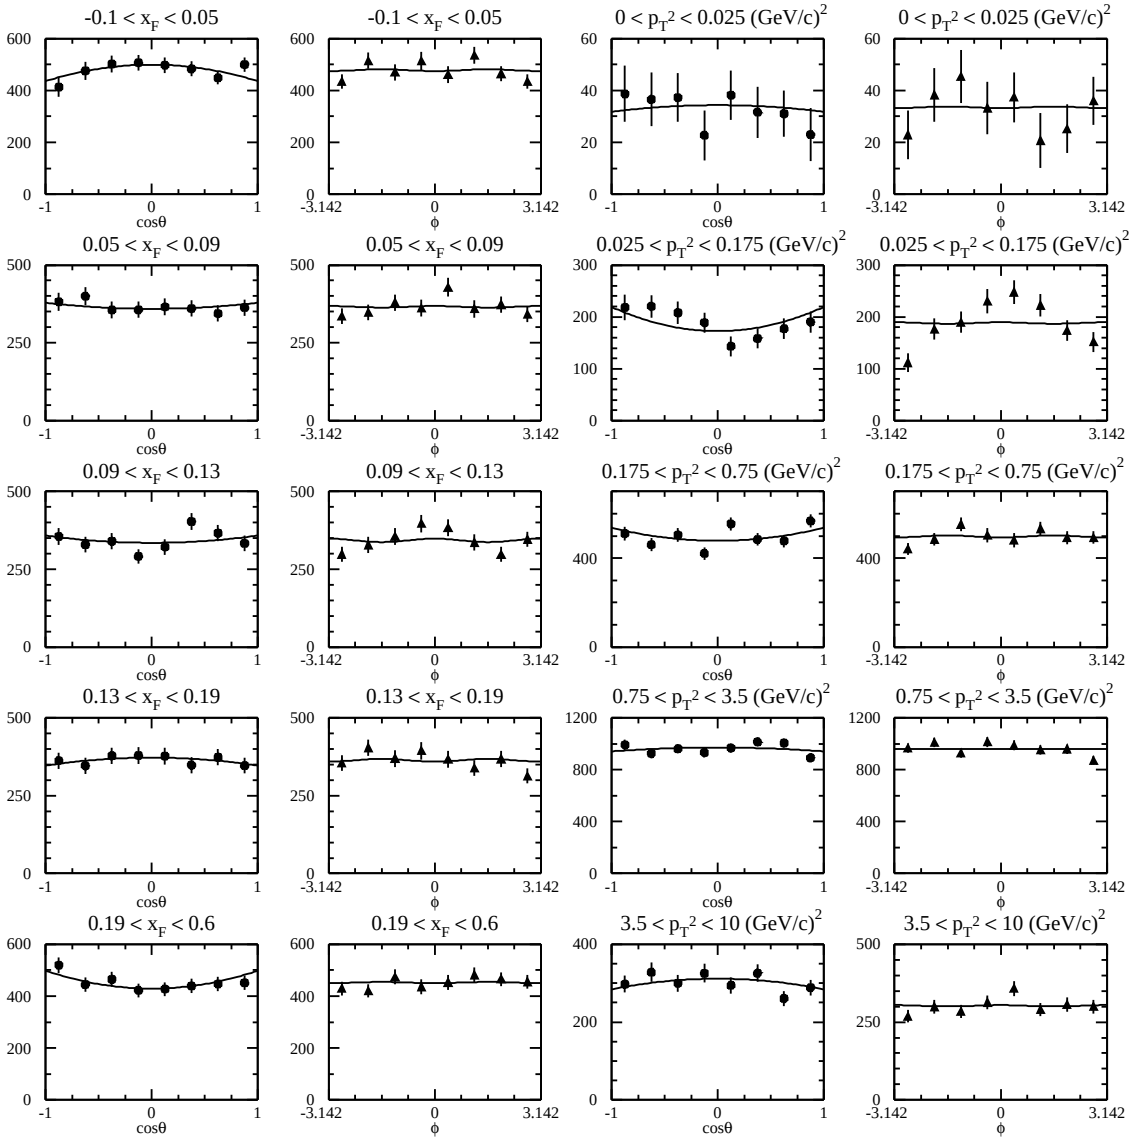

Figure 9.6: Projections of acceptance corrected 2-d histograms on $\cos\theta$ and ϕ axes along with projections of fits of Equation (9.5).

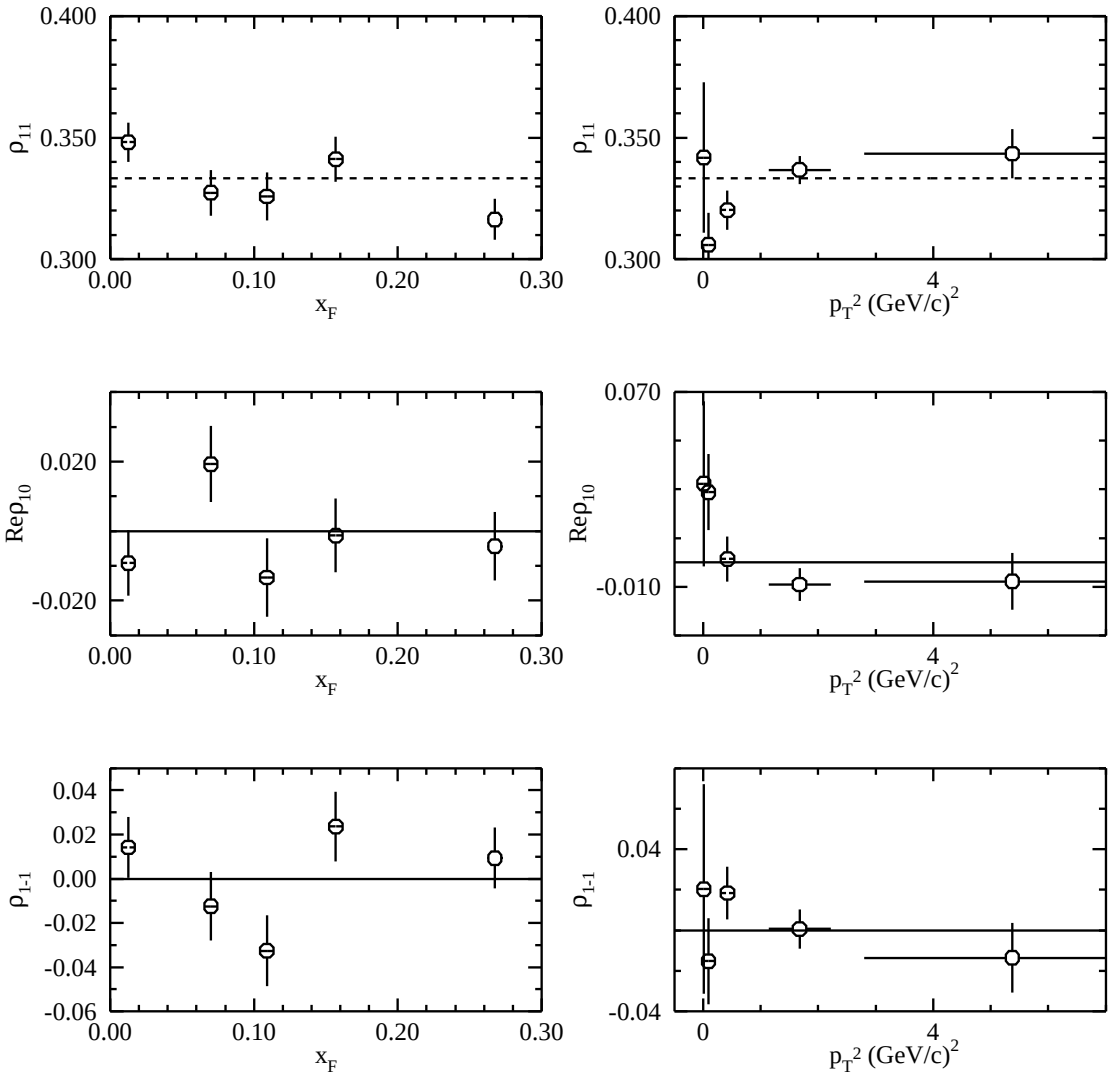

Figure 9.7: Spin density matrix elements as functions of x_F and p_T^2 for full data sample. Error bars include statistical and systematic errors added in quadrature.

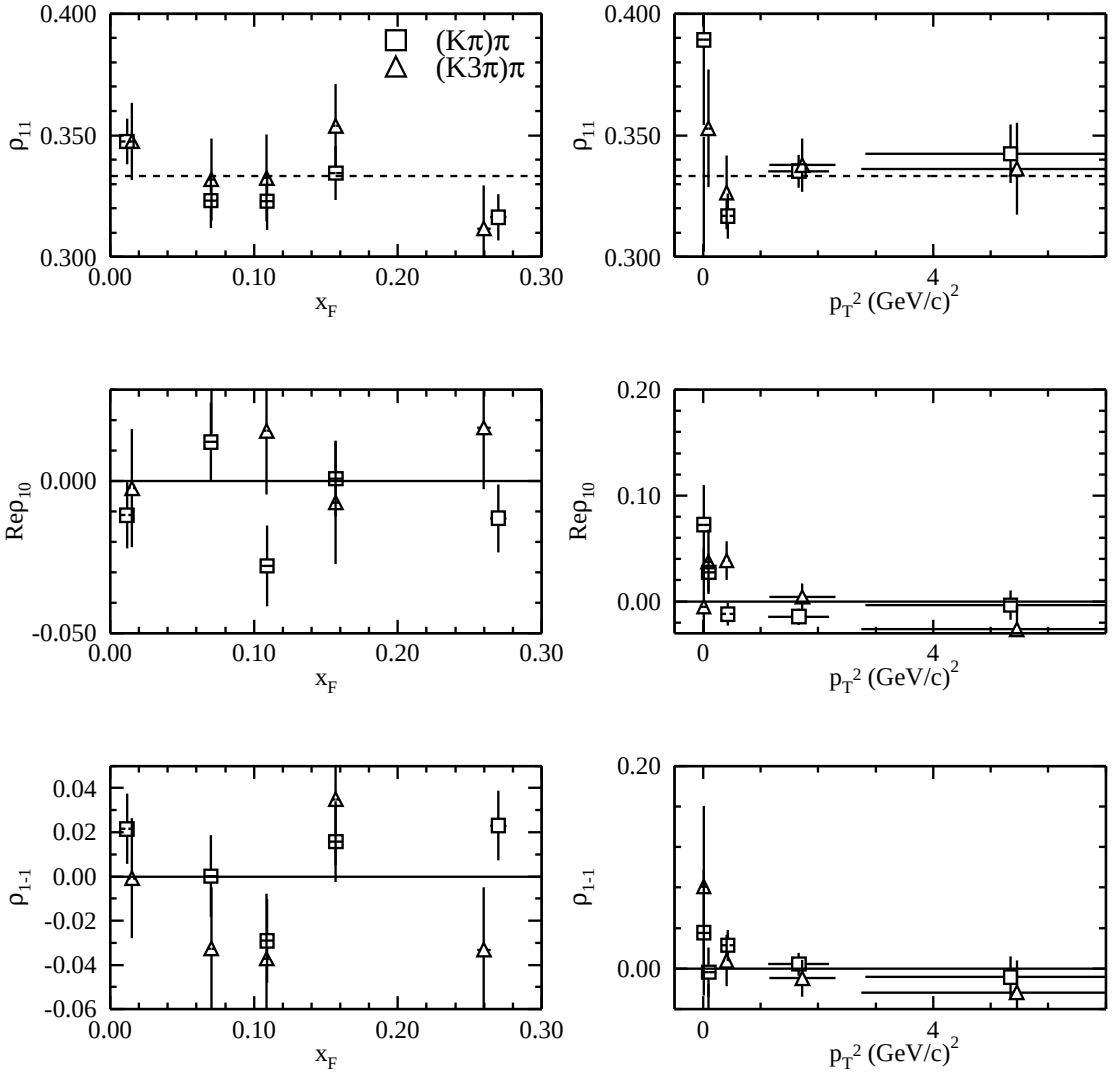

Figure 9.8: Spin density matrix elements as functions of x_F and p_T^2 comparing $(K3\pi)\pi$ and $(K\pi)\pi$ modes. Error bars include statistical and systematic errors added in quadrature.

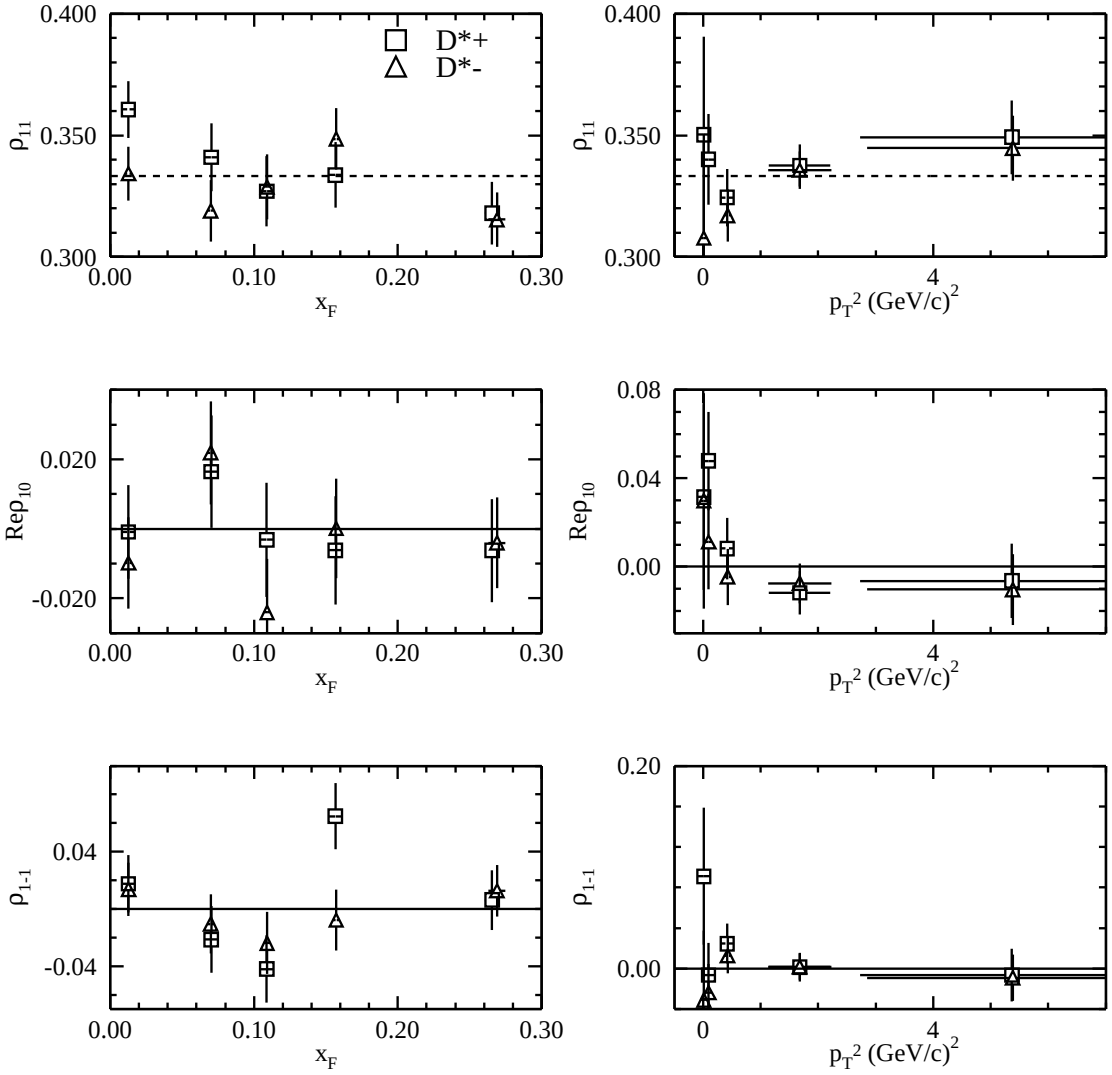

Figure 9.9: Spin density matrix elements as functions of x_F and p_T^2 with D^{*+} and D^{*-} shown separately. Error bars include statistical and systematic errors added in quadrature.

9.3 Spin Alignment

In the previous section, it was shown that no polarization is observed at the level of sensitivity of the experiment. It is useful to be able to quantify that result, that is to put a numerical upper limit on the polarization.

A quantity employed by other experiments [Ku91] is the *spin alignment parameter*. Using our notation for the spin density matrix elements, it has the form:

$$\eta = 1 - 3\rho_{11} \tag{9.8}$$

Clearly $\eta = 0$ in the unpolarized case, so the spin alignment can be thought of as a measurement of the amount of polarization. More specifically, a positive value of η corresponds to an excess of particles in the $J_z = 0$ helicity state above that found in the $J_z = \pm 1$ states.

η was computed by averaging the values of ρ_{11} in both the x_F and p_T^2 histograms which yielded the same central values and error bars, an important consistency check. Because there is no reason to assume the same values for leading and non-leading particles, separate calculations were done for D^{*+} and D^{*-}. The results obtained for the average values of the spin alignment are: $\langle\eta\rangle = 0.01 \pm 0.02$ for the D^{*-} and $\langle\eta\rangle = -0.01 \pm 0.02$ for the D^{*+}.

Although this is a first attempt at such a measurement of D^* spin alignment in a hadroproduction experiment, numerous measurements have been performed on D^*s produced in e^+e^- collisions [Br98], most of which are consistent with no polarization. Some experiments [Ab87, Ac97], however, have measured a positive value for η.

Chapter 10

Systematic Errors

10.1 Introduction

Associated with any measurement of a physical quantity is a range of uncertainty, or *error*. Such errors consist of two components: a *statistical error* and a *systematic error*.

Interactions in high energy physics are governed by the laws of quantum mechanics. The fundamental laws give rise to probability distributions, and thus by making individual measurements, one is sampling probability density functions. The physical quantities ultimately measured thus have errors arising from sampling statistics. These errors are proportional to $1/\sqrt{n}$ where n is the number of measurements taken. In high energy physics experiments, n corresponds to the number of signal events used in the final analysis.

But there are many other sources of error which may scale differently than $1/\sqrt{n}$. These systematic errors arise from imperfections in the measurement process itself. Examples can include inaccuracies in the acceptance functions resulting from imperfect detector simulations, fitting inaccuracies resulting from imperfect modeling of the

signal or background shape, incorrect modeling of the time dependence of detector performance, etc. These errors can be more difficult to quantify than statistical ones and require careful study.

10.2 Determination of Systematic Errors

The technique employed to quantify systematic errors was to compare results computed in different ways and examine the differences. Differences were studied bin by bin, and reported as fractions of the statistical errors on these bins. The root mean squares of these fractions over all bins were used as an approximation of the systematic errors in order to reduce the effect of bin by bin fluctuations. This technique is only valid, of course, if the ratio of systematic error to statistical error is approximately constant over all bins. It will be shown that this assumption is justified for nearly all types of systematic errors studied. Systematic errors are thus reported as percentages of statistical errors. Finally, these percentages are added in quadrature to determine an overall factor by which to scale the statistical errors to obtain the overall error.

In subsequent subsections, different sources of systematic errors are discussed. For each of these, the differences in results divided by the statistical error are plotted in Figures 10.1 and 10.2. The RMS values of these ratios taken over all x_F and p_T^2 bins are listed in Table 10.1.

10.2.1 D^0 Mass Resolution

In order to reduce the background level, a cut was placed on the candidate D^0 mass at 2.5σ on either side of the mean. (See Section 7.3.) If the Monte Carlo modeling of the D^0 mass resolution wasn't perfect, this could lead to errors in the acceptance

function and hence the resulting distributions. To ascertain the sensitivity to this cut, results were compared using a cut at 2.5σ and 3σ.

10.2.2 Monte Carlo Production Model

As discussed in Section 6.4.3, it was necessary to weight the Monte Carlo events to correct for the discrepancy between the observed x_F and p_T^2 distributions of D^*s in data and that in the Monte Carlo. The correction curves could not be determined perfectly, due both to limited statistics in data, and to simplified models (low order polynomials) for the curve shapes.

To estimate the effect on the results, the p_T^2 weighting curve parameters were varied by 1σ from the values obtained by the fit. Results were compared with the weighting curve as used in the analysis.

Since the x_F weighting curve was less significant, results were simply compared to results where no x_F weighting curve was employed.

10.2.3 Tracking Efficiency

As discussed in Section 7.5, a comparison of results for the two decay modes of the D^0 ($K\pi$ and $K3\pi$), provided a means of checking the modeling of tracking efficiency. Results obtained from the subsamples corresponding to the two modes were therefore compared.

This comparison was somewhat complicated by the fact that the sample sizes are different and the samples are statistically independent. To address the first problem, corrected yields in the $K3\pi$ mode were multiplied by the branching ratio $\Gamma(D^0 \to K^-\pi^+\pi^+\pi^-)/\Gamma(D^0 \to K^-\pi^+) = 1.96$ obtained from the Particle Data Book [Ha02,

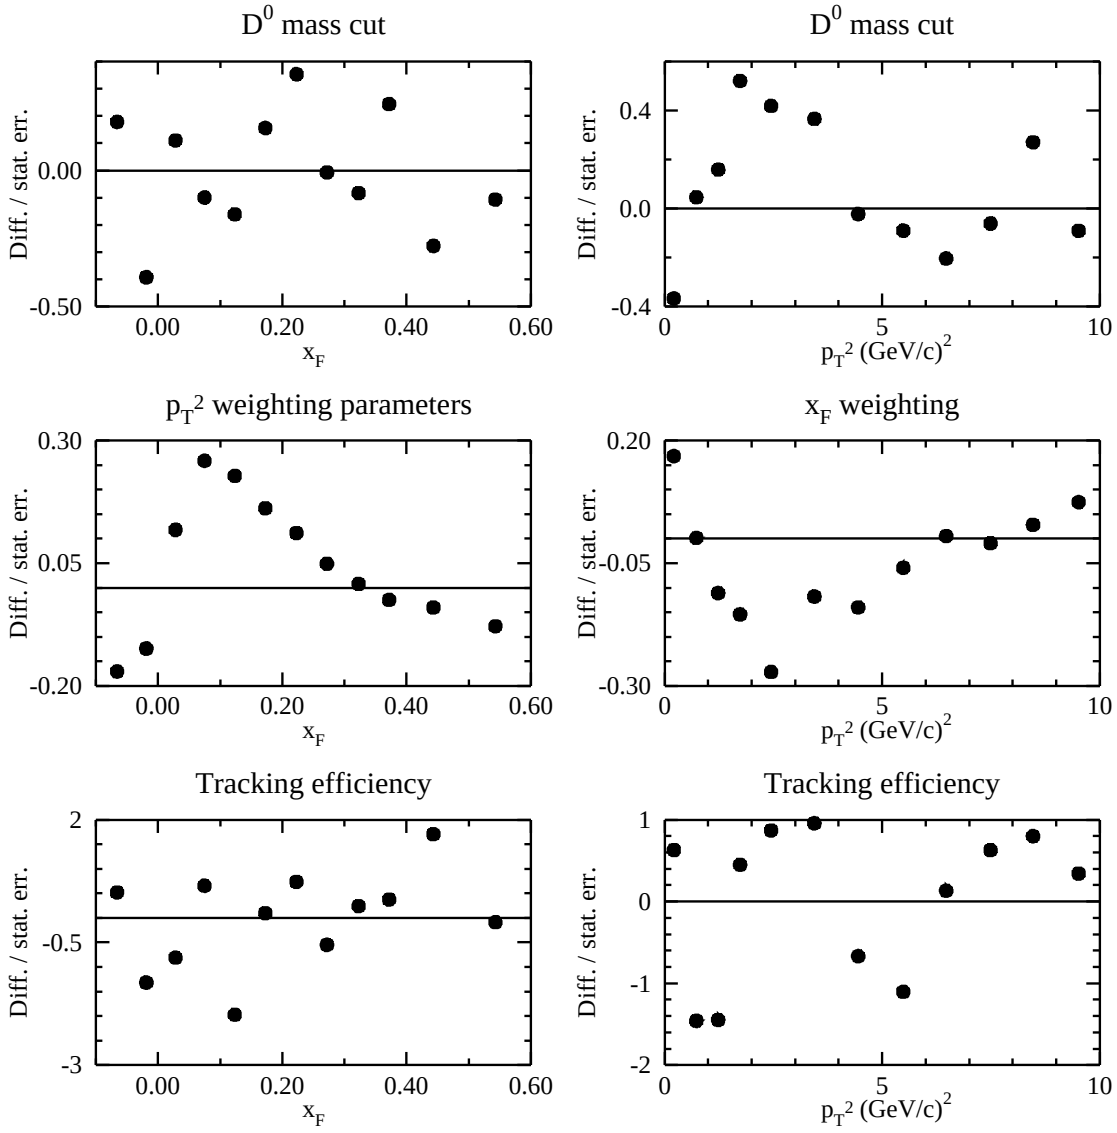

Figure 10.1: Differences divided by statistical errors for various systematic checks as functions of x_F and p_T^2. (Plot 1 of 2)

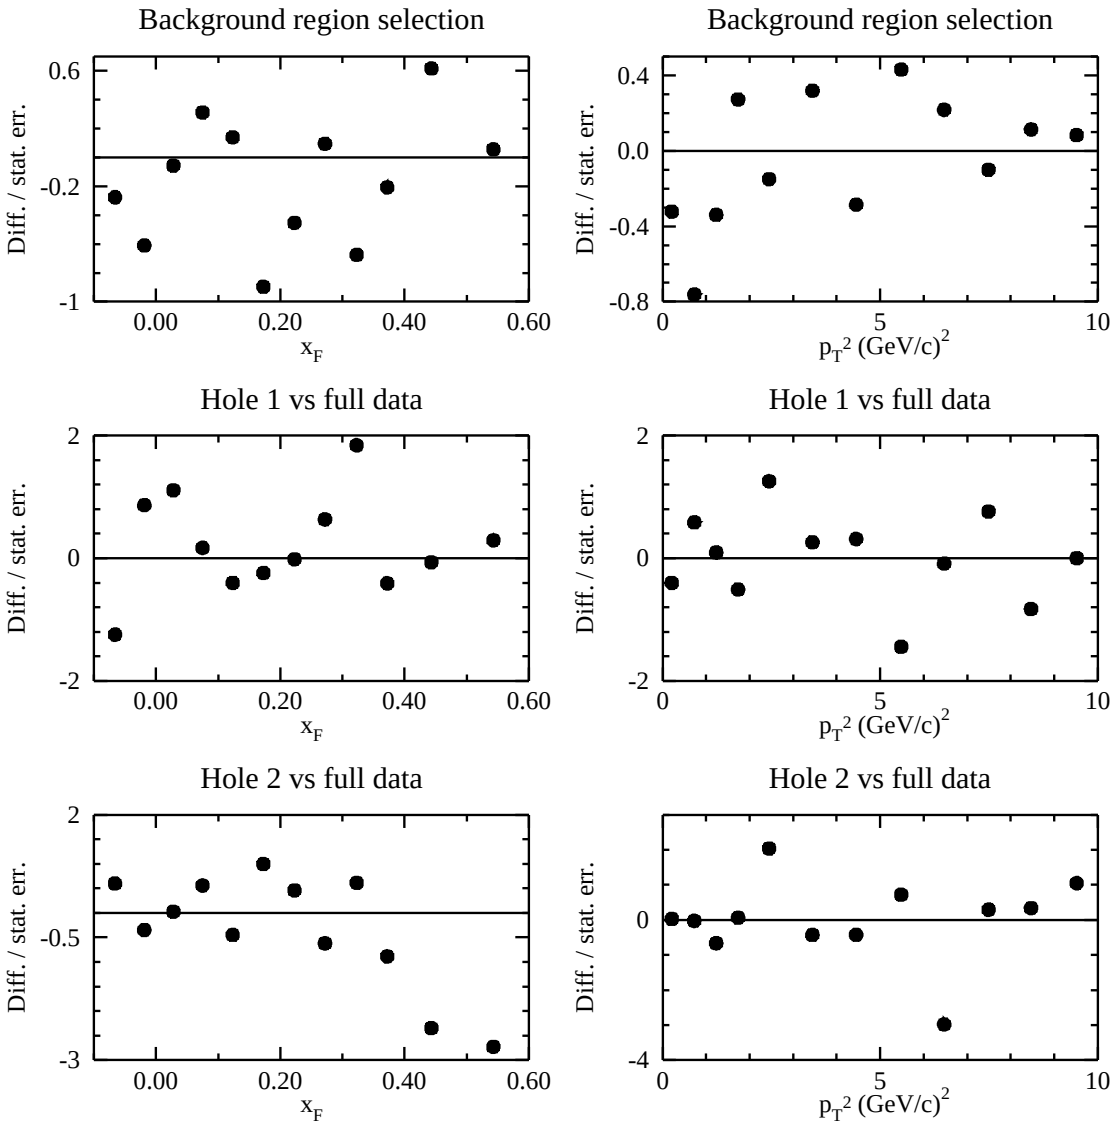

Figure 10.2: Differences divided by statistical errors for various systematic checks as functions of x_F and p_T^2. (Plot 2 of 2)

p. 613]. Comparisons yielded RMS ratio values of 0.955 and 0.881 over x_F and p_T^2 bins, respectively.

Unlike other comparisons discussed in this chapter which involved comparisons of the same data set analyzed in different ways, this comparison involved two statistically independent data samples. For differences of the former kind the RMS values obtained in the differences divided by statistical errors would hopefully be smaller than one, and indeed this is the case for all such errors studied. The situation is more complicated in the present case, however.

If there is no systematic difference between two statistically independent samples, we would expect this RMS value to be close to unity, and if there is a systematic difference, the RMS value should be larger. What then do we make of the present result, where the samples are independent and the RMS values are both less than one? It might appear that this suggests our method of computing statistical errors is incorrect and they are being overestimated, so this was investigated as follows.

The sample was split into two halves, based on a random number generator with a uniform distribution. For each event a random number between 0 and 1 was generated, and if it was less than 0.5 the event was included in sample A, otherwise it was included in sample B. This created two statistically separate samples of equal size with no systematic differences between them. A comparison between the x_F and p_T^2 distributions was done in the same way used to generate the plots in this section. The root mean square values of the differences divided by the statistical errors over all bins of x_F and p_T^2 was calculated, again as in the other errors studied. This experiment was repeated for a total of 100 times, and the resulting RMS values were plotted. The resulting plots are shown in Figure 10.3.

The results have a mean of about 1.0 as expected, and the spread is large enough

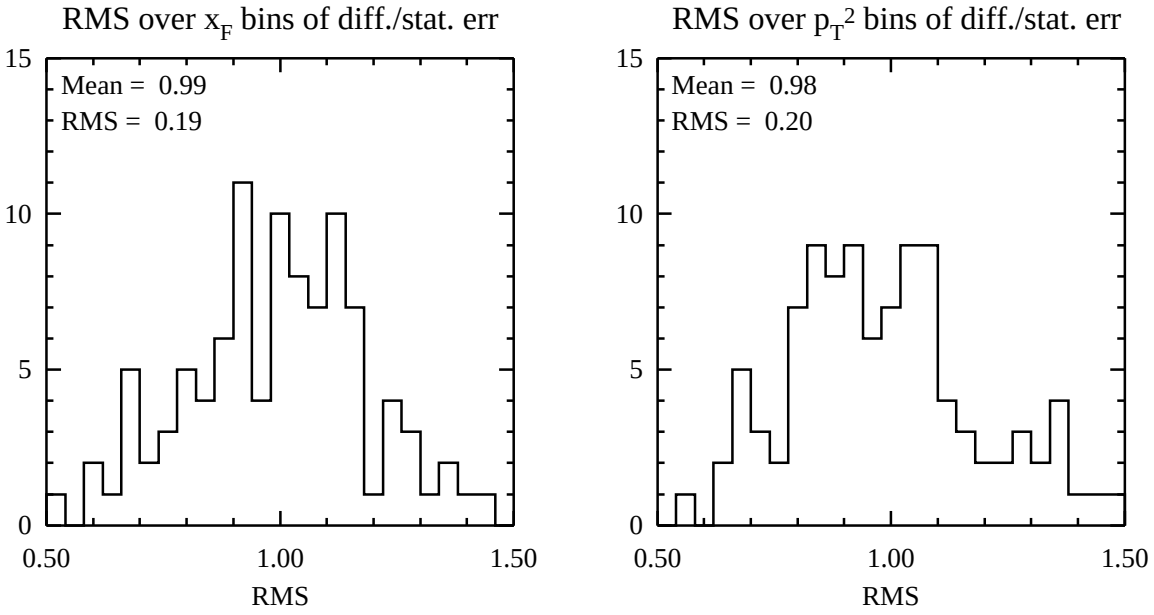

Figure 10.3: Distributions of RMS values of differences divided by statistical errors over x_F and p_T^2 bins between statistically independent subsamples of equal size.

to explain the low RMS values obtained in the tracking efficiency study. Therefore, there doesn't seem to be a problem with our statistical error calculations.

This large width in the RMS plots does, however, make it difficult to extract a systematic error when comparing two independent samples for which the systematic error does not dominate the statistical error. In these cases, a systematic error equal to 25% of the statistical error was assigned, as it is about the largest error which could produce comparison results consistent with no error. Thus a value of 0.25 is shown in Table 10.1 for the tracking efficiency.

10.2.4 Fitting Procedure

Since event mixing was used to model the background, and the signal was not fit at all, the only variable involved was the definition of the signal and background regions.

Recall that the signal region was defined as $Q < 12$ MeV/c^2 and the background region was defined as $Q > 12$ MeV/c^2. The choice of 12 MeV/c^2 as the cutoff point was arbitrary, so results were compared with those obtained using a cutoff of 14 MeV/c^2.

10.2.5 Hole Modeling

As explained in Section 3.7, over the period of E791's run, the 500 GeV π^- beam caused a significant amount of ionization in certain regions of the drift chambers, which caused deposits to gradually form on the drift chamber wires in the regions through which the beam passed. This eventually caused a loss in tracking efficiency for decay tracks which passed through those same regions. These drift chamber "holes" were modeled by the Monte Carlo in a time dependent fashion by dividing the run period into five parts and modeling the hole separately for these parts. In practice the full data set was corrected by the average hole position and size. As a systematic check, the full data set corrected with the average hole was compared to various subsets corresponding to one of the five run periods and corrected with the appropriate hole model. The first two comparisons are plotted in Figure 10.2. A scaled version of the hole 2 difference was used in assigning a systematic error. This technique was motivated by previous E791 analyses which did similar studies.

10.3 Summary

Table 10.1 summarizes the systematic errors as fractions of statistical errors. Recall that the assumption made was that these fractions are approximately constant, independent of the statistical error itself. Examination of the plots of these ratios (Figures 10.1 and 10.2) shows that for all but one of the sources of error, the x_F and p_T^2 Monte Carlo weighting, the points seem to be randomly scattered about the

Source	Syst / Stat err ratio (x_F)	Syst / Stat err ratio (p_T^2)
D^0 mass cut	0.211	0.27
p_T^2/x_F weighting	0.137	0.123
Tracking efficiency	0.25	0.25
D^* fitting	0.453	0.334
Hole modeling	0.62	0.55
Total	0.85	0.75

Table 10.1: Summary of systematic error fractions.

horizontal axis and show no trends as functions of statistical error. The x_F weighting and even more so the p_T^2 weighting show a strong dependence. The ratios become larger as the statistical error becomes smaller, consistent with a constant error rather than one which scales with the statistical error. The overall magnitude of this error, however, is small compared with the other errors, and is in fact the smallest error reported in Table 10.1, and therefore it was not considered problematic to treat it like the other sources of error.

An overall error of 80% of the statistical error was assigned to the systematic error. Adding this in quadrature to the statistical error means that all statistical error bars need to scaled by a factor of 1.3 to take into account the systematic error. All the plots in the analysis chapters which indicate that systematic errors are included used this scale factor. Errors reported on fits to these plots thus include a systematic as well as statistical component.

Chapter 11

Conclusions

Differential cross sections for the production of $D^{*\pm}(2010)$ mesons in 500 GeV/c π^-–nucleon interactions as functions of Feynman-x (x_F) and square transverse momentum (p_T^2) have been reported. Fits to the traditional forms given by Equations (7.4) and (7.5), show poor agreement with data, except in limited regions of x_F, but nowhere for p_T^2. The PYTHIA model shows good agreement with the x_F distribution, but with default parameters, poor agreement with our p_T^2 distribution which is much stiffer than the model. A tuned PYTHIA model with an increased charm quark mass, increased intrinsic k_T, and other changes fits the data much better. The best model for our p_T^2 distribution, however, is Equation 7.6, which fits our data over the entire range with a chi-square per degrees of freedom of almost exactly one. This equation was suggested as it fit next-to-leading order QCD predictions very well, so the good fit suggests that perturbative QCD is very effective at describing p_T^2 distributions. Also, the good fit of PYTHIA to the x_F distributions suggests this as well, since PYTHIA employs leading order QCD along with parton showers as a production model. Fragmentation is believed to have a greater effect on x_F distributions than p_T^2 distributions, which

makes simple comparisons with direct NLO QCD models as in Equation (7.6) difficult in the x_F case.

The charge asymmetry parameter (Eq. (8.1)) was also reported as functions of both x_F and p_T^2. As expected, there is strong evidence for the leading particle effect with the asymmetry rising steeply at high x_F. PYTHIA models, both tuned and untuned, predict a leading particle effect considerably greater than is seen in our data. The $D^{*\pm}$ asymmetry data is consistent with E791 $D^{\pm}$ asymmetry data, whereas PYTHIA predicts more asymmetry in the D^* case. A model incorporating both fragmentation and recombination provides a much better fit to both the D^* and D data. This suggests the Lund string fragmentation model alone is not sufficient to describe these asymmetries. The asymmetry is approximately constant as a function of p_T^2; while it dips to zero near $p_T^2 = 0$, data points in this region are consistent with the average value.

Spin density matrix elements for the D^*s have been reported also as functions of x_F and p_T^2, and the average value of the spin alignment parameter η was calculated as well. We see no evidence of polarization, either on average or as a function of x_F or p_T^2. An upper limit on the polarization is provided by our measurements of the spin alignment parameter: $\eta = 0.01 \pm 0.02$ for leading particles, and $\eta = -0.01 \pm 0.02$ for non-leading particles. This suggests that mechanisms which give rise to polarization in the hadroproduced hyperons are not present in the meson case, or the effects are possibly washed out by fragmentation.

With its large data set, E791 has proven an effective testing ground for charm production models, and for the applicability of perturbative QCD in heavy quark production. Having the largest hadroproduced charm sample in the world enabled us to significantly reduce error bars over those in previous measurements of the same

phenomena. Also, it allowed a first investigation of possible polarization in hadroproduced heavy quark mesons.

The attention of the high energy physics community is now focused elsewhere, particularly on B physics. E791 was, in a sense, the last of its kind. In addition to providing interesting physics in its own right, studies like the present one can also be used to tune Monte Carlo models which describe the backgrounds in B experiments, where significant amounts of charm may be produced. Understanding these backgrounds accurately is crucial to successful analysis in the B sector. We believe that the results presented here, as well as other E791 analyses, and the E791 data set itself, will be very useful in a rich variety of ways for many years to come.

Appendix A

Angular Distributions in the Production and Decay of a Spin-1 Particle

In this appendix, the angular distribution of decay products of a spin-1 particle produced in a parity-conserving interaction and decaying into two spin-0 particles, is derived and related to the spin density matrix elements of the spin-1 particle in a helicity basis. This result is applicable to the following system

$$\pi^- + N \rightarrow D^{*+} + X \tag{A.1}$$

$$D^{*+} \rightarrow D^0 + \pi^+ \tag{A.2}$$

studied in the analysis described in this document. The derivation follows that given by Chung [Ch71], but it can be found in many texts.

A.1 Notation and Definitions

The behavior of one and two particle states under rotations, boosts, and parity will be needed, so the notation used here to describe these operations must first be defined.

A.1.1 Rotations

The symbol $R(\alpha, \beta, \gamma)$ denotes a rotation by the Euler angles (α, β, γ).[1] Clearly, any unit vector $\hat{\beta}$ having direction (θ, ϕ) in a spherical polar coordinate system (with polar angle θ and azimuthal angle ϕ), is given by:

$$\hat{\beta} = R(\phi, \theta, 0)\hat{z} \tag{A.3}$$

where $\hat{z}$ is a unit vector pointing along the positive z-axis.

Corresponding to every rotation R is a unitary operator $U[R]$ which acts on state vectors. Clearly the group property

$$U[R_1 R_2] = U[R_1]U[R_2] \tag{A.4}$$

must hold for arbitrary rotations $R_1 = R(\alpha_1, \beta_1, \gamma_1)$ and $R_2 = R(\alpha_2, \beta_2, \gamma_2)$. This unitary operator is given by:

$$U[R(\alpha, \beta, \gamma)] = e^{-i\alpha J_z} e^{-i\beta J_y} e^{-i\gamma J_z} \tag{A.5}$$

Where J_i denotes the i-th component of the angular momentum operator.

Let $|jm\rangle$ denote a system of total angular momentum j and z-component of

[1] In this formalism, the system is first rotated by an angle γ about the z axis, followed by a rotation β about the y axis, and finally a rotation α about the z axis. Note that these are "active" rotations: the system is rotated, not the coordinate axes.

angular momentum m. Clearly j is invariant under any rotations, and so a rotated system must be a linear combination of states $|jm'\rangle$. Using the notation of Rose[Ro57]:

$$U[R(\alpha,\beta,\gamma)]|jm\rangle = \sum_{m'=-j}^{j} D^j_{m'm}(\alpha,\beta,\gamma)\,|jm'\rangle \tag{A.6}$$

where

$$\begin{aligned} D^j_{m'm}(\alpha,\beta,\gamma) &= \langle jm'|U[R(\alpha,\beta,\gamma)]|jm\rangle \\ &= \langle jm'|e^{-i\alpha J_z}e^{-i\beta J_y}e^{-i\gamma J_z}|jm\rangle \\ &= e^{-im'\alpha}\,d^j_{m'm}(\beta)\,e^{-im\gamma} \end{aligned} \tag{A.7}$$

and

$$d^j_{m'm}(\beta) = \langle jm'|e^{-i\beta J_y}|jm\rangle. \tag{A.8}$$

Tables of d-functions can be found in the literature[Ha02, p. 245]. Two properties of the D-functions will be needed. The first follows immediately from the group property (A.4):

$$\begin{aligned} D^j_{m'm}(R_2R_1) &= \langle jm'|U[R_2]U[R_1]|jm\rangle \\ &= \sum_k \langle jm'|U[R_2]|jk\rangle\langle jk|U[R_1]|jm\rangle \\ &= \sum_k D^j_{m'k}(R_2)D^j_{km}(R_1) \end{aligned} \tag{A.9}$$

Secondly, from the unitarity of $U[R]$:

$$\begin{aligned} D^{j*}_{m'm}(R) &= \langle jm'|U[R]|jm\rangle^* \\ &= \langle jm|U[R]^\dagger|jm'\rangle \end{aligned}$$

$$= \langle jm|U[R^{-1}]|jm'\rangle$$
$$= D^{j}_{mm'}(R^{-1}) \tag{A.10}$$

A.1.2 Boosts

The notation $L(\vec{p})$ is used to denote a Lorentz transformation (boost) that takes a particle of mass m at rest to a frame in which the particle has momentum $\vec{p}$. In the special case where the boost direction is along the z-axis, the notation $L_z(p)$ is used. Clearly, if $\vec{p}$ has direction (θ, ϕ) and magnitude p, then

$$L(\vec{p}) = R(\phi, \theta, 0)\, L_z(p)\, R^{-1}(\phi, \theta, 0). \tag{A.11}$$

As before, the unitary operator corresponding to this operator is $U[L(\vec{p})]$ which satisfies a group property similar to that for rotations. Thus,

$$U[L(\vec{p})] = U[R(\phi, \theta, 0)]\, U[L_z(p)]\, U^{-1}[R(\phi, \theta, 0)]. \tag{A.12}$$

A.1.3 Parity

Finally, we define the parity operator $\mathcal{P}$, which reverses the direction of the spatial coordinates:

$$\mathcal{P} : (x, y, z) \rightarrow (-x, -y, -z) \tag{A.13}$$

Clearly position and momentum vectors change sign under $\mathcal{P}$:

$$\mathcal{P}(\vec{r}) = -\vec{r} \qquad \mathcal{P}(\vec{p}) = -\vec{p} \tag{A.14}$$

but angular momentum $\vec{J} = \vec{r} \times \vec{p}$ doesn't change:

$$\mathcal{P}(\vec{J}) = \vec{J}. \tag{A.15}$$

And thus, if we let $\Pi = U[\mathcal{P}]$ denote the unitary operator associated with $\mathcal{P}$, then for states of definite angular momentum

$$\Pi|jm\rangle = \eta|jm\rangle \tag{A.16}$$

where $\eta = \pm 1$ is the intrinsic parity of the state.

A.2 Helicity

Helicity states are angular momentum states in which the quantization axis is chosen along the direction of the particle's momentum. Helicity states can transform differently than the canonical states under transformations which change the momentum, as the quantization axis can change. The letter λ is conventionally used to denote the component of angular momentum along $\vec{p}$. For a particle at rest, the quantization axis is generally chosen along z, so the helicity states and canonical states are identical.

A helicity state of momentum $\vec{p}$ (where $\vec{p}$ has direction (θ, ϕ)), can be constructed from a rest state $|j\lambda\rangle$ in two ways. The rest state can be rotated to the direction of $\vec{p}$ and then boosted by $L(\vec{p})$, or it can first be boosted along the z-axis and then rotated. This follows from Equation (A.12):

$$\begin{aligned} |\vec{p}, j\lambda\rangle &= U[L(\vec{p})]\, U[R(\phi, \theta, 0)]\, |j\lambda\rangle \\ &= U[R(\phi, \theta, 0)]\, U[L_z(p)]\, |j\lambda\rangle. \end{aligned} \tag{A.17}$$

Next, we consider the problem of constructing two particle helicity states, which will be needed later. Let s_i, w_i, and $\vec{p}_i$ denote the spin, mass, and momentum of the i-th particle where $i \in \{1, 2\}$. We work in the center of mass frame, where $\vec{p}_1 = -\vec{p}_2 = \vec{p}$. Let (θ, ϕ) represent the direction of $\vec{p}$ in a spherical polar coordinate system.

The state is constructed by first boosting the rest states of the two particles along the $\pm z$ axes, and then applying the appropriate rotation:

$$\begin{aligned} |\phi\theta\lambda_1\lambda_2\rangle &= aU[R(\phi,\theta,0)]\,\{U[L_z(p)]|s_1\lambda_1\rangle\,U[L_{-z}(p)]|s_2\,-\lambda_2\rangle\} \\ &= U[R(\phi,\theta,0)]\,|00\lambda_1\lambda_2\rangle \end{aligned} \tag{A.18}$$

where $|s_i\lambda_i\rangle$ is the rest state of particle i and the normalization constant a is chosen so that:

$$\langle\Omega'\lambda_1'\lambda_2'|\Omega\lambda_1\lambda_2\rangle = \delta^{(2)}(\Omega'-\Omega)\delta_{\lambda_1'\lambda_1}\delta_{\lambda_2'\lambda_2}. \tag{A.19}$$

Also note the minus sign appearing in the rest state for particle 2 is necessary so that its helicity quantum number is $+\lambda_2$.

Equation (A.18) gives a basis of states of definite direction (a "plane wave" basis). States of definite angular momentum J (a "spherical wave" basis) can be constructed from these as follows:

$$|JM\lambda_1\lambda_2\rangle = \frac{N_J}{2\pi}\int dR\,D^{J*}_{M\mu}(R)\,U[R]\,|00\lambda_1\lambda_2\rangle \tag{A.20}$$

where N_J is a normalization constant, and μ is to be determined later, and the integral is over all three Euler angles of R. To show that (A.20) represents a state of definite angular momentum, it suffices to show that it transforms as such under an arbitrary

rotation. So applying a rotation R':

$$U[R']|JM\lambda_1\lambda_2\rangle = \frac{N_J}{2\pi}\int dR\, D^{J*}_{M\mu}(R)\, U[R'']\,|00\lambda_1\lambda_2\rangle \tag{A.21}$$

where $R'' = R'R$. Now, using Equations (A.9) and (A.10):

$$\begin{aligned} D^{J*}_{M\mu}(R) &= D^{J*}_{M\mu}(R'^{-1}R'') \\ &= \sum_{M'} D^{J*}_{MM'}(R'^{-1})\, D^{J*}_{M'\mu}(R'') \\ &= \sum_{M'} D^{J}_{M'M}(R')\, D^{J*}_{M'\mu}(R''). \end{aligned}$$

Also, since R' is fixed, $dR = dR''$, so that

$$U[R']|JM\lambda_1\lambda_2\rangle = \sum_{M'} D^{J}_{M'M}(R')|JM'\lambda_1\lambda_2\rangle, \tag{A.22}$$

and thus the state (A.20) does transform as a state of angular momentum J.

To determine μ, we explicitly write $R = R(\phi,\theta,\gamma)$, then

$$\begin{aligned} U[R(\phi,\theta,\gamma)]\,|00\lambda_1\lambda_2\rangle &= U[R(\phi,\theta,0)]\,U[R(0,0,\gamma)]\,|00\lambda_1\lambda_2\rangle \\ &= e^{-i(\lambda_1-\lambda_2)\gamma}\,U[R(\phi,\theta,0)]\,|00\lambda_1\lambda_2\rangle \end{aligned} \tag{A.23}$$

In obtaining the last line, use is made of the fact that rotations about the z-axis commute with boosts along the same axis:

$$[R(0,0,\gamma), L_{\pm z}(p)] = 0. \tag{A.24}$$

Substituting Equation (A.23) into (A.20):

$$
\begin{aligned}
|JM\lambda_1\lambda_2\rangle &= \frac{N_J}{2\pi}\int dR\, D^{J*}_{M\mu}(\phi,\theta,0)\, e^{i\mu\gamma}\, e^{-i(\lambda_1-\lambda_2)\gamma}\, U[R(\phi,\theta,0)]\,|00\lambda_1\lambda_2\rangle \\
&= N_J\int d\Omega\, D^{J*}_{M\lambda}(\phi,\theta,0)\,|\phi\theta\lambda_1\lambda_2\rangle \qquad \text{(A.25)}
\end{aligned}
$$

where $\lambda = \lambda_1 - \lambda_2$. In the second line we have integrated over γ and in the process find $\mu = \lambda$, and in the first line, Equation (A.7) was used to obtain the dependence on γ of $D(R)$.

A.3 Spin Density Matrix

Consider the parity conserving interaction:

$$a + b \rightarrow c + J, \qquad \text{(A.26)}$$

where c may represent a multi-particle system, and J denotes a particle of spin J for which we will construct the spin density matrix. Initial state particles a and b are assumed to be in an unpolarized state.

Recall that a mixture of states $|\psi_i\rangle$, having fractional populations w_i can be characterized by the density operator:

$$\rho = \sum_i w_i |\psi_i\rangle\langle\psi_i|. \qquad \text{(A.27)}$$

We first write the density operator for the initial state of Equation (A.26) in terms of two-particle helicity states in the center of mass frame. We fix the initial momentum $\vec{p}_i = \vec{p}_a = -\vec{p}_b$ with direction $(\theta = 0, \phi = 0)$, and the final momentum $\vec{p}_f = \vec{p}_J = -\vec{p}_c$ with direction $\Omega_0 = (\theta_0, \phi_0)$. Since the initial state is unpolarized, all possible initial

helicities occur with equal weight, so the density operator becomes:

$$\rho_{(i)} \propto \sum_{\lambda_a \lambda_b} |\vec{p}_i \lambda_a \lambda_b\rangle \langle \vec{p}_i \lambda_a \lambda_b|. \tag{A.28}$$

(The proportionality constant can be chosen so that $\mathrm{Tr}(\rho) = 1$, but it will not be needed here.)

Now, let T denote the transition operator corresponding to the interaction (A.26). The final state density operator can be constructed from $\rho_{(i)}$ by applying the operator to the states appearing in the sum:

$$\rho \propto \sum_{\lambda_a \lambda_b} T\, |\vec{p}_i \lambda_a \lambda_b\rangle \langle \vec{p}_i \lambda_a \lambda_b|\, T^\dagger. \tag{A.29}$$

Now, to construct the spin density matrix corresponding to J we take the trace over all other final state variables:

$$\begin{aligned} \rho^J_{\Lambda\Lambda'} &\propto \int d\Omega_0 \sum_{\lambda_a \lambda_b \lambda_c} \langle \vec{p}_f \lambda_c \Lambda | T | \vec{p}_i \lambda_a \lambda_b\rangle \langle \vec{p}_i \lambda_a \lambda_b | T^\dagger | \vec{p}_f \lambda_c \Lambda'\rangle \\ &= \int d\Omega_0 \sum_{\lambda_a \lambda_b \lambda_c} \langle \vec{p}_f \lambda_c \Lambda | T | \vec{p}_i \lambda_a \lambda_b\rangle \langle \vec{p}_f \lambda_c \Lambda' | T | \vec{p}_i \lambda_a \lambda_b\rangle^*. \end{aligned} \tag{A.30}$$

We now use the parity conservation of T to find a symmetry relation in $\rho^J_{\Lambda\Lambda'}$. First we define the coordinate system so the interaction (A.26) takes place in the xz plane. Next, we define the y-parity operator by

$$\mathcal{P}_y = \mathcal{P} R(0, \pi, 0). \tag{A.31}$$

Clearly, $\mathcal{P}_y$ represents a reflection in the xz-plane:

$$\mathcal{P}_y : (x, y, z) \rightarrow (x, -y, z). \tag{A.32}$$

The corresponding unitary operator is denoted:

$$\Pi_y = U[\mathcal{P}_y] = \Pi e^{-i\pi J_y}. \tag{A.33}$$

Clearly Π_y commutes with both rotations about the y axis and boosts in the xz-plane:

$$[\Pi_y, U[R_y]] = 0 \qquad [\Pi_y, U[L(\vec{p})]] = 0, \quad \text{where } \vec{p} \cdot \hat{y} = 0 \tag{A.34}$$

And thus, when applied to the helicity states in (A.30), Π_y acts directly on the rest states $|s_i \lambda_i\rangle$:

$$\begin{aligned} \Pi_y |s_i \lambda_i\rangle &= \eta_i \sum_{\lambda} d^{s_i}_{\lambda \lambda_i}(\pi) \, |s_i \lambda\rangle \\ &= \eta_i (-1)^{s_i - \lambda_i} |s_i \, -\lambda_i\rangle \end{aligned} \tag{A.35}$$

where we have used the identity[Ch71, p. 77]:

$$d^j_{m'm}(\pi) = (-1)^{j-m} \delta_{m', -m}. \tag{A.36}$$

Now, using these results in Equation (A.30),

$$\begin{aligned} \rho^J_{-\Lambda \, -\Lambda'} &= \int d\Omega_0 \sum_{\lambda_a \lambda_b \lambda_c} \langle \vec{p}_f \lambda_c \; -\Lambda | T | \vec{p}_i \lambda_a \lambda_b \rangle \, \langle \vec{p}_f \lambda_c \; -\Lambda' | T | \vec{p}_i \lambda_a \lambda_b \rangle^* \\ &= (-1)^{s_c - \lambda_c} \eta_c (-1)^{J - \Lambda} \eta (-1)^{s_c - \lambda_c} \eta_c (-1)^{J - \Lambda'} \eta \times \end{aligned}$$

$$
\begin{aligned}
& \int d\Omega_0 \sum_{\lambda_a\lambda_b\lambda_c} \langle \vec{p}_f\ -\lambda_c\Lambda|\Pi_y T|\vec{p}_i\lambda_a\lambda_b\rangle\, \langle \vec{p}_f\ -\lambda_c\Lambda'|\Pi_y T|\vec{p}_i\lambda_a\lambda_b\rangle^* \\
= \ & (-1)^{\Lambda-\Lambda'} \int d\Omega_0 \sum_{\lambda_a\lambda_b\lambda_c} \langle \vec{p}_f\ -\lambda_c\Lambda|T\Pi_y|\vec{p}_i\lambda_a\lambda_b\rangle\, \langle \vec{p}_f\ -\lambda_c\Lambda'|T\Pi_y|\vec{p}_i\lambda_a\lambda_b\rangle^* \\
= \ & (-1)^{\Lambda-\Lambda'}(-1)^{s_a-\lambda_a}\eta_a(-1)^{s_b-\lambda_b}\eta_b(-1)^{s_a-\lambda_a}\eta_a(-1)^{s_b-\lambda_b}\eta_b \times \\
& \int d\Omega_0 \sum_{\lambda_a\lambda_b\lambda_c} \langle \vec{p}_f\ -\lambda_c\Lambda|T|\vec{p}_i\ -\lambda_a\ -\lambda_b\rangle\, \langle \vec{p}_f\ -\lambda_c\Lambda'|T|\vec{p}_i\ -\lambda_a\ -\lambda_b\rangle^* \\
= \ & (-1)^{\Lambda-\Lambda'}\rho_{\Lambda\Lambda'}
\end{aligned}
\tag{A.37}
$$

In the last line, we have used the fact that a sum over $-\lambda_i$ is the same as a sum over λ_i.

Now, we consider the case of a spin-1 particle, $J = 1$. The spin density matrix is thus a 3×3 matrix. The fact that it is Hermitian and normalized so that $\mathrm{Tr}(\rho) = 1$ means it will have eight independent real components. The symmetry relation (A.37), however, reduces this number to four:

$$
\rho = \begin{pmatrix} \rho_{11} & \rho_{10} & \rho_{1-1} \\ \rho_{10}^* & 1-2\rho_{11} & -\rho_{10}^* \\ \rho_{1-1} & -\rho_{10} & \rho_{11} \end{pmatrix}
\tag{A.38}
$$

with ρ_{1-1} and ρ_{11} real. Here we have used the convention that the first row or column of the matrix correspond to helecity state $\lambda = +1$ and the last to $\lambda = -1$.

A.4 Angular Distribution of Decay Products

Now we consider both the production and decay of a resonance with spin-J

$$
a + b \to c + J \qquad J \to 1 + 2
\tag{A.39}
$$

We shall also use the same notations as used in the previous section. If we let $\mathcal{M}_d$ denote the decay operator of the resonance J, then the overall amplitude for (A.39) can be written:

$$\mathcal{M}_{fi} \propto \sum_{\Lambda} \langle \vec{p}\,\lambda_1\lambda_2|\mathcal{M}_d|J\Lambda\rangle\langle \vec{p}_f\lambda_c\Lambda|T|\vec{p}_i\lambda_a\lambda_b\rangle \tag{A.40}$$

Let $\Omega = (\theta, \phi)$ denote the decay angle of particle 1 in the rest frame of the resonance J. Integrating over the other variables, the differential cross section becomes:[2]

$$\frac{d\sigma}{d\Omega} \propto \int d\Omega_0 \sum |\mathcal{M}_{fi}|^2. \tag{A.41}$$

Combining (A.41), (A.40), and substituting (A.30) into the result, we obtain:

$$\frac{d\sigma}{d\Omega} \propto \sum_{\Lambda\Lambda'\lambda_1\lambda_2} \langle \vec{p}\,\lambda_1\lambda_2|\mathcal{M}_d|J\Lambda\rangle\rho^J_{\Lambda\Lambda'}\langle J\Lambda'|\mathcal{M}_d^\dagger|\vec{p}\,\lambda_1\lambda_2\rangle. \tag{A.42}$$

For the decay amplitude of the resonance of the spin-J particle with $J_z = M$, we can write

$$\langle\phi\theta\lambda_1\lambda_2|\mathcal{M}_d|JM\rangle = \langle\phi\theta\lambda_1\lambda_2|JM\lambda_1\lambda_2\rangle\langle JM\lambda_1\lambda_2|\mathcal{M}_d|JM\rangle \tag{A.43}$$

since the decay operator necessarily conserves angular momentum. From (A.25) and the normalization (A.19), it follows that

$$\langle\theta\phi\lambda_1\lambda_2|JM\lambda_1\lambda_2\rangle = N_J D^{J*}_{M\lambda}(\phi, \theta, 0) \tag{A.44}$$

where $\lambda = \lambda_1 - \lambda_2$. Now, since $\mathcal{M}_d$ is a rotational invariant, the second bracket in (A.43) must depend only on the rotationally invariant quantities J, λ_1, and λ_2. So

[2]We assume here that the transition operator T is approximately constant over the width of the resonance, and so we don't explicitly show the integration over the Breit-Wigner distribution.

we can write,

$$\langle \phi\theta\lambda_1\lambda_2|\mathcal{M}_d|JM\rangle = N_J\, F^J_{\lambda_1\lambda_2} D^{J*}_{M\lambda}(\phi,\theta,0). \tag{A.45}$$

If we let

$$g^J_{\lambda_1\lambda_2} \propto |F^J_{\lambda_1\lambda_2}|^2 \tag{A.46}$$

then we can substitute this result into (A.42) to obtain:

$$\frac{d\sigma}{d\Omega} \propto N_J^2 \sum_{\Lambda\Lambda'\lambda_1\lambda_2} \rho^J_{\Lambda\Lambda'}\, D^{J*}_{\Lambda\lambda}(\phi,\theta,0)\, D^J_{\Lambda'\lambda}(\phi,\theta,0) g^J_{\lambda_1\lambda_2} \tag{A.47}$$

where $\lambda = \lambda_1 - \lambda_2$.

In our case, we have $J = 1$, and particles 1 and 2 each have spin-0. Therefore, the factor $g^J_{\lambda_1\lambda_2} = g^1_{00}$ is a constant, and $\lambda = \lambda_1 - \lambda_2 = 0$, and so we have

$$\frac{d\sigma}{d\Omega} \propto \sum_{\Lambda\Lambda'} \rho_{\Lambda\Lambda'} D^{1*}_{\Lambda 0}(\phi,\theta,0)\, D^1_{\Lambda' 0}(\phi,\theta,0) \tag{A.48}$$

From Equation (A.7), we have $D^1_{M0}(\phi,\theta,0) = e^{iM\phi} d^1_{M0}(\theta)$. From tables of d-functions[Ha02, p. 245], we find:

$$d^1_{0\,0}(\theta) = \cos\theta \tag{A.49}$$

and

$$d^1_{-1\,0}(\theta) = -d^1_{1\,0}(\theta) = \frac{\sin\theta}{\sqrt{2}} \tag{A.50}$$

Substituting the above, along with (A.38) into (A.48), we obtain

$$\frac{d\sigma}{d\Omega} \propto \begin{pmatrix} -e^{-i\phi}\sin\theta & \sqrt{2}\cos\theta & e^{i\phi}\sin\theta \end{pmatrix} \begin{pmatrix} \rho_{11} & \rho_{10} & \rho_{1-1} \\ \rho_{10}^* & 1-2\rho_{11} & -\rho_{10}^* \\ \rho_{1-1} & -\rho_{10} & \rho_{11} \end{pmatrix} \begin{pmatrix} -e^{i\phi}\sin\theta \\ \sqrt{2}\cos\theta \\ e^{-i\phi}\sin\theta \end{pmatrix} \tag{A.51}$$

Performing the matrix multiplication and simplifying, we obtain

$$\begin{aligned} \frac{d\sigma}{d\Omega} \propto I(\theta,\phi) &= \frac{1}{4\pi}[1+(1-3\cos^2\theta)(3\rho_{11}-1) \\ &\quad -(3\sqrt{2}\sin 2\theta\cos\phi)\mathrm{Re}\,\rho_{10} \\ &\quad -(3\sin^2\theta\cos 2\phi)\rho_{1-1}] \end{aligned} \tag{A.52}$$

where $I(\theta,\phi)$ has been normalized so that

$$\int I(\theta,\phi)\,d\Omega = 1. \tag{A.53}$$

Bibliography

[Ab87] S. Abachi *et al.*, Phys. Lett. B **199** (1987) 585.

[Ac97] K. Ackerstaff *et al.*, Z. Phys. C **74** (1997) 437.

[Ad99] M. Adinolfi *et al.*, Nucl. Phys. B **547** (1999) 3.

[Ag86] M. Aguilar-Benitez *et al.*, Phys. Lett. B **169** (1986) 106.

[Ai02] E. M. Aitala *et al.*, Phys. Lett. B **539** (2002) 218.

[Ai96] E. M. Aitala *et al.*, Phys. Lett. B **371** (1996) 157.

[Ai97] E. M. Aitala *et al.*, Phys. Lett. B **411** (1997) 230.

[Ai99] E. M. Aitala *et al.*, Phys. Lett. B **462** (1999) 225.

[Al77] G. Altarelli and G. Parisi, Nucl. Phys. B **126** (1977) 298.

[Al92] G. A. Alves *et al.*, Phys. Rev. Lett. **69** (1992) 3147.

[Al94a] G. A. Alves *et al.*, Phys. Rev. Lett. **72** (1994) 812.

[Al94b] G. A. Alves *et al.*, Phys. Rev. D **49** (1994) 4317.

[Am93] S. Amato *et al.*, Nucl. Instrum. Meth. **A324** (1993) 535.

[An83] B. Andersson, G. Gustafson, and B. Söderberg, Z. Phys. C **20** (1983) 317.

[An97] J. C. Anjos, G. Herrera, J. Magnin, and F. R. A. Simão, Phys. Rev. D **56** (1997) 394.

[Ap86] J. A. Appel *et al.*, Nucl. Instrum. Meth. **A 243** (1986) 361.

[Au74a] J. Aubert, *et al.*, Phys. Rev. Lett. **33** (1974) 1404.

[Au74b] J. Augustin, *et al.*, Phys. Rev. Lett. **33** (1974) 1406.

[Av03] C. Avila, J. Magnin, L. M. Mendoza-Navas. E-print: hep-ph/0307358 v2 (2003).

[Ba87] D. Bartlett *et al.*, Nucl. Instrum. Meth. **A 260** (1987) 55.

[Ba91] S. Barlag *et al.*, Z. Phys. C **49** (1991) 555.

[Ba88] S. Barlag *et al.*, Z. Phys. C **39** (1988) 451.

[Bh78] V. Bharadwaj *et al.*, Nucl. Instrum. Meth. **155** (1978) 411.

[Bh85] V. Bharadwaj *et al.*, Nucl. Instrum. Meth. **A 228** (1985) 283.

[Bj67] J. Bjorken, Phys. Rev. **163** (1967) 1767.

[Bl96] K. Blum, *Density Matrix Theory and Applications.* 2nd Ed. Plenum Press, 1996. (ISBN 0-306-45341-X)

[Br02] S. Bracker and S. Hansen, E-print: hep-ex/0210034.

[Br96] S. Bracker, K. Gounder, K. Hendrix, and D. Summers, IEEE Trans. Nucl. Sci. **43** (1996) 2457.

[Br98] G. Brandenburg *et al.*, Phys. Rev. D **58** (1998) 052003.

[Bo81] M. G. Bowler, Z. Phys. C **11** (1981) 169.

[Bu76] G. Bunce *et al.*, Phys. Rev. Lett. **36** (1976) 1113.

[Ch71] S. U. Chung, *Spin Formalisms*. CERN 71-8 (1971).

[Da77] K. Das and R. Hwa, Phys. Lett. B **68** (1977) 459; Erratum *ibid.* **73** (1978) 504.

[Fa94] A. F. Falk and M. E. Peskin, Phys. Rev. D **49** (1994) 3320.

[Fe86] R. Fernow, *Introduction to Experimental Particle Physics*. Cambridge University Press, 1986. (ISBN 0-521-37940-7)

[Fr94] S. Frixione, M. L. Mangano, P. Nason, and G. Ridolfi, Nucl. Phys. B **431** (1994) 453.

[Ge64] M. Gell-Mann, Phys. Lett. **8** (1964) 214.

[Gl70] S. Glashow, J. Iliopoulos, and L. Maiani, Phys. Rev. D **2** (1970) 1285.

[Gr64] O. Greenberg, Phys. Rev. Lett. **13** (1964) 598.

[Ha02] K. Hagiwara *et al.*, Phys. Rev. D **66** (2002) 010001.

[Ha87a] R. Hance *et al.*, IEEE Trans. Nucl. Sci. **34** (1987) 878.

[Ha87b] S. Hansen *et al.*, IEEE Trans. Nucl. Sci. **34** (1987) 1003.

[Ha90] P. N. Harriman, A. D. Martin, R. G. Roberts, and W. J. Stirling, Phys. Rev. D **42** (1990) 798.

[Jo78] L. M. Jones and H. W. Wyld, Phys. Rev. D **17** (1978) 1782.

[Ku91] Y. Kubota *et al.*, Phys. Rev. D **44** (1991) 593.

[Lo92] L. Lönnblad, C. Peterson, and T. Rögnvaldsson, Comput. Phys. Commun. **70** (1992) 167.

[Mc83] J. G. McEwen *et al.*, Phys. Lett. B **121** (1983) 198.

[Na89] P. Nason, S. Dawson, R. Ellis, Nucl. Phys. B **327** (1989) 49.

[Pe83] C. Peterson, D. Schlatter, I. Schmitt, and P. M. Zerwas, Phys. Rev. D **27** (1983) 105.

[Pe87] D. Perkins, *Introduction to High Energy Physics.* 3rd Ed. Addison-Wesley, 1987. (ISBN 0-201-12105-0)

[Pl95] H. Plothow-Besch, Int. J. Mod. Phys. **A10** (1995) 2901.

[Ra88] J. R. Raab *et al.*, Phys. Rev. D **37** (1988) 2391.

[Ro57] M. E. Rose, *Elementary theory of angular momentum.* John Wiley & Sons, Inc., New York, 1957. (Reprinted by Dover Publications Inc., 1995. ISBN 0-486-68480-6)

[Sj94] T. Sjöstrand, Comput. Phys. Commun. **82** (1994) 74. (PYTHIA version 5.7 and JETSET version 7.4 were used.)

[Sl85] K. Sliwa *et al.*, Phys. Rev. D **32** (1985) 1053.

[Su85] D. J. Summers, Nucl. Instrum. Meth. **A 228** (1985) 290.

[Su86] M. Suzuki, Phys. Rev. D **33** (1986) 676.

[Su92a] D. J. Summers *et al.*, XXVII Rencontre de Moriond, Electroweak, 15–22 March, 1992, 417. (E-print: hep-ex/0009015.)

[Su92b] P. J. Sutton, A. D. Martin, R. G. Roberts, and W. J. Stirling, Phys. Rev. D **45** (1992) 2349.

[Tr03] S. M. Troshin and N. E. Tyurin, AIP Conf. Proc. **675** (2003) 579. (E-print: hep-ph/0201267)

[Vo95] R. Vogt and S. J. Brodsky, Nucl. Phys. B **438** (1995) 261.

[Wi94] J. Wiener, Ph.D. Dissertation. Princeton University (1994).

[Wi96] N. Witchey, Ph.D. Dissertation. Ohio State University (1996).