THE SILICON DRIFT DETECTOR AND ITS USE IN A VERTEX TRACKER

APPROVED BY
DISSERTATION COMMITTEE:

Supervisor: ______________________________

____

____

____

____

For Martha Lane Orr, who waited

and

For Annemarie Roth, who understood

THE SILICON DRIFT DETECTOR AND ITS USE IN A VERTEX TRACKER

by

DAVID M. READ, B.S.

DISSERTATION

Presented to the Faculty of the Graduate School of
The University of Texas at Austin
in Partial Fulfillment
of the Requirements
for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN
May 1996

Acknowledgements

The assistance and guidance of Dr. G. W. Hoffmann is gratefully acknowledged. The generous help of Drs. K. Wilson and C. Pruneau in resolving software difficulties is also acknowledged. Finally, the patient explanations and helpful suggestions of Dr. R. L. Ray are appreciated.

THE SILICON DRIFT DETECTOR
AND ITS USE IN A VERTEX TRACKER

Publication No. ____________________

David M. Read, Ph.D.
The University of Texas at Austin, 1996

Supervisor: G.W. Hoffmann

Silicon Drift Detectors (SDDs) represent a step forward in the detection of fast charged particles in heavy-ion experiments. They exhibit excellent spatial and charge resolution, and tolerate high occupancies and event rates. The design of silicon drift detectors is discussed, prototype detectors are tested, computer code to simulate the performance and analysis of data is developed, and a working detector composed of SDDs is discussed and simulated.

Table of Contents

List of Tables

List of Figures

Chapter 1

STAR physics with an SVT

The material in this chapter has appeared previously [1].

1.1 STAR goals

The primary physics goal of the RHIC project is to collide heavy nuclei with sufficient energy to produce a transition from hot, dense, hadronic matter to a color de-confined, chirally symmetric plasma of quarks and gluons (QGP). STAR will focus mainly on the investigation of soft processes resulting in hadron production at central rapidities and transverse momenta below 2 GeV/c but, in addition, will study hard QCD processes such as the production of jets, mini-jets and hard photons. STAR will also investigate nucleon structure functions and effects of quark and gluon shadowing through proton-proton and proton-nucleus interactions. The intent of the experimental program is to measure many observables, some of which are obtained on an event-by-event basis, and study their dependence on the size of the colliding system, the collision impact parameter and total energy.

The time evolution of nucleus-nucleus collisions is thought to involve an initial parton scattering, pre-equilibrium stage, an intermediate thermalization stage in which the QGP might be formed, and a final expansion and hadronization stage. Observables can be associated with each stage which convey different information about the collision dynamics. For example, hard parton scatterings in the initial pre-equilibrium stage result in high $p_\perp$ particles, jets, high energy γs, and charmed mesons. Conditions during the intermediate thermalization stage (*i.e.,* hot hadronic gas or QGP) affect the production of mini-jets, strangeness, anti-baryons, isospin anomalies and energy/entropy fluctuations. Multi-strange baryon and ϕ meson production are strongly dependent on strangeness density during this stage and are therefore sensitive indicators of possible chiral symmetry restoration in the QGP. Global properties such as size and temperature during the final hadronization stage can be determined by the $p_\perp$ spectra and by $\pi\pi$ and KK HBT interferometry.

The baseline configuration of the STAR detector consists of a large room temperature solenoidal magnet, a time projection chamber (TPC) for charged particle tracking, and trigger counters (vertex position detectors, central trigger barrel). Charged particle momenta will be measured with the TPC at mid-rapidity ($|\eta| \leq 2$) with full azimuthal coverage and a transverse momentum, $p_\perp$, threshold of about 150 MeV/c. Particle identification through measurement of ionization density (dE/dx) will be accomplished for particles emitted with $|\eta| \leq 2$ and $150 \leq p_\perp \leq 600$MeV/c. With the TPC alone, STAR is limited to the study of hadron production with $p_\perp$ above 150 MeV/c and has a modest efficiency and accuracy for the study of short lived particles decaying before they reach the inner layer of the TPC. Multi-strange baryons and anti-baryons cannot be studied with the STAR baseline detector.

The purpose of the STAR Silicon Vertex Tracker (SVT) upgrade is to provide unique capabilities for the measurement of several of the preceding list of observables as well as to enhance the overall performance of STAR. Unique capabilities of the SVT include the detection of multi-strange baryons (Ξ^- and Ω^-) and measurement of low $p_\perp$ ($40 \leq p_\perp \leq 200$ MeV/c) spectra. Measurements of rare particles such as the D mesons and strange composite objects, which are impossible to observe with the TPC alone, might become accessible with the SVT. The SVT will provide substantial improvement in the momentum resolution for high $p_\perp$ particles, in the reconstruction efficiency for short lived neutral particles such as K_s^0, ϕ, Λ, and $\overline{\Lambda}$, and in two-track resolution for HBT studies owing to the excellent position resolution achievable with Si drift chamber technology. The SVT will also enhance track reconstruction and particle identification for soft charged particles such as $\pi^\pm$, $K^\pm$, p, $\overline{p}$, d and $\overline{d}$, will enable better determination of the primary vertex, and will help distinguish particles resulting from secondary decays.

It is unlikely that conclusive evidence of QGP formation at RHIC can be based on a single observable or signature; instead we will require the simultaneous observation of many signatures. It is expected that only a consistent analysis of many signals coming from all stages of the collision may lead to an unambiguous identification of the phase transition. STAR complemented with the SVT represents a powerful instrument for studying this possible transition to a new state of matter.

A detailed description of the integrated STAR system (baseline detector plus upgrades) can be found in the original STAR conceptual design report (CDR) [2].

In this document, we outline a proposal for the construction of an SVT for STAR which uniquely addresses low-$p_\perp$ and flavor physics otherwise unsatisfactorily covered by other RHIC detectors.

1.2 Physics Capabilities of the Silicon Vertex Tracker (SVT)

The SVT detector will enhance STAR capabilities for particle identification and momentum resolution and will extend the detector acceptance for low $p_\perp$ phenomena. It will offer unique capabilities for the study of strange and multi-strange particle production and in addition will allow HBT interferometry studies using neutral particles. The SVT will also expand the STAR third level trigger capabilities. These topics are discussed in more detail in the following.

1.2.1 Flavor physics

The different stages of nucleus-nucleus collisions (parton scattering, thermalization, expansion and hadronization) will be studied by measuring the yield and spectral distribution of produced hadrons. The formation of a QGP might be revealed through an enhancement of the relative abundance of strange to non-strange quarks [3]. Considerable reduction of the ratio of strange to anti-strange baryon production is also expected with the formation of a QGP [4]. Measurements of multiply strange baryon ratios (*e.g.* Ξ^-/Λ) are very sensitive to enhanced strange quark density [5], [6] which might result from chiral symmetry restoration in the QGP. The addition of the SVT with its high position resolution close to the production vertex allows the elimination of most of the combinatoric background and drastically improves the detection efficiency for strange particles. Figure 1.1 shows the tracking efficiency for secondary tracks in the SVT alone as a function of momentum.

To illustrate the positive impact of the excellent SVT position resolution we show in fig. 1.2 an Armenteros scatter plot for K_s^0 and Λ reconstructed from combined SVT and TPC tracks. Although no particle identification information is used in this plot, a clean separation of K_s^0 and Λ is obtained.

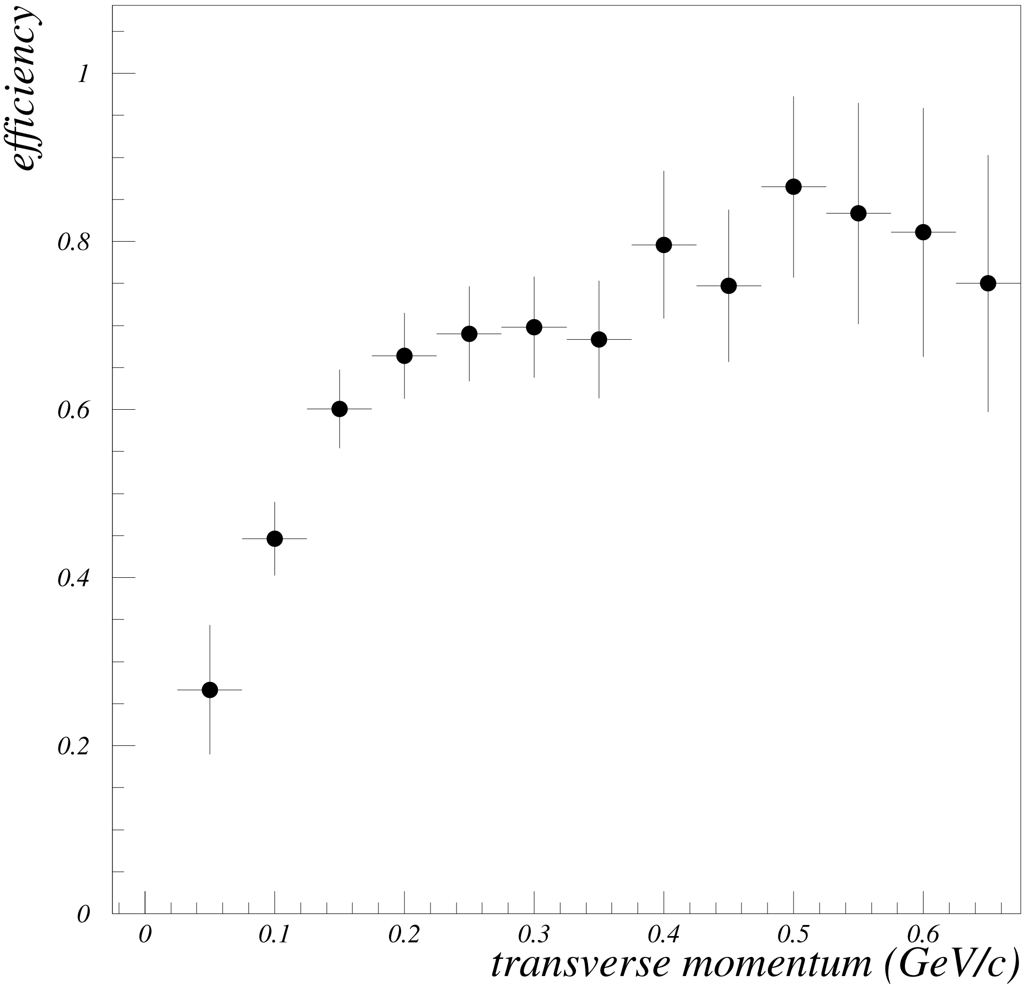

Figure 1.1: Secondary track reconstruction efficiency for the SVT alone

The capability of the combined SVT and TPC detector for reconstructing strange and multistrange particles is indicated by the invariant mass spec-

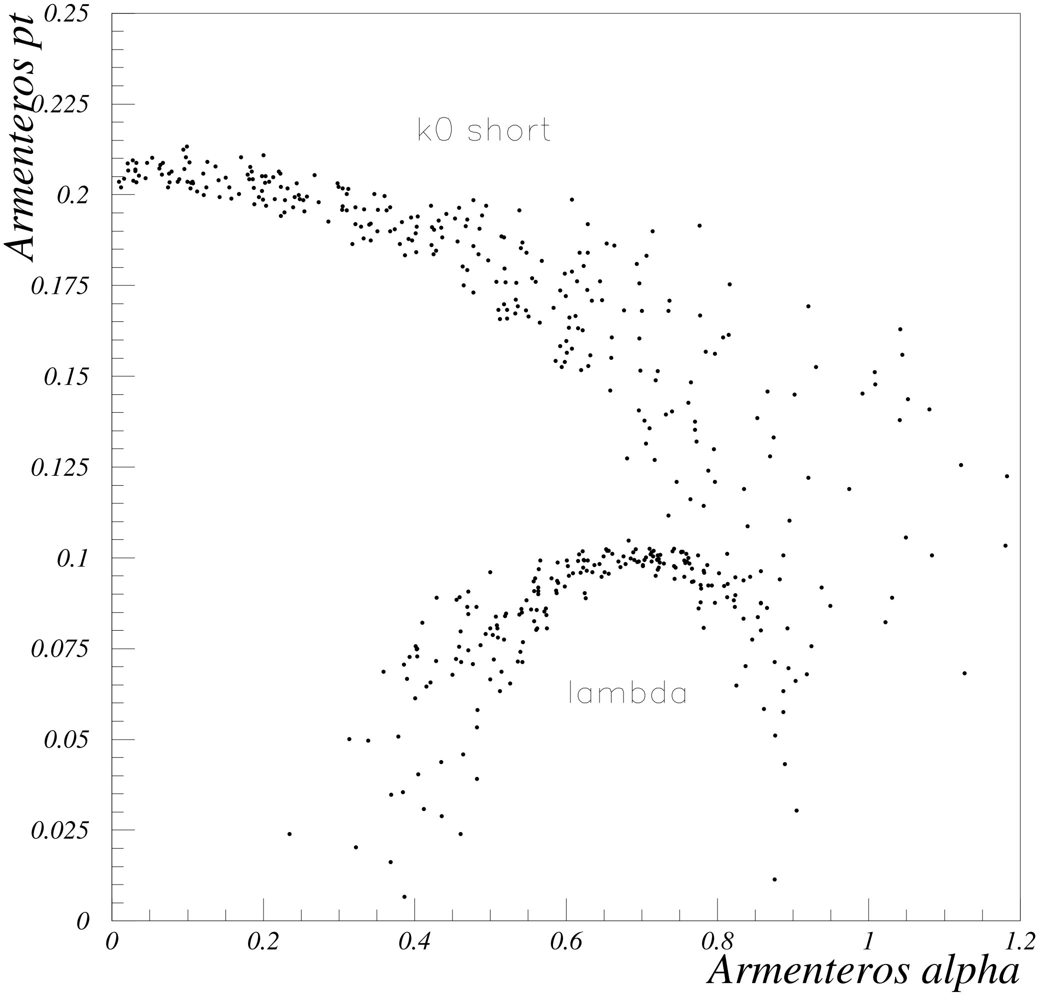

Figure 1.2: Armenteros plot showing K_s^0 and Λ separation based on SVT information

tra in figures 1.3 and 1.4 (see also [7]). To obtain the Λ spectrum a complete tracking algorithm in the SVT was employed, including momentum information from the TPC, matching efficiency, and detector resolution factors derived from simulations. The decay particles are identifiable via energy loss in the combined tracking system. The Ξ spectrum assumes ideal tracking. Based on the simulation for the Λ, we expect a modest reduction of the signal to noise ratio for the Ξ in the case of realistic tracking.

The ability of the SVT+TPC system to reconstruct multi-strange baryons coupled with the ability of STAR to study global event observables makes the proposed detector system unique among the RHIC experiments.

Modification of the mass and decay widths of certain very short lived particles has been predicted to occur in extremely dense hadronic matter and might also signal the formation of a QGP and chiral symmetry restoration [8]. Because of its narrow width and simple decay schemes the ϕ meson has received particular attention [9]. The excellent momentum resolution and wide $p_\perp$ coverage of combined SVT and TPC detectors will allow the study of possible medium modifications of the ϕ meson mass and width through the $K^+ + K^-$ decay channel.

Charm and in particular open charm production is currently believed to be an excellent indicator of the initial parton scattering stage of heavy ion collisions [10]. Because such particles must propagate through the medium they might also be sensitive to the production of the QGP. As the D-meson production cross section (σ_c/σ_s) and lifetime ($c\tau \approx 200\ \mu$m) are small and considering that most D decay channels are complex, the detection capability of open charm with STAR SVT and TPC is unclear at present. Strong enhancements in the

D meson production, predicted recently [11], will produce a cleaner signal than that expected from previous estimates of the production cross section. It is apparent however that such a measurement cannot be achieved without the excellent position resolution afforded by the SVT [12].

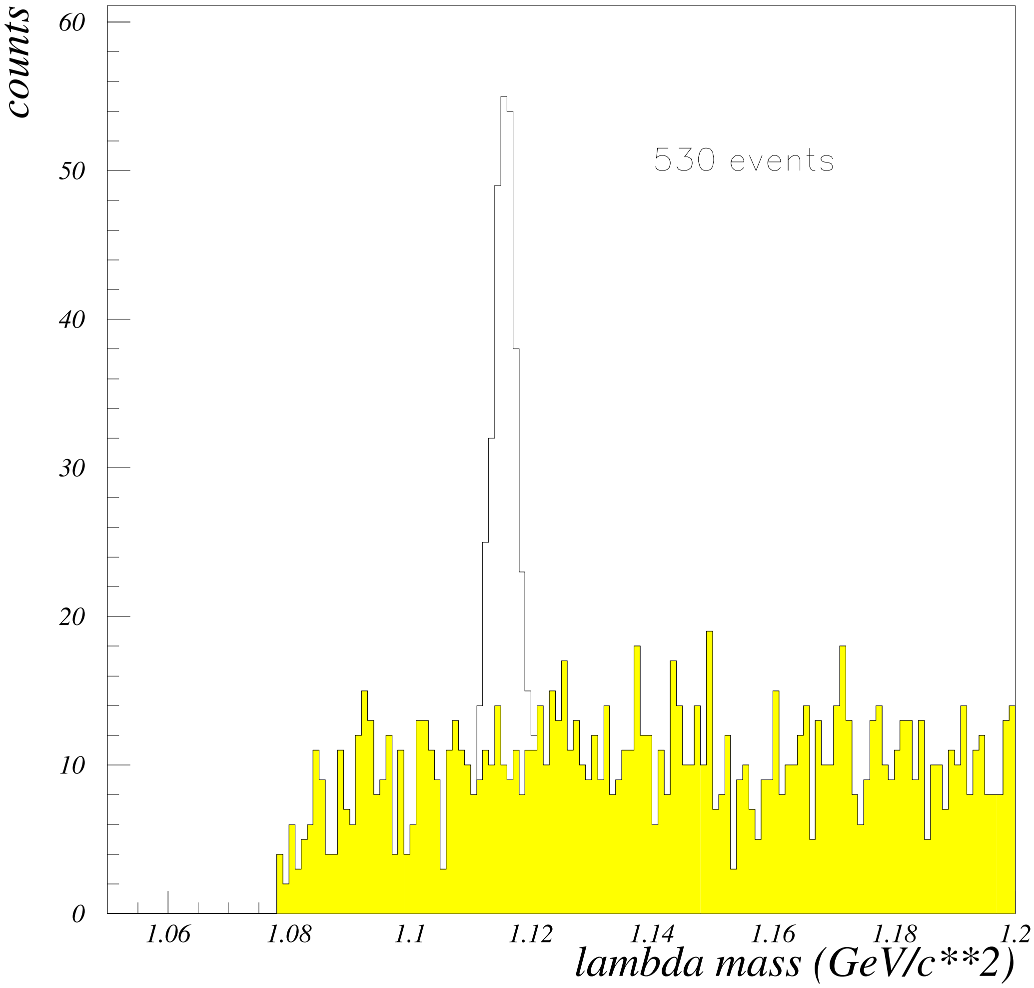

Figure 1.3: Invariant mass spectra of reconstructed Λ particles based on SVT and TPC tracking information.

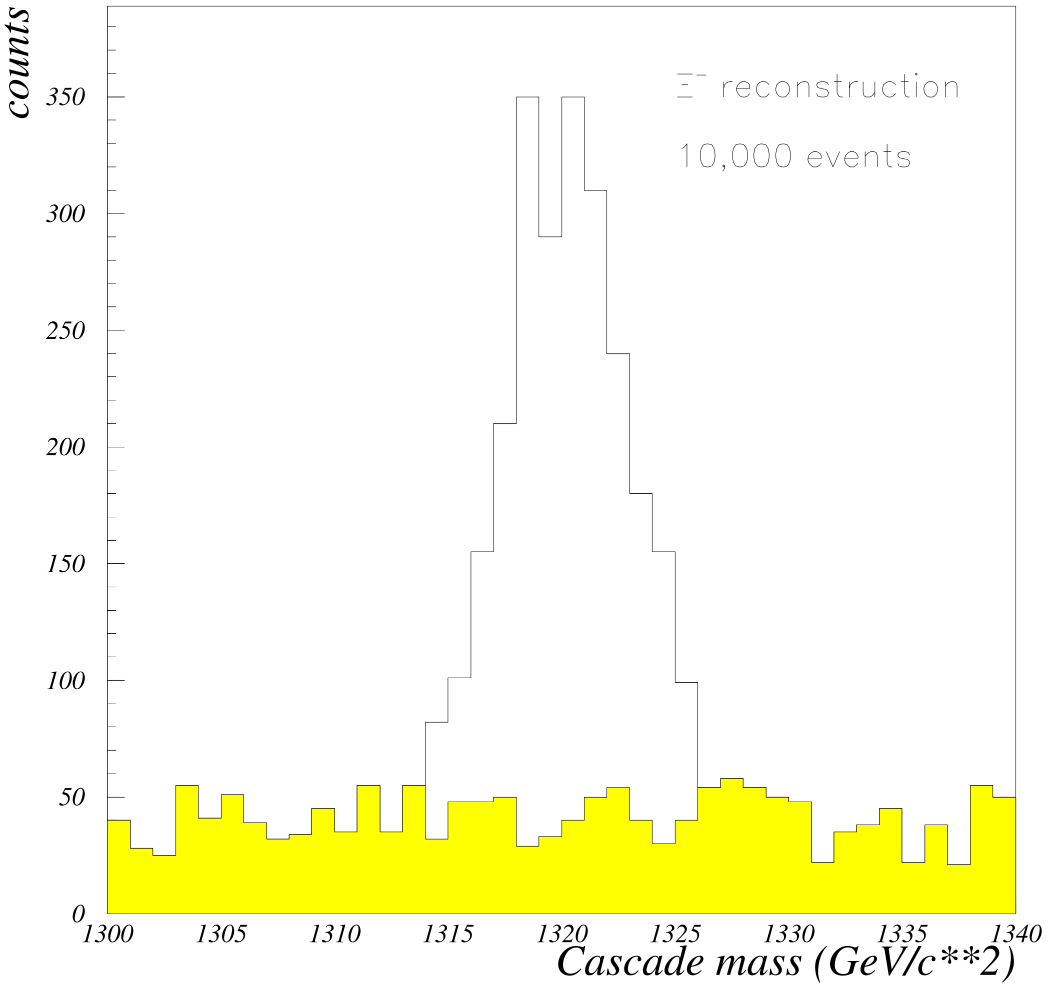

Figure 1.4: Invariant mass spectra of reconstructed Ξ particles based on SVT and TPC tracking information.

1.2.2 Low $p_\perp$ physics

Enhancements in the low momentum part of invariant cross section spectra were discovered in heavy ion reactions at CERN and AGS. Although several explanations were suggested by theorists, the phenomena related to pion $p_\perp$ spectra under 200 MeV/c is not fully understood [13], [14], [15], [16]. First results from the AGS for low $p_\perp$ kaons also reveal a very steep component at very low $p_\perp$ (below 50 MeV/c) which might be linked to a QGP phenomenon [17]. If these enhancements persist at the higher incident RHIC energies and in the baryon free rapidity region studied with STAR, explanations based on resonance models [18] could be excluded in favor of more exotic models based on some form of phase transition [16]. The high position resolution and low $p_\perp$ acceptance of the SVT for pions, kaons and protons is mandatory for such studies at RHIC. Several collective effects have been hypothesized which predict changes (enhancements) in the low $p_\perp$ particle spectra as well as correlations among the low $p_\perp$ particles. 'Van Hove bubbles' [18], localized plasma formation, might lead to spatially correlated hadronization of low $p_\perp$ particles and can be verified by measuring fluctuations of the yield in rapidity. Chiral condensate models [16] predict enhancements in the low $p_\perp$ particle spectra; disoriented chiral condensate models [19] in addition predict anomalous ratios of charged to neutral pions on an event-by-event basis. Hints of this behavior may have been observed in cosmic ray collisions [20]. Information on the yield of low $p_\perp$ charged pions on an event-by-event basis, as well as their correlation behavior as a function of centrality may provide information about the hadronization of a chiral condensate [21]. Bose-Einstein condensation [22] of pions is predicted to occur if a critical pion density can be reached at RHIC energies. If so, the majority of the pions will be emitted with very low momenta. In the case of coherent emission the mean momentum depends only on the source radius. SVT capabilities for detecting Bose-Einstein condensation effects in the RHIC environment are detailed in [23]. Figure 1.5 shows the magnitude of the effect.

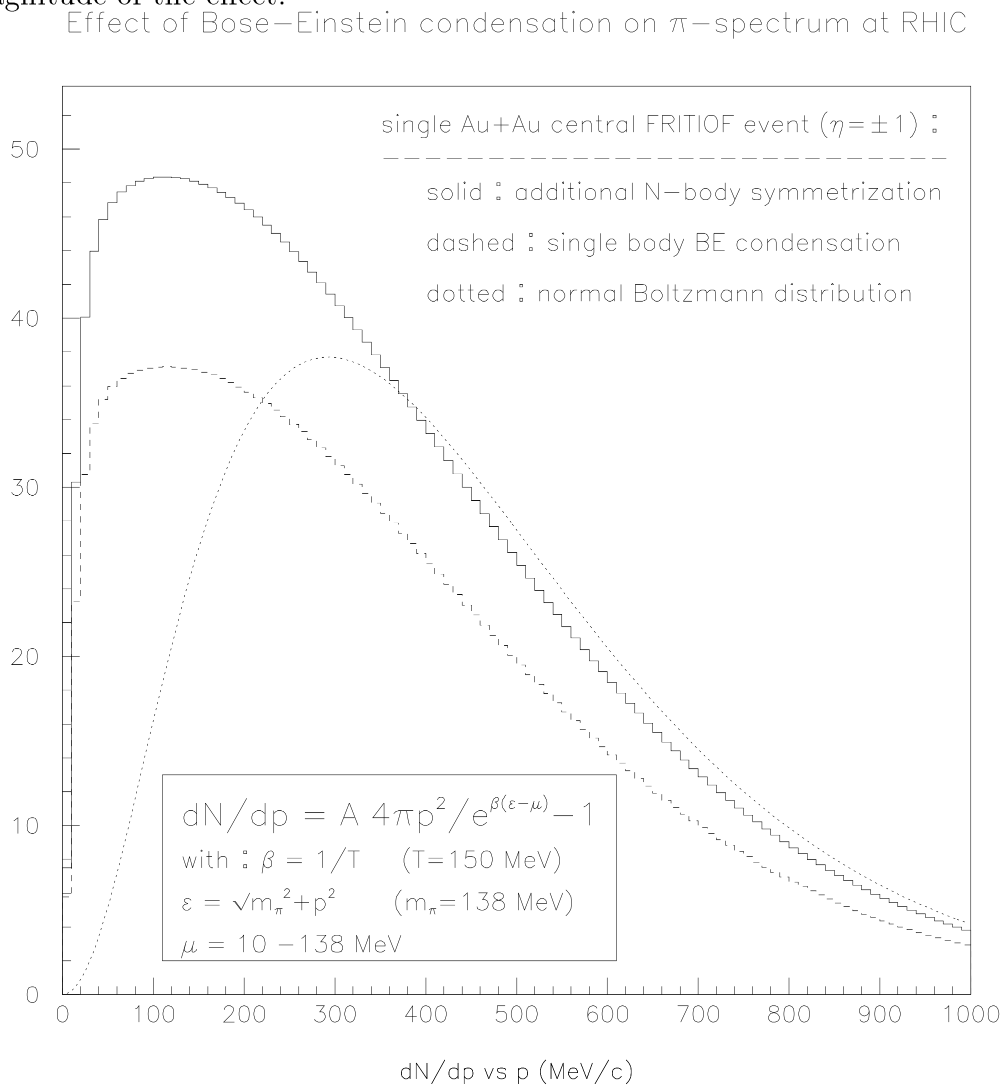

Figure 1.5: The effect of Bose Einstein condensation and n-body symmetrization on the π spectrum.

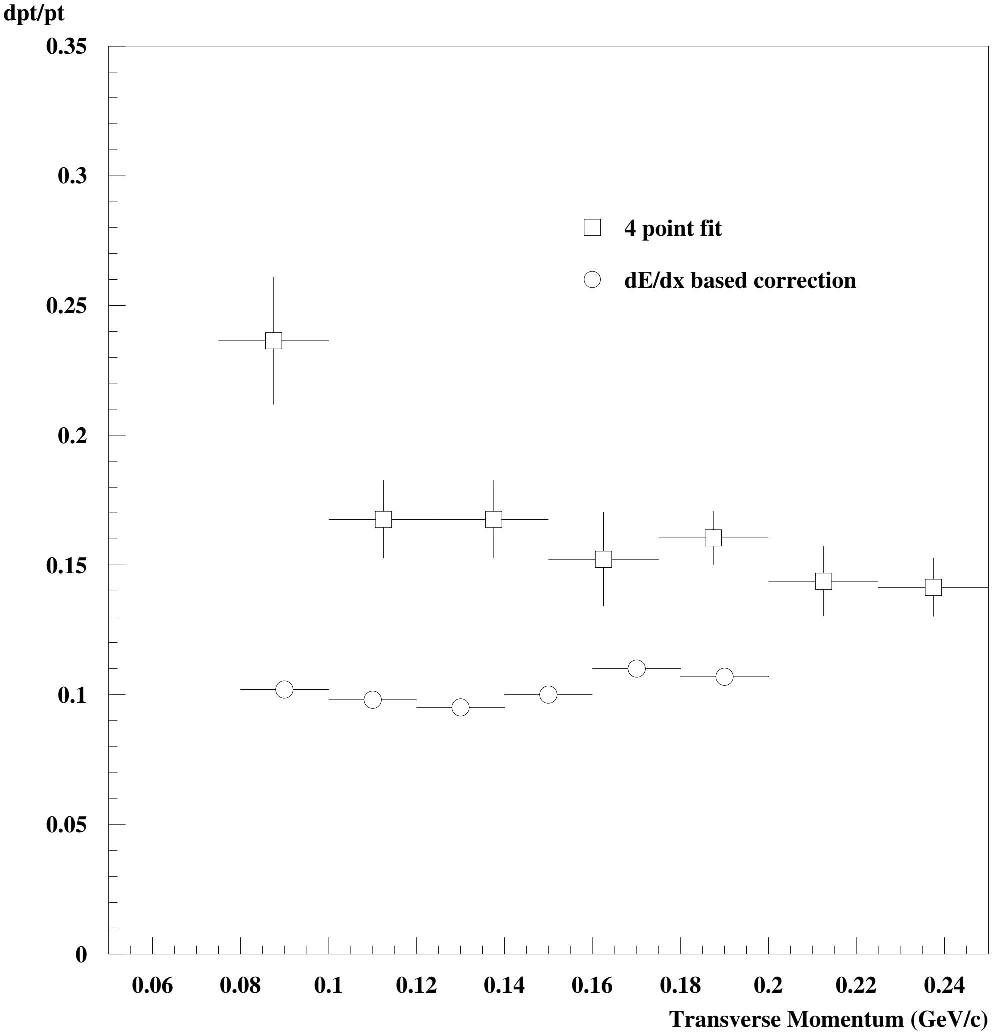

Figure 1.6: Momentum resolution for π reconstructed in the SVT alone

Figure 1.6 shows the pion momentum resolution for reconstructed tracks between 60 MeV/c and 250 MeV/c. The open dots represent the improvement achievable by including de/dx information from the SVT in the calculation. Although moderate, the resolution in fig. 3b is sufficient to permit momentum reconstruction, which preserves the shape of the spectrum at low $\mathrm{p}_\perp$. The presence of a low $\mathrm{p}_\perp$ enhancement will not only be determined for a large ensemble of events but, given the charged pion multiplicities expected at RHIC, it is statistically feasible to measure a shift in the mean of the momentum spectrum, event-by-event, which could be used as an exotic third level trigger (see section 1.2.4).

1.2.3 HBT Interferometry with the SVT

Measurements of two particle Bose-Einstein correlations can be interpreted in terms of the space-time structure of the emitting source (HBT interferometry) and may provide a signature of QGP formation [24], [25]. Theoretical studies show that kaon interferometry offers significant advantages over pions with respect to final state interactions and resonance formation, which obscure the interpretation of the data [26], [27]. Also, studies of interferometry based on hydrodynamic models [28] reveal that in order to best determine the structure of the emitting system the two-particle spectrum should include a wide range of the total two particle momenta including very low momenta. We discuss two applications of HBT interferometry for STAR in which the SVT makes a vital contribution: (1) by extending the acceptance to low $p_\perp$ and (2) by enabling neutral kaon (K_s^0) interferometry.

Increasing the $p_\perp$ acceptance improves the two pion statistics and might enable STAR to perform $\pi - \pi$ HBT on an event-by-event basis. Also, the study of two-particle correlations over an extended range of two-particle mean $p_\perp$ will provide better constraints on collision dynamics.

Neutral kaon interferometry via the SVT is unique to STAR. Central Au+Au collisions produce on the order of 200 short lived neutral kaons per event and of these approximately 10% are expected to be reconstructed using SVT tracking of the decay pions. The K_s^0 correlation function offers an obvious advantage due to the absence of Coulomb effects. A comparison between the results from charged and neutral kaon HBT will quantitatively determine the Coulomb effect. K_s^0 HBT is also not affected by two-track resolution, which is a limiting factor in charged particle correlations. K_s^0 are identified by their decay into charged pions. These pions are emitted back-to-back along a line oriented randomly in the K_s^0 rest frame. Even for two K_s^0 at exactly zero relative momentum the decay pions are well separated in configuration and momentum space and can be measured with the same efficiency as if the relative momentum of the kaons was large. Moreover, because the K_s^0 is not a strangeness eigenstate, the K_s^0-K_s^0 correlation includes a unique interference term that provides additional space-time information [29] as well as insight into possible strangeness distillation effects [30] in cases where the source baryon density is appreciable.

We simulated the K_s^0 HBT performance of STAR with the SVT for central Au+Au HIJING events including an imposed Bose-Einstein correlation enhancement as expected from a 10 fm radius spherical source. Considering that the kaons can be expected to freeze out earlier than the pions, this radius is plausible. The event reconstruction was simplified by assuming perfect matching between SVT and TPC tracks. Tracking inefficiencies and detector resolution were taken into account however. Details of the procedure can be found in [31]. Figures 1.7 and 1.8 show the correlation functions C(Q_{out}) and C(Q_{side}) [24], [25]. Q_{out}, which refers to the relative transverse momentum along the direction of the average pair momentum is predicted to offer the best discrimination between plasma and hadron gas scenarios [25], [32]. Our simulation suggests that with the SVT a measurement of Q_{out} from K_s^0 HBT is achievable. The SVT is essential for K_s^0 interferometry, because of its tracking capability which enables efficient secondary vertex reconstruction.

1.2.4 Other SVT capabilities which enhance the physics performance of the STAR baseline detector

The SVT will further enhance the capabilities of the STAR baseline detector in areas where the TPC performance is improved by combining SVT and TPC information. For the simulations presented in this section, full tracking within the SVT and matching with TPC tracks were performed as in the analysis of real events, unless otherwise mentioned.

Primary track reconstruction: The SVT not only extends the range of primary particle reconstruction in STAR to lower momenta, but also enhances the tracking efficiency for momenta below 400 MeV/c compared to the TPC alone. Figure 1.9 shows the primary track reconstruction efficiency for SVT and TPC as a function of the momentum. The plot corresponds to the present level of tracking capabilities. Improvements in both tracking software packages are expected in the future. Figure ref:track-eff shows the number of ghost tracks in the SVT tracking as a function of momentum. This number is greatly reduced by matching tracks between SVT and TPC.

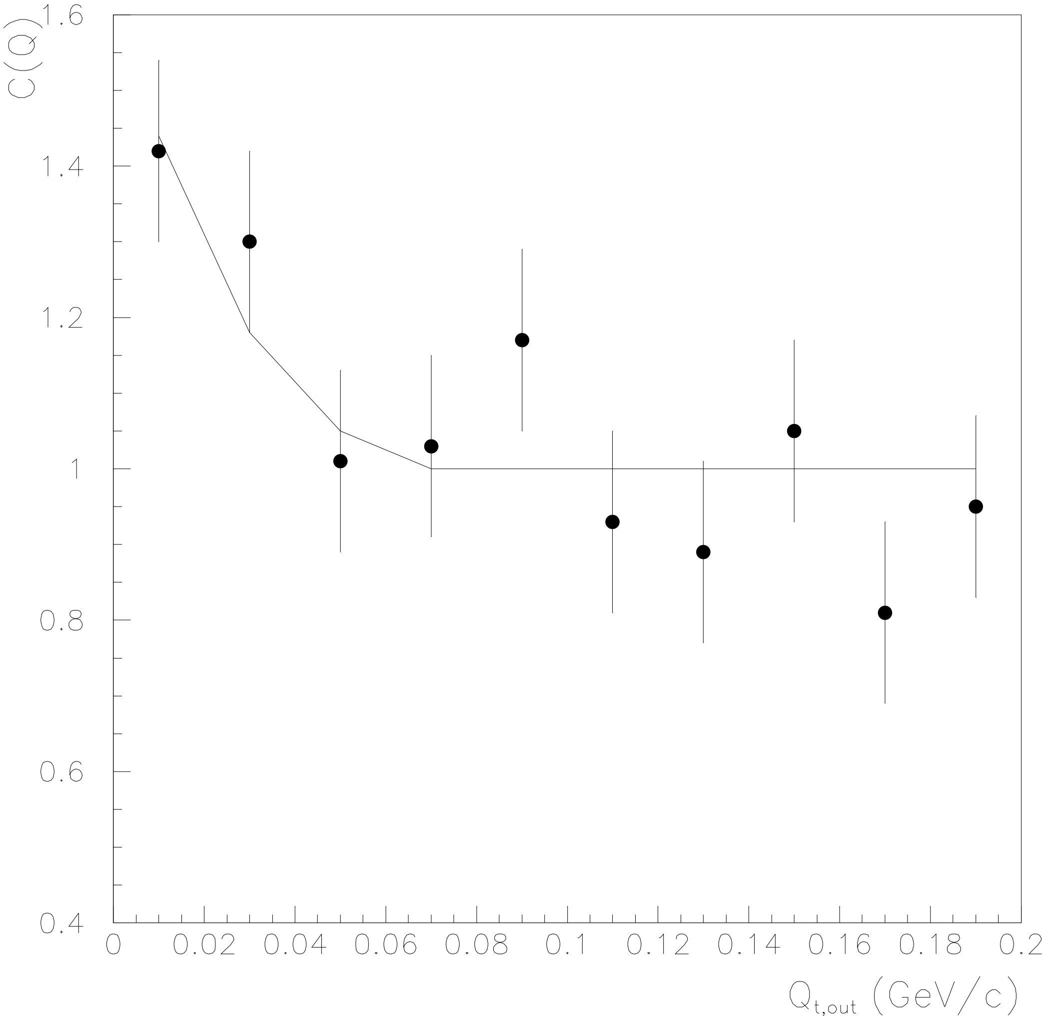

Figure 1.7: Measurements of correlation function C(Q$_{out}$) from K$_s^0$ pair reconstruction with SVT and TPC

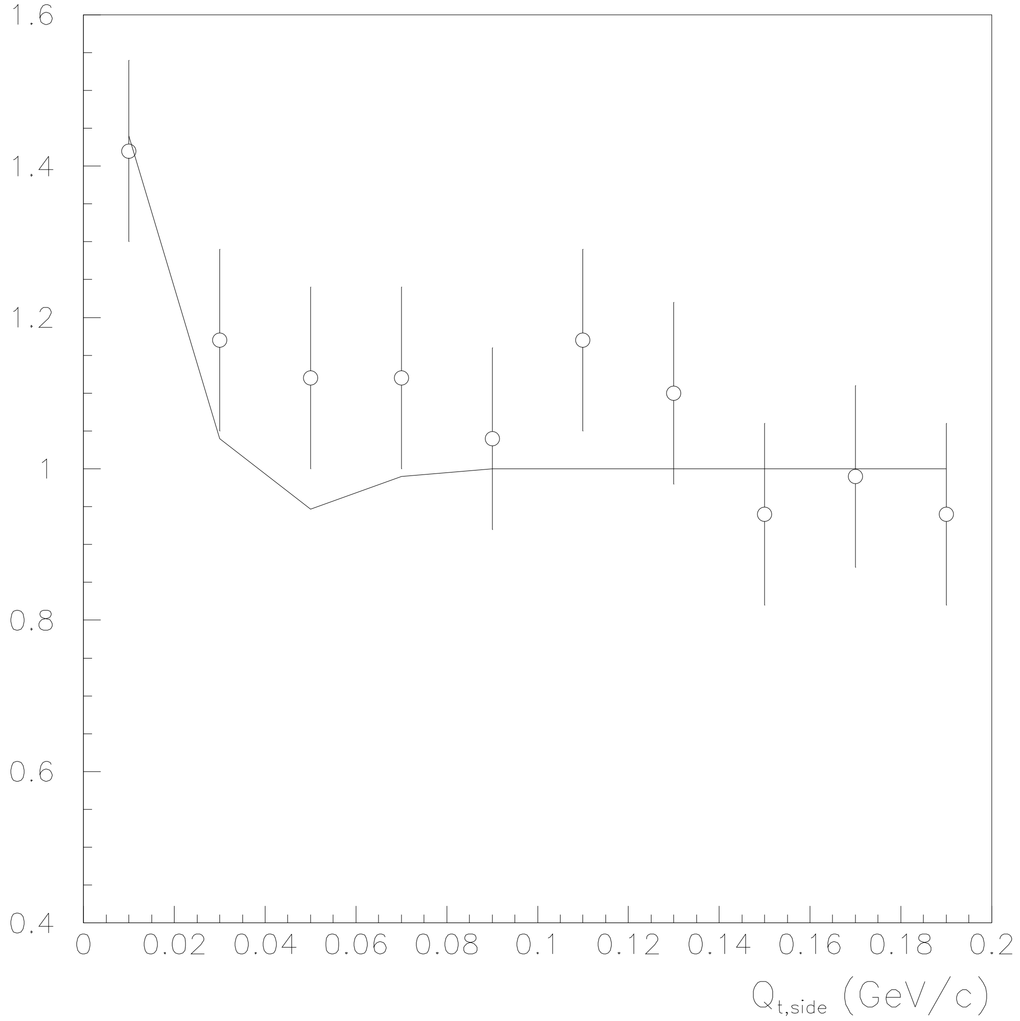

Figure 1.8: Measurements of correlation function C(Q_{side}) from K_s^0 pair reconstruction with SVT and TPC

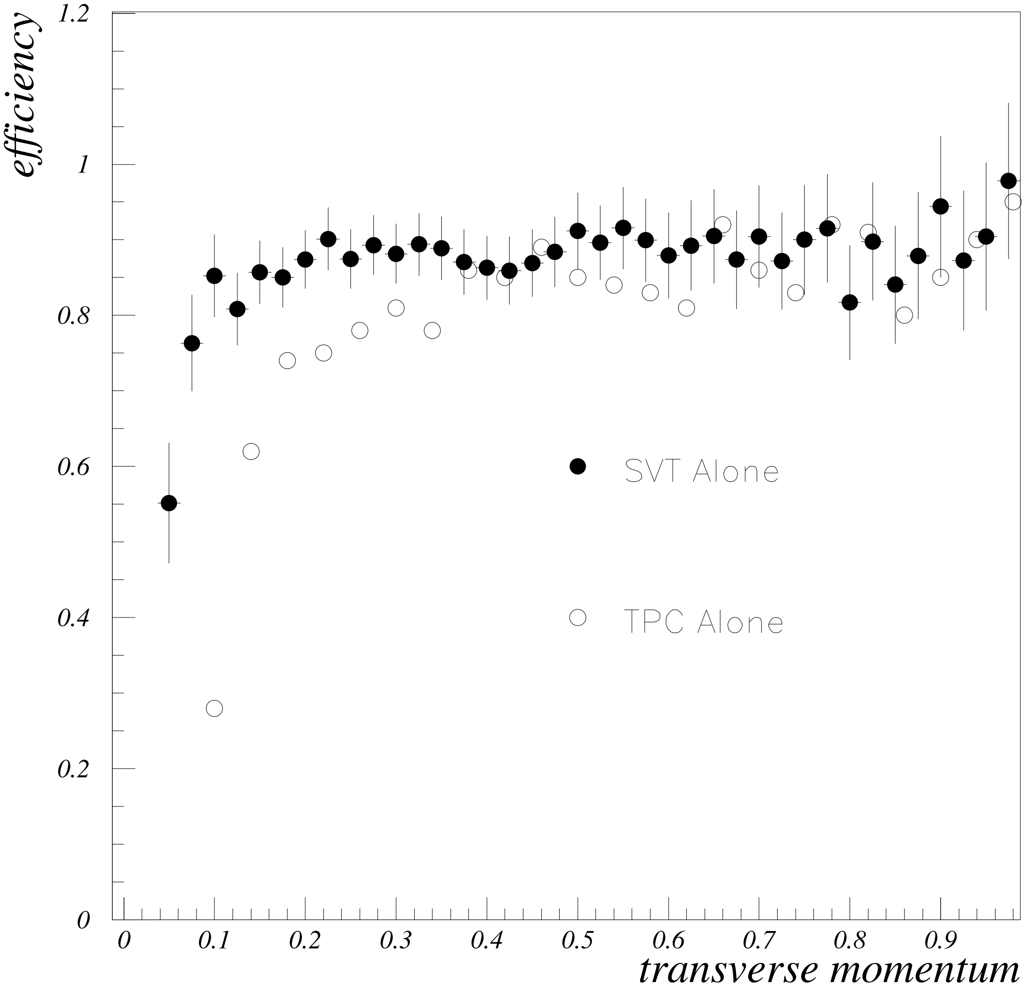

Figure 1.9: Primary track reconstruction efficiency for SVT and TPC as a function of momentum

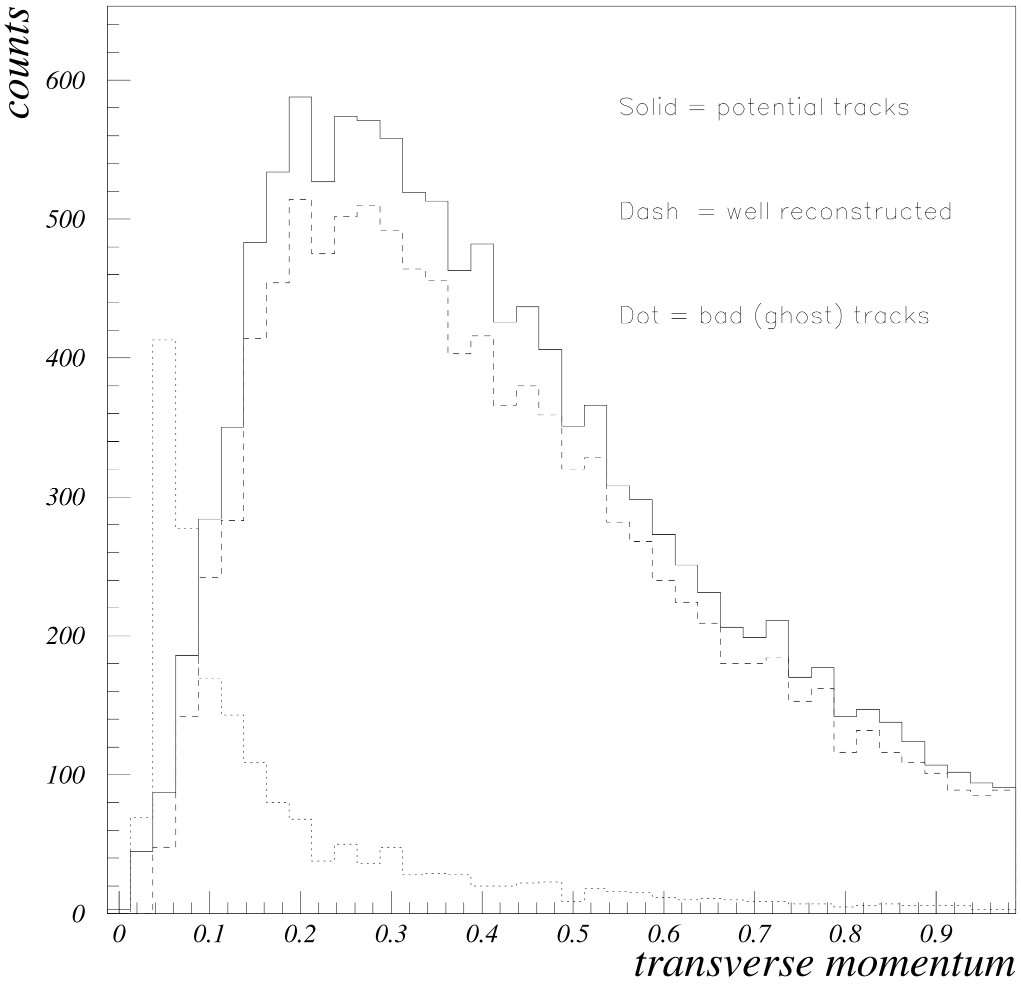

Figure 1.10: Classification of tracks in the SVT as a function of momentum based on the tracking efficiency

Particle Identification: The good de/dx resolution of the SVT will complement significantly the particle identification capabilities of the TPC. Figure 1.11 shows the average energy loss as a function of $p_{\perp}$ for the combined SVT+TPC system. Figure 1.12 shows a slice in momentum space. The addition of the SVT extends the upper limit for $K - \pi$ separation to about 600 MeV/c and for $p - \pi$ to about 800 MeV/c. The improvement in particle identification will aid in HBT analysis, strangeness production studies, and in overall event characterization.

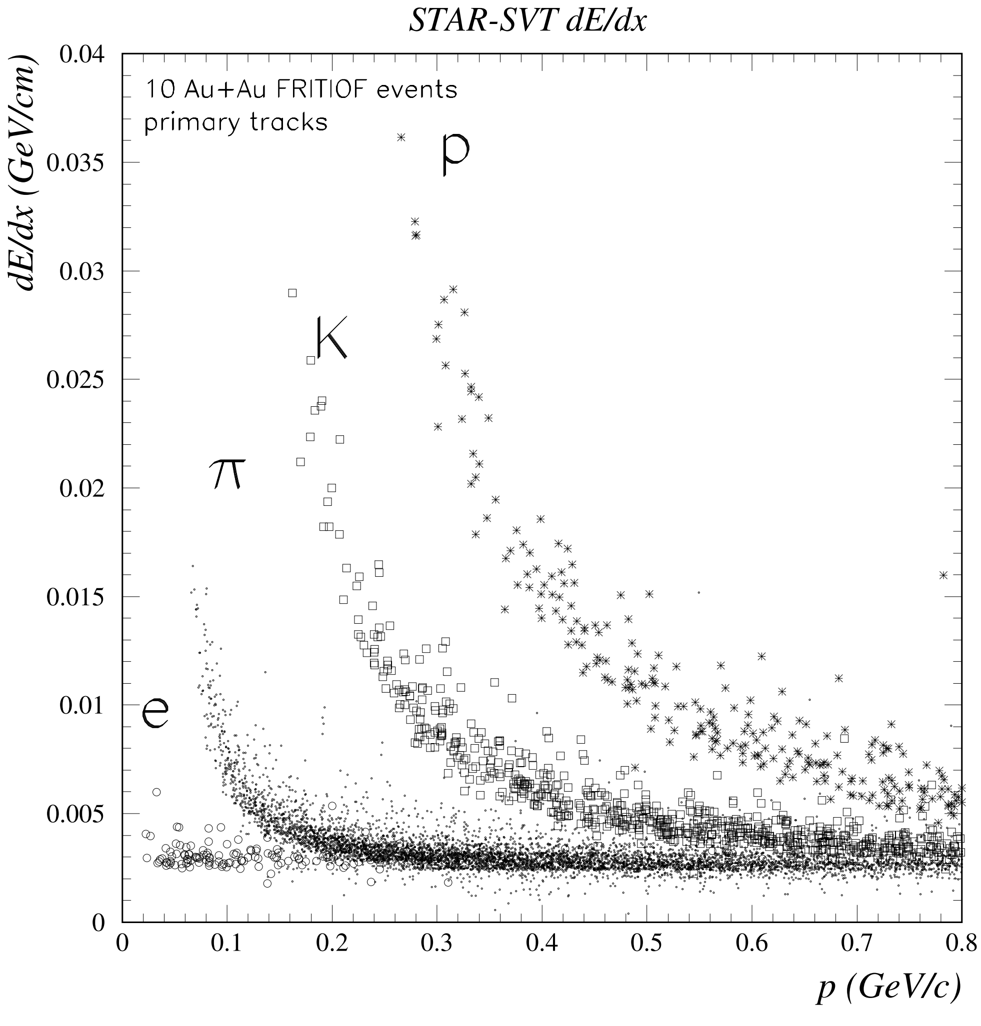

Figure 1.11: dE/dx vs. momentum for several particle species in STAR

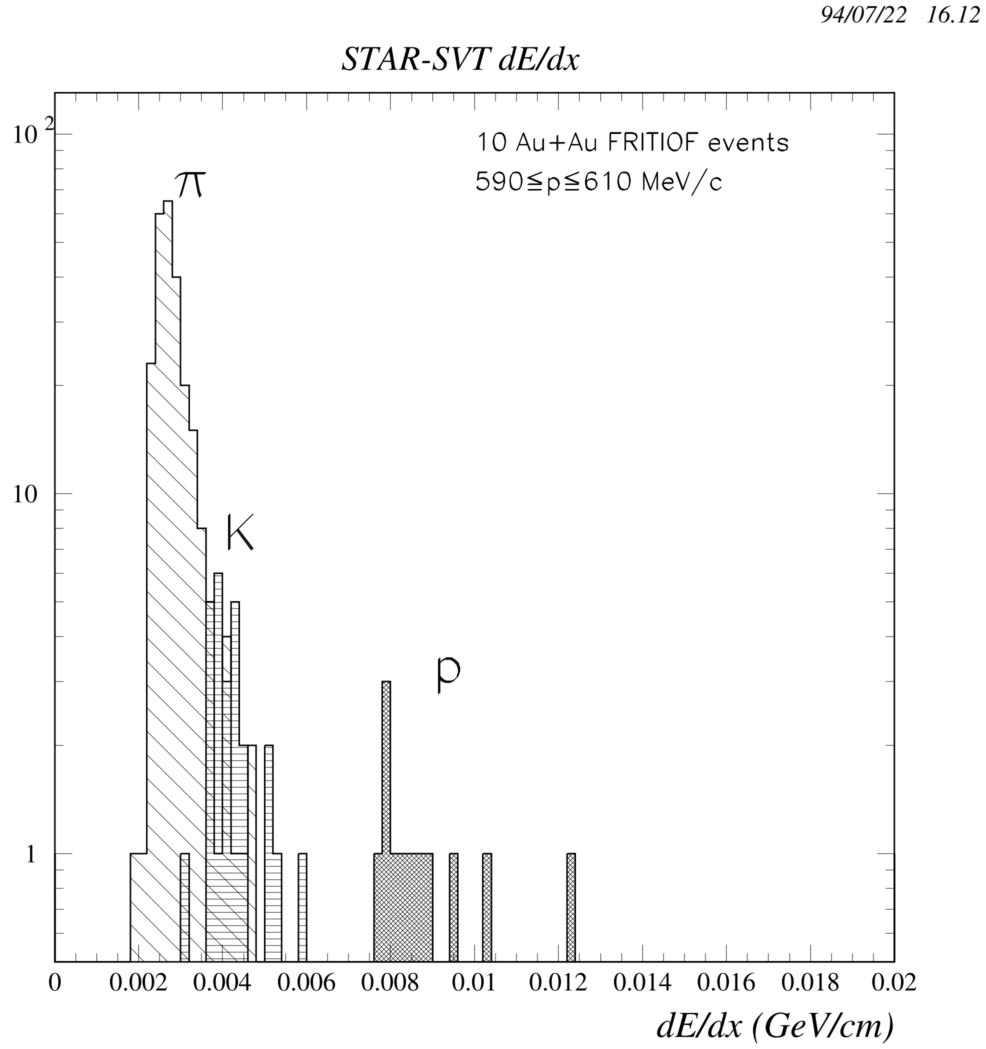

Figure 1.12: Energy loss in a small range of particle momentum

Two-track resolution: HBT analysis depends strongly on the ability to separate tracks that have small relative momenta. We show in fig. 1.13(a) the result of a laboratory measurement of the minimum relative distance needed for two hit separation on a Silicon Drift Chamber as a function of the size of the signal [33] . The SVT also contributes to enhance the pion pair statistics on an event by event basis, by extending the momentum range for reconstruction of pion tracks. Figure 1.13(b) shows the uncertainty of the RMS radius as a function of two track separation and lower momentum cutoff. The dashed curves correspond to the TPC separation range projected back to the outer SVT radius, the solid curve reflects the SVT separation power at the same point. The improved two track resolution of the SVT-TPC system will increase the maximum source size to which STAR is sensitive. Figure 1.13(c) shows that the SVT might enable HBT analysis on an event-by-event basis.

Secondary track reconstruction: The superior position resolution of the SVT enhances the identification and reconstruction of decay vertices, thus considerably reducing the combinatoric background to obtain a statistically significant sample of short lived particles such as ϕ, K_s^0 and Λ. This improvement results from the enhanced impact parameter resolution (see fig. 1.14). It is augmented by a constant SVT-TPC track matching efficiency of 86% in the momentum range from 200 MeV/c to 40 GeV/c. The high detection efficiency thus achieved will enable STAR to examine the correlation between the production cross section of these particles and other observables such as the total multiplicity, the transverse energy production of the particle's mean $p_\perp$ hence providing much more information than is available from inclusive distributions alone.

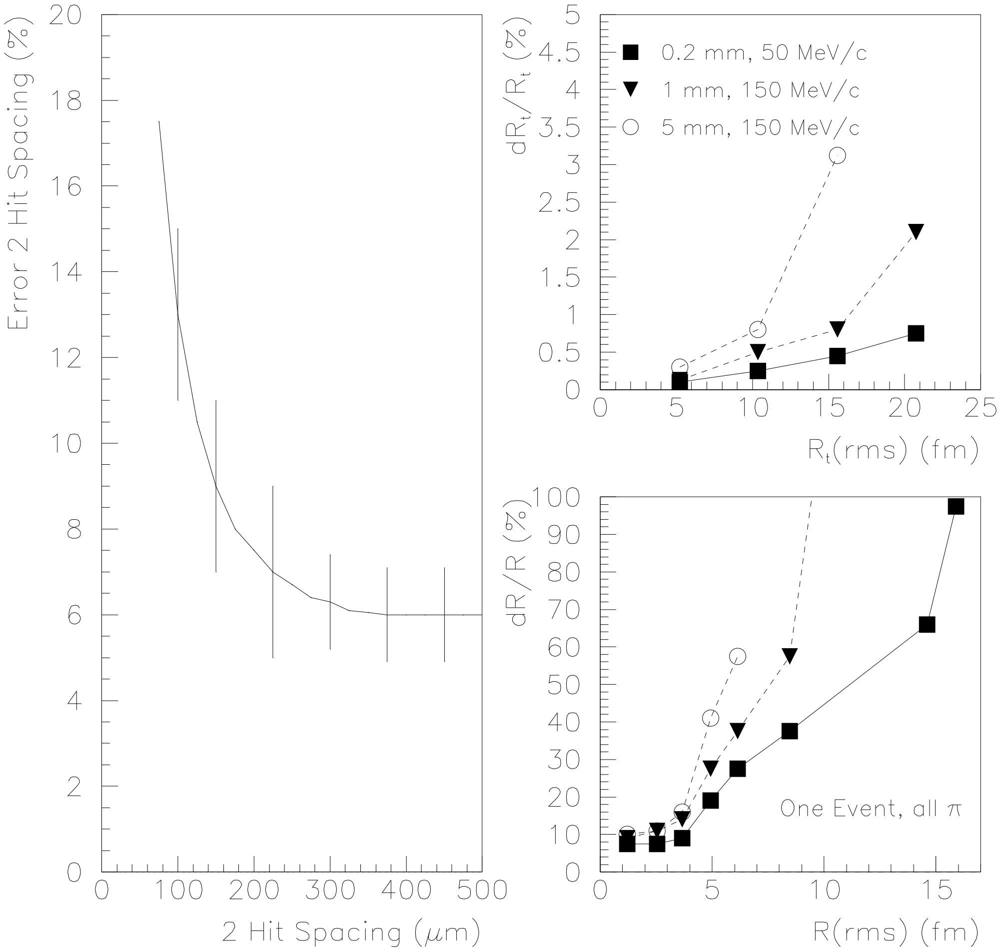

Figure 1.13: a.)Two hit separation as a function of signal width. b.) Uncertainty in rms radius for inclusive pion pair spectra as a function of lower momentum cutoff and two track separation at the outer radius of the SVT. Dashed curves reflect a range for the TPC performance depending on the drift gas, the solid curve is the result of the SVT analysis newline c.) Same as b.) for single event pion pair spectra.

High $p_\perp$: Momentum resolution for tracks with $p_\perp$ above 2 GeV/c is important for studies of gluon shadowing and jet quenching in the heavy-ion program as well as for some aspects of the pp spin physics program, in particular, the high momentum electron-hadron separation which is crucial for $W^\pm$ and Z^0 reconstruction and the study of Drell-Yan e^+e^- pairs. However, at high $p_\perp$ the momentum resolution of the TPC alone is seriously impaired by its modest position resolution ($\approx 700\ \mu$m), which limits the determination of the curvature of these nearly straight tracks. The addition of the SVT will provide vector information near the main vertex, greatly improving the momentum resolution for high $p_\perp$. Figure 1.15 shows the transverse momentum resolution as a function of $p_\perp$ for the TPC alone, and the TPC+SVT system. Note that at 50 GeV/c the improvement provided by the addition of the SVT is more than a factor of 5. Inclusion of the main vertex, as determined by the combined tracking system, will further improve the momentum resolution.

Primary vertex location: The SVT improves the main vertex position resolution and in turn augments primary particle momentum resolution, secondary particle identification, and high $p_\perp$ momentum resolution. A plot of primary vertex z resolution as a function of the mass of symmetric systems is shown in fig. 1.16 for the TPC and TPC+SVT systems. The addition of the SVT is of primary importance for the lighter systems. There is also a corresponding improvement in the two vertex separation from applying the Raleigh criteria to the results of the previous plot. This improved vertex resolution will be extremely useful in p-p collision studies at high luminosity where recorded events will include multiple primary vertices. It should also enhance STAR's ability to reject multi-vertex events when searching for events with rare multi-

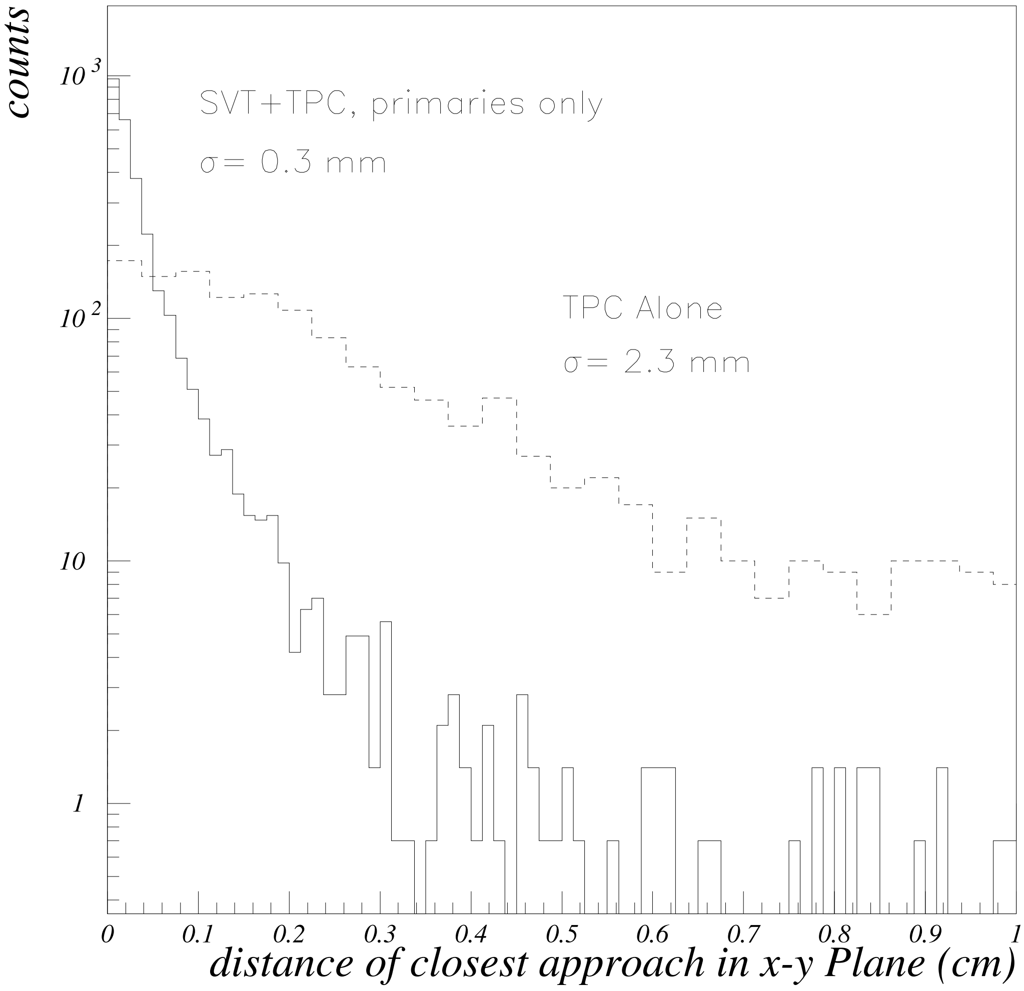

Figure 1.14: Impact parameter resolution from the SVT and TPC

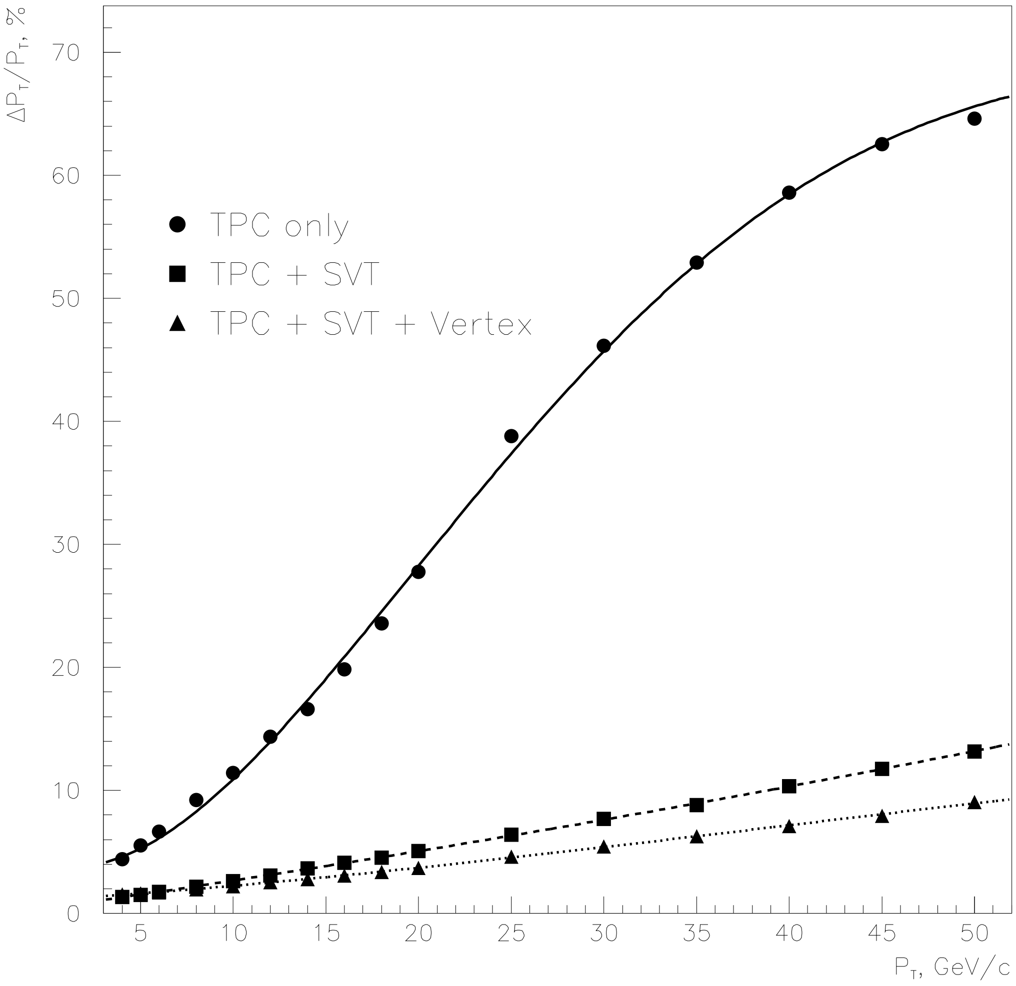

Figure 1.15: High $p_\perp$ resolution from the TPC alone and the SVT+TPC system

plicity fluctuations.

Trigger: The SVT will also provide third level trigger capabilities for STAR. Using existing algorithms, it should be feasible to achieve track reconstruction and particle identification via dE/dx with the SVT within a few tens of milliseconds. Thus, charged particle spectra are accessible on the third trigger level. Conditions could be envoked on e.g. particle ratio fluctuations and/or spectral shapes. One example of particular interest is the low momentum pion enhancement. Figure 1.17 shows the pion mean $p_\perp$ calculated event-by-event for 200 events for two different production mechanisms, as discussed in the previous chapter. The mean momentum is shifted by varying the chemical potential μ in the Boltzmann distribution. A μ near the pion mass corresponds to a condensation scenario in which the low momentum component of the spectrum is enhanced. Figure 1.17 shows the difference in the mean of the two spectra, which can be determined in time for a third level trigger decision.

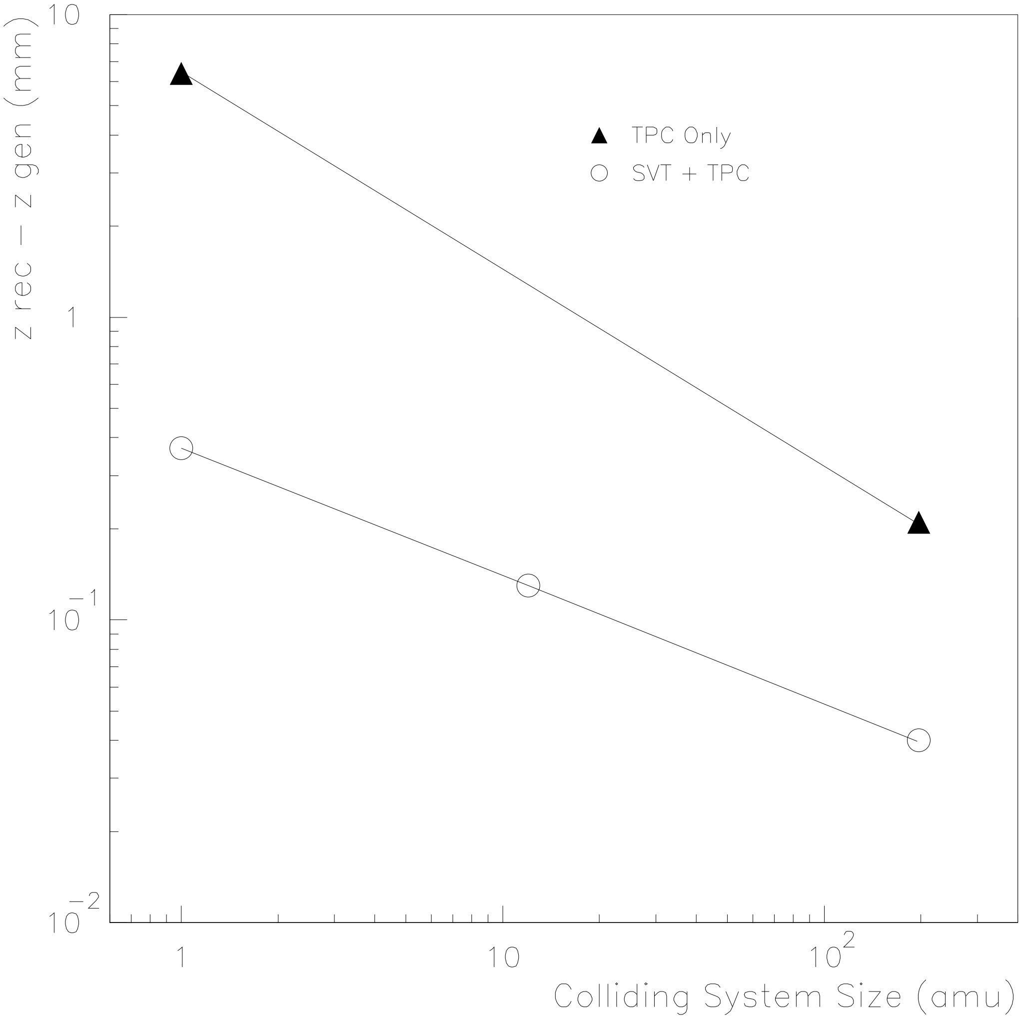

Figure 1.16: Main vertex resolution measured by the SVT and TPC as a function of system size

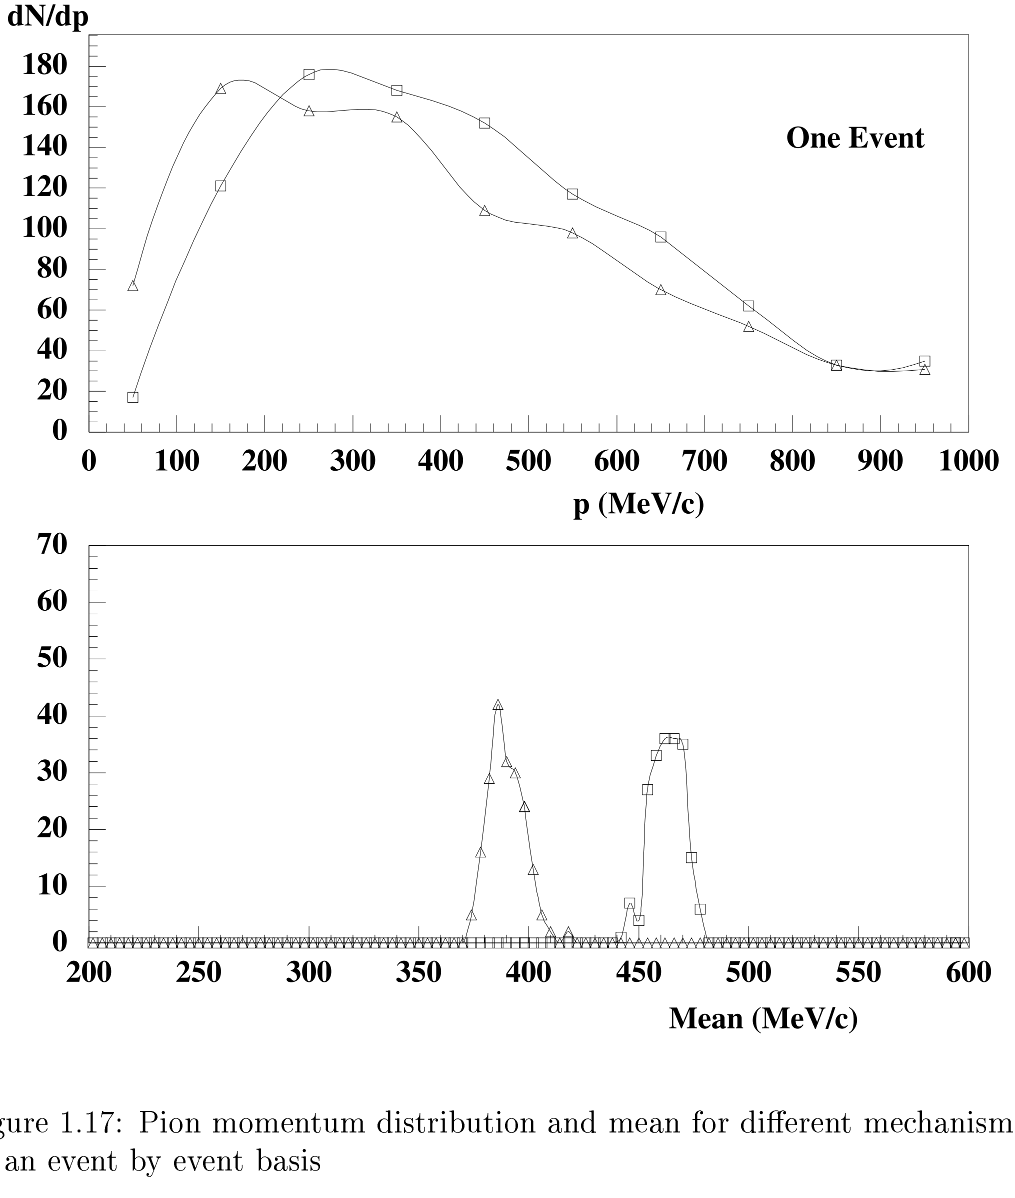

Figure 1.17: Pion momentum distribution and mean for different mechanisms on an event by event basis

Chapter 2

Principles and design of the Silicon Drift Detector

2.1 Principles of operation

2.1.1 Introduction

Silicon Drift Detectors (SDDs) combine some of the best features of conventional gas drift chambers (such as the relatively small number of instrumented channels) with some of the best features of solid state detectors (such as excellent position and energy resolution).

Conventional gas drift chambers consist of thin high-voltage wires surrounded by a gas. When a fast charged particle passes through the gas, ion pairs are created; the electrons drift one direction in the field applied by the high-voltage wires, while the ions drift oppositely. The electrons are collected on the wires (anodes), and the resulting signal is amplified. The drift time of the electrons is measured, and a velocity map can be used to determine the drift distance from the drift time. These detectors have a relatively small number of instrumented channels, but their position resolution is often as large as 250 μm.

Solid state detectors, such as silicon microstrip detectors, work in a similar fashion. A thin wafer of n-type silicon with a continuous n^+n junction on one side has implanted, upon the other side, a series of very closely spaced p-type cathodes. A reverse-bias voltage is applied to deplete the bulk silicon and create a drift field. When a fast charged particle passes through the silicon, ion pairs are created. The electrons drift towards the n^+n junction, while the holes drift towards the p^+ cathodes.

Silicon microstrip detectors have excellent position resolution, due mostly to the fine cathode pitch (on the order of 25 μm) allowed by modern silicon processing techniques. However, since each cathode requires its own amplifier, the number of channels of amplifiers required can be quite large. For instance, for a modest microstrip detector, with 25 μm pitch, a 5-cm detector would require 2000 amplifiers. Additionally, this setup would yield only one-dimensional position information; a second layer of microstrip detector would be required to determine two-dimensional position.

Silicon drift detectors, in comparison, operate in a manner similar to convetional gas drift chambers, except that the drift gas is replaced by n-type silicon [34]. When the passage of fast charged particles ionizes a small region in the SDD, the resulting free electrons drift in the presence of an externally applied electric field, until they are collected at an anode. The drift velocity of the electrons in the silicon is given by

$$v_{drift} = \mu E \tag{2.1}$$

where v_{drift} is the drift velocity, μ is constant known as the *electron mobility*, and E is the local electric field. Since the mobility is a constant (unless the bulk silicon of the detector suffers from non-uniformities in the doping), a constant electric field implies a constant drift velocity. Thus, measuring the drift time allows calculation of the drift distance. Also, the charge cloud spreads as it drifts, due to diffusion and coulomb repulsion. This spreading causes signals to appear not only on the anode closest to the intersection point, but also on anodes adjacent to that central anode. By analyzing signals from several anodes, it is possible to determine a position transverse to the drift direction with substantially greater accuracy than the anode pitch.

Figure 2.1 shows a schematic side view of a silicon drift detector. In this detector, the electric field is generated by a series of cathodes which are held at successively decreasing (negative) potentials. This decreasing potential (a potential gradient) creates an electric field. The same potential is applied to each pair of cathodes across the thickness of the wafer, which is held at ground potential by the collection anode (see below). The effect of the cathodes is fully to deplete the bulk silicon, creating a conduction region. This conduction region is the region where the signal electrons drift.

At the end of the drift region, the collection anode is held at ground, while the cathode directly across the thickness of the wafer is held at a negative potential larger than its nearest neighbor cathode. This is done so that all of the electrons in a cloud are channeled to the collection anode with reasonable efficiency and minimal spreading of the charge cloud in the direction normal to the wafer.

By segmenting the collection anodes, a position transverse to the drift direction can be obtained. The position resolution of the SDD is comparable to that of a silicon microstrip detector, but the number of channels of amplification and readout is comparable to that of a gas drift chamber. For example, the SDDs designed for STAR have a 6 cm x 6 cm active area, with 240 anodes,

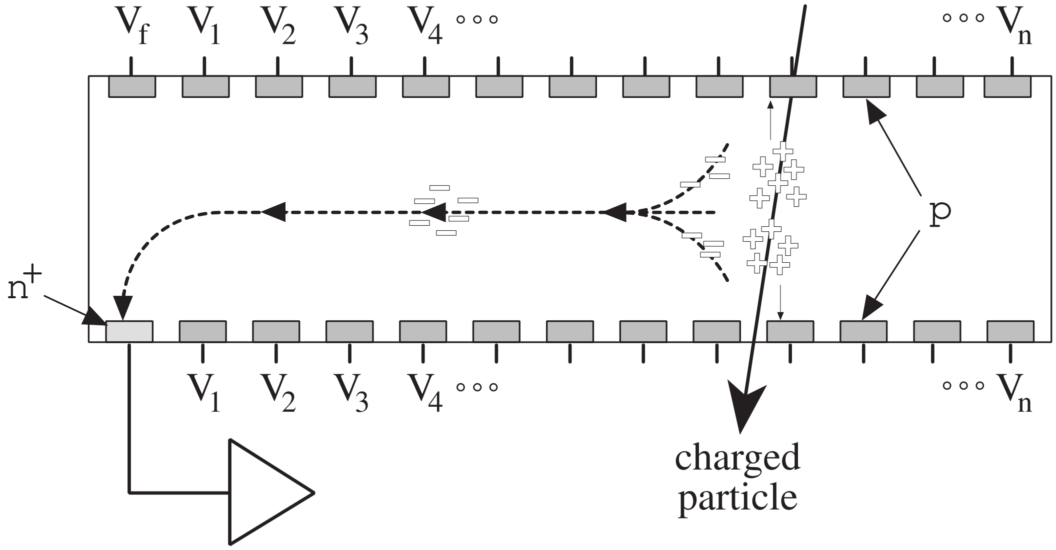

Figure 2.1: Cross-sectional schematic view of a silicon drift detector

spaced with 250 μm pitch. The position resolution of these devices is expected to be about 25 μm, which is similar to the resolution of the microstrip detector (although not quite as good), but much better than that of most gas drift chambers. However, the silicon drift detector requires much more time for the signal to drift to the anode (about 6 μs for a 3 cm drift in the STAR SDDs) than in silicon microstrip detectors. This means that the SDDs cannot be used for extremely high rate experiments.

2.1.2 Silicon drift detectors as a p-n-p junction

In order to understand how the cathodes deplete the bulk silicon and how this affects the operation of SDDs, it is useful to understand the physics of a p-n junction.

Figure 2.2 shows a schematic view of a p-n junction, after [36]. P-type silicon is produced by “doping” pure silicon with impurities of an element which has one less valence electron than silicon. Boron is the element usually chosen. N-type silicon contains impurities with one *more* valence electron than silicon (usually phosphorus). When contact is made between a sample of p-type silicon and a sample of n-type silicon, recombination of electrons from the n-type region with holes in the p-type region takes place. This recombination is possible because the electrons are thermally excited into the conduction band of the silicon, and become mobile.

Since both silicon samples were electrically neutral before being brought into contact, the result of this recombination is that the “extra” electrons from the n-type silicon end up on the ‘p’ side of the junction, leaving a net negative charge on the ‘p’ side and a net positive charge on the ‘n’ side. This local “space charge” creates a local electric field which eventually grows until further recombination of holes and electrons is energetically forbidden. The space charge region is known as the *depletion layer,* since all mobile charge carriers in the region have been removed, rendering it non-conductive. This has practical implications for the operation of the silicon drift detector, which will be discussed below.

N-type silicon is typically doped with a much lower impurity density than the p-type region; charge conservation then requires that the depletion layer extends farther into the n-type region than in the p-type region. This relationship can be quantified [35]:

$$N_a x_p = N_d x_n \tag{2.2}$$

where N_a is the density of acceptor impurities in the p-type region (in cm^{-3}),

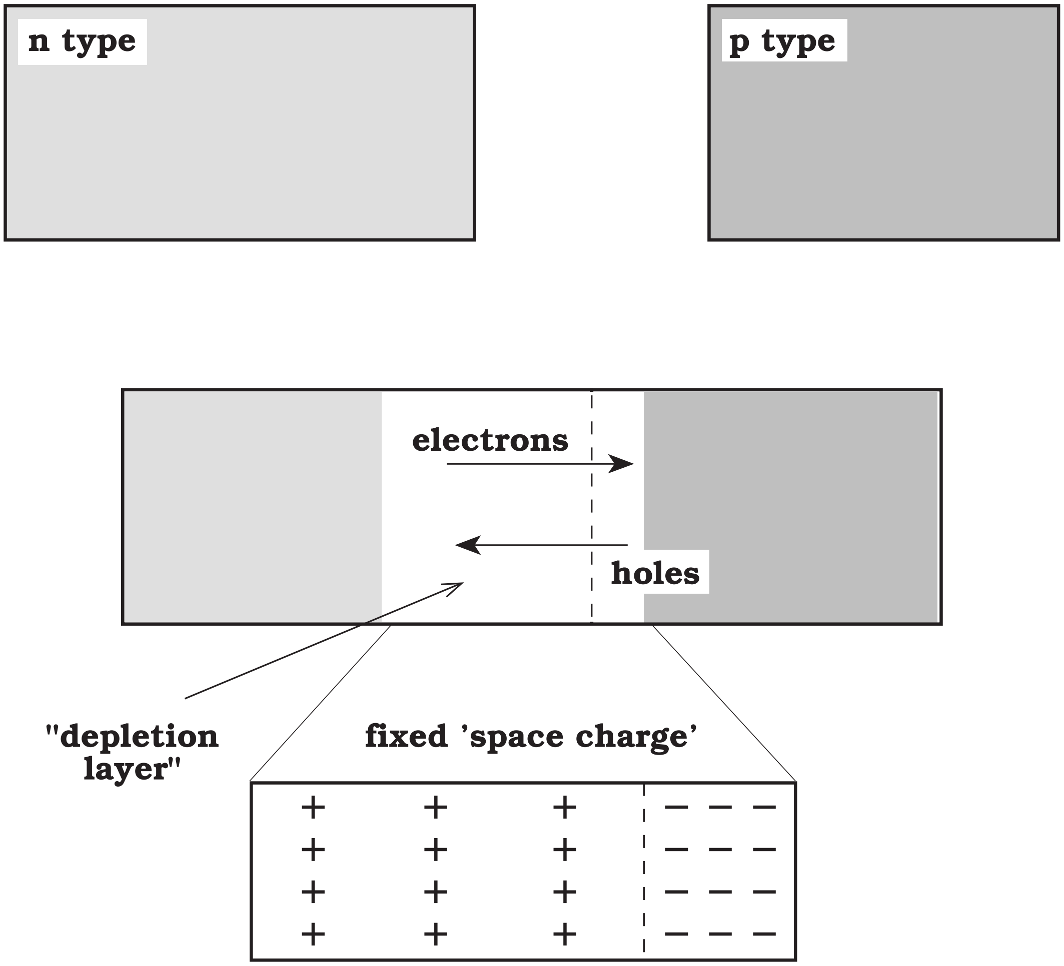

Figure 2.2: P-N junctions and depletion layers

x_p is the extent of the depletion layer in the p-type region, N_d is the density of donor impurities in the n-type region, and x_n is the extent of the depletion layer in the n-type region.

We can calculate the local electric field described above. In the region far away from the junction, the p-type silicon will be charge-neutral. This can be written as

$$N_a - N_d + p - n = 0 \tag{2.3}$$

where p and n are the concentration of holes and free electrons, respectively.

Since $N_a >> N_d$ in the p-type region (essentially, one can consider N_d to be zero), and $n = 0$, we can identify

$$N_a = p \quad . \tag{2.4}$$

This is reasonable, considering that the ionization energy for the holes in the p-type region is a few tenths of an eV at most, while the band gap for silicon is roughly 1.12 eV. From this, one concludes that most of the holes (and free electrons) come from the impurities, rather than the intrinsic silicon.

Then we can write down an expression for the hole concentration

$$p = n_i e^{(E_i - E_F)/k_B T} \tag{2.5}$$

where E_i is the intrinsic energy of the silicon, E_F is the Fermi energy of the silicon, and n_i, the carrier (hole or electron) concentration of the silicon, is a tempera-ture-dependent density characteristic of the material. For room temperature (300 K) silicon, $n_i = 1.45 \times 10^{10}$ cm^{-3} . Finally, we can write down an expression [41] for the potential in the bulk p-type silicon relative to the Fermi level by combining equations 2.4 and 2.5:

$$\phi_p = -\frac{1}{q}(E_i - E_F) = -\frac{k_B T}{q} \ln \frac{N_a}{n_i} \quad . \tag{2.6}$$

Similarly, the potential in the n-type silicon, far from the junction, can be derived as

$$\phi_n = -\frac{1}{q}(E_i - E_F) = \frac{k_B T}{q} \ln \frac{N_d}{n_i} \quad . \tag{2.7}$$

Finally, we can solve for the so-called *built-in potential* Φ_{bi}:

$$\Phi_{bi} = \phi_n - \phi_p = \frac{k_B T}{q} \ln \frac{N_a N_d}{{n_i}^2} \quad . \tag{2.8}$$

For the silicon used in the STAR SDDs, the doping of the n-type bulk silicon is 1.0×10^{12} cm^{-3}, while the doping of the p-type field cathodes is 1.0×10^{18} cm^{-3}, yielding a built-in junction potential of approximately 0.58 V.

It is possible to calculate the depth of the depletion layer for a given junction. For this calculation, we will use the *abrupt junction* approximation, which considers the p-n junction to be infinitesimally thin; the silicon is doped only with donors when $x < 0$, and only with acceptors when $x > 0$.

First it is useful to note that the depth of the full depletion layer can be written as

$$W = x_p + x_n \quad . \tag{2.9}$$

In the depletion layer there is a space charge equal to the density of impurity ions (assuming that all impurities have been ionized). This, in turn, yields a differential equation for the electric potential (from Poisson's equation):

$$\frac{d^2\phi}{dx^2} = +q\frac{N_a}{\epsilon_{Si}} \qquad \text{for } -x_p \le x < 0 \tag{2.10}$$

where ϵ_{Si} is the absolute dielectric constant in the silicon, approximately 1.05×10^{-14} F/cm. Similarly,

$$\frac{d^2\phi}{dx^2} = -q\frac{N_d}{\epsilon_{Si}} \qquad \text{for } 0 < x \leq x_n \ . \tag{2.11}$$

The electric field $\mathcal{E}$ can be obtained by integrating equations 2.10 and 2.11, yielding

$$\mathcal{E}(x) = -\frac{d\phi}{dx} = -\frac{qN_a}{\epsilon_{Si}}(x + x_p) \qquad \text{for } -x_p \leq x < 0 \tag{2.12}$$

and

$$\mathcal{E}(x) = -\frac{d\phi}{dx} = +\frac{qN_d}{\epsilon_{Si}}(x - x_n) \qquad \text{for } 0 < x \leq x_n \ . \tag{2.13}$$

Note that the maximum electric field is found at $x = 0$, where

$$\mathcal{E}_{max} = \frac{qN_a}{\epsilon_{Si}}x_p = \frac{qN_d}{\epsilon_{Si}}x_n \quad . \tag{2.14}$$

Integrating equations 2.12 and 2.13 yields the total potential difference from one side of the depletion layer to the other (namely the built-in potential Φ_{bi}) :

$$\begin{aligned}
\Phi_{bi} \quad &= -\int_{-x_p}^{x_n} \mathcal{E}(x)dx \\
&= -\int_{-x_p}^{0} \mathcal{E}(x)dx - \int_{0}^{-x_n} \mathcal{E}(x)dx \\
&= \frac{qN_a x_p^2}{2\epsilon_{Si}} + \frac{qN_d x_n^2}{2\epsilon_{Si}} \\
&= \frac{1}{2}\mathcal{E}_{max}W \quad .
\end{aligned} \tag{2.15}$$

Recalling the definition of $\mathcal{E}_{max}$, one can cast x_p and x_n in terms of $\mathcal{E}_{max}$:

$$x_p = \frac{\epsilon_{Si}}{q}\mathcal{E}_{max}\frac{1}{N_a} \tag{2.16}$$

and

$$x_n = \frac{\epsilon_{Si}}{q}\mathcal{E}_{max}\frac{1}{N_d} \quad . \tag{2.17}$$

Then we can rewrite equation 2.9 as

$$W = \frac{\epsilon_{Si}}{q}\mathcal{E}_{max}\left[\frac{1}{N_d} + \frac{1}{N_a}\right] \quad . \tag{2.18}$$

Thus

$$W = \frac{\epsilon_{Si}}{q}\mathcal{E}_{max}\left[\frac{N_a + N_d}{N_a N_d}\right] \tag{2.19}$$

and

$$\mathcal{E}_{max} = W\frac{q}{\epsilon_{Si}}\left[\frac{N_a N_d}{N_a + N_d}\right] \quad . \tag{2.20}$$

Inserting eq. 2.20 back into eq. 2.15 yields

$$\Phi_{bi} = W^2\frac{q}{2\epsilon_{Si}}\left[\frac{N_a N_d}{N_a + N_d}\right] \quad . \tag{2.21}$$

This lets us solve for the depletion depth W:

$$W = \sqrt{\frac{2\epsilon_{Si}}{q}\left[\frac{N_a + N_d}{N_a N_d}\right]\Phi_{bi}} \quad . \tag{2.22}$$

Using this formula for the STAR SDDs, with the values of N_a and N_d given earlier, the unbiased depletion layer has a depth of approximately 28 μm.

By applying a bias potential to the p-n junction, it is possible to allow electrons from the n-type region to reach the p-type region, thus increasing the depth of the depletion layer. When the bias potential is large enough, electrons are free to move anywhere within the silicon, and thus the depletion layer grows to include all of the silicon. In this state, the silicon is “fully depleted;” there are no free charge carriers available.

The depth of this layer can be calculated using an approximation to eq. 2.22. The first observation is that the p-type regions are doped much more heavily than the n-type region (see the values for N_a and N_d given earlier).

Also, we will be applying an external bias voltage, which will replace the Φ_{bi} term in eq. 2.22. Thus we arrive at an expression for the depth of the depletion layer in terms of the bias voltage V_{bias} and the doping density N_d:

$$W \simeq \sqrt{\frac{2\epsilon_{Si}V_{bias}}{qN_d}} \quad . \tag{2.23}$$

Thus the depth of the depletion layer is a constant times the square root of the bias voltage.

Silicon drift detectors can be considered to be two back-to-back p-n junctions. There is a depletion layer just underneath each p-type region (cathode), extending for some distance into the bulk silicon. Beyond this, the n-type silicon in the center of the wafer remains conducting. Conducting silicon can carry currents from thermal generation of electrons, which would show up as noise in the detector. Further, the conducting region would reduce the active area of the detector. Clearly it is desirable to deplete the entire thickness of the wafer, in order to reduce the volume of noise-generating silicon and increase the volume of active detector.

In order to do this, the bulk silicon is held at ground. Then identical negative potentials are applied to cathodes on either side of the silicon wafer. As the applied potential grows, the depletion layer also grows, until the silicon between the cathodes has become fully depleted. The potential in the bulk silicon can be calculated explicitly by integrating eq. 2.13, but clearly it has a quadratic dependence on the distance from the p-n junction interface.

Figure 2.3 shows a schematic view of a cross-section of silicon wafer, containing just one cathode pair. With no bias potential applied, small depletion layers form just beneath each cathode. The majority of the bulk silicon is

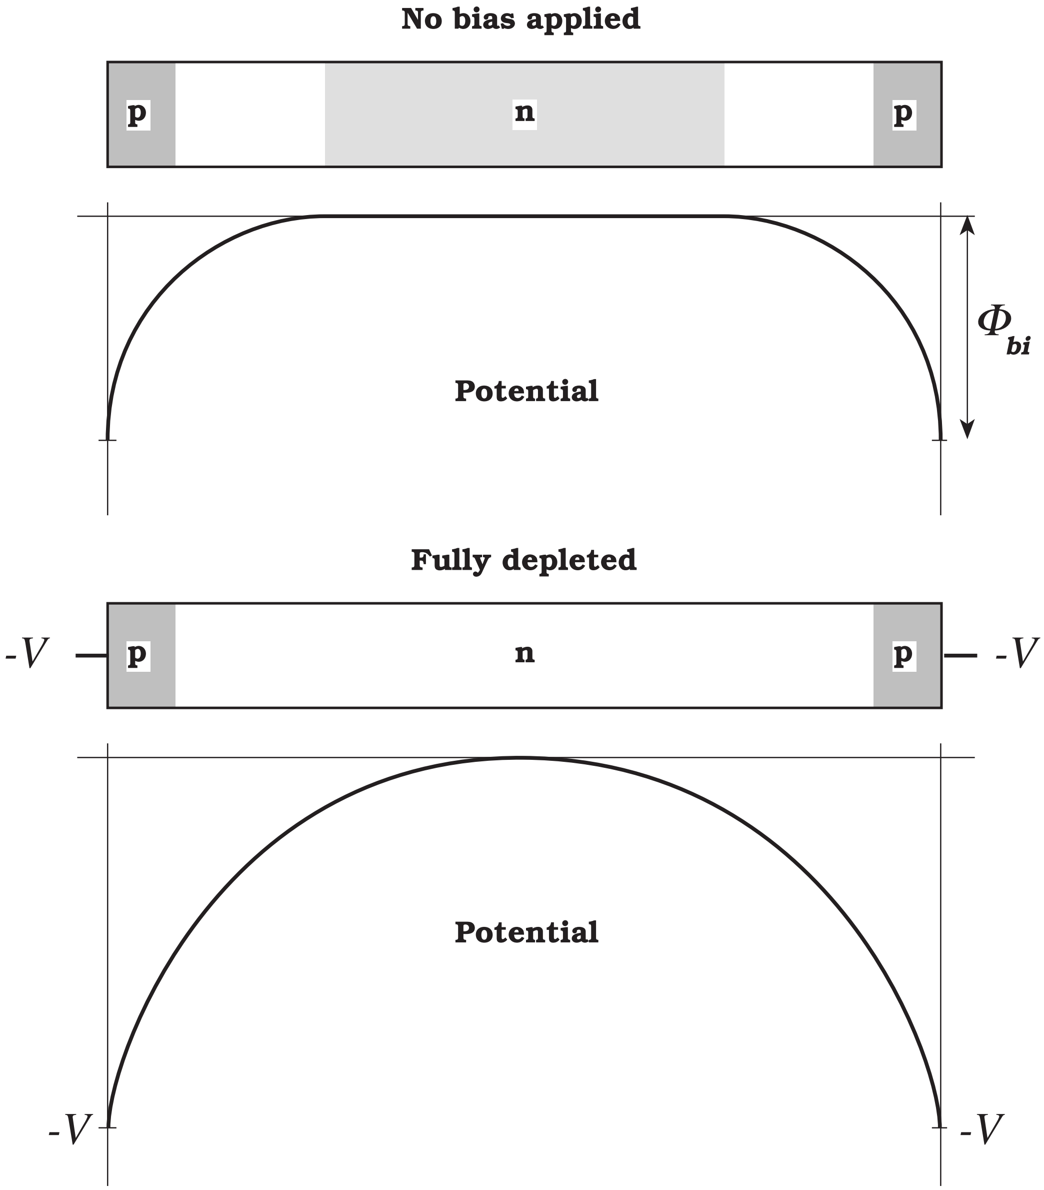

Figure 2.3: The silicon drift detector interpreted as a partially or fully depleted p-n-p junction

undepleted, leaving a conduction layer in the center of the wafer. If sufficient negative potential is applied to the cathodes, the depletion layer grows until the bulk silicon is fully depleted, leaving no conducting silicon except at the surface.

The "transport" electric field is applied as described earlier, by increasing the (negative) potential applied to the cathodes with distance from the anode.

In the case of the STAR SDDs, it is desirable to estimate the depth (in Volts) of the potential well in the depleted bulk silicon. We can start with eq. 2.23 and solve for V_{bias}, obtaining the maximum potential difference (between the field cathodes and the minimum potential in the bulk silicon):

$$\Delta V = \frac{qW^2}{2\epsilon_{Si}} N_d \quad . \tag{2.24}$$

To use this formula to calculate the potential difference, we must note that W in eq. 2.24 is one-half the thickness of the wafer (since we are modelling the SDD wafer as two back-to-back p-n junctions). It is also helpful to note that the expression shown is in the form of a constant times the doping density.

The design specification for the silicon used in the prototype STAR SDDs was $N_d = 1.0 \times 10^{-12}$ cm^{-3}, which would yield a potential "valley" with a depth of approximately 17.1 V. However, the silicon delivered from the manufacturor had a lower doping, of only $N_d \simeq 0.1 - 0.3 \times 10^{-12}$ cm^{-3}, which yielded a potential valley of at most 5.1 V. This left the bulk silicon only partially depleted, but there was an additional side effect. The potential minimum is in the center of the wafer thickness only if the field cathodes have identical applied potentials on each side of the wafer. The potential valley as designed was 17 V deep, which allowed for several volts of difference between facing cathodes. However, the delivered silicon, with its small potential valley (as little as 1.7 V!) required much more stringent control over the applied potentials than was possible. Small potential differences (as little as 2 V) between facing cathodes caused the electron cloud to drift to the surface of the silicon rather than the collection anode. As a result, the prototype STAR SDDs performed inadequately, with complete signal loss in some of the detectors. The next batch of silicon delivered by the vendor had a doping density within the specified limits, and the detectors performed as designed.

Figure 2.4 shows a 2-dimensional map of the electric potential over a small region close to the collection anode. This map was calculated by solving Poisson's equation (eqs. 2.10 and 2.11), given the known bulk quantities of the bulk silicon of the SDD wafers (dielectric constant, doping density, etc.) and knowing the boundary conditions (the applied potentials at the cathodes and the collection anode). For illustrative purposes, the quantity mapped is actually the negative of the electric potential.

The level, flat line segments along the near and far sides (0 and 275 μm depth) are the field cathodes. The large level line at 275 μm depth is the "focussing cathode." The gutter-like valley at 0 μm depth, between 950 μm and 1150 μm length is the collection anode. Over most of the cathodes the potential is identical on either side of the wafer; however, the 20 cathodes closest to the collection anode (the region pictured here) are biased with slightly unequal potentials so that the electron cloud will be "steered" towards the collection anode.

It is easy to see how the electron cloud from a charged-particle pas-

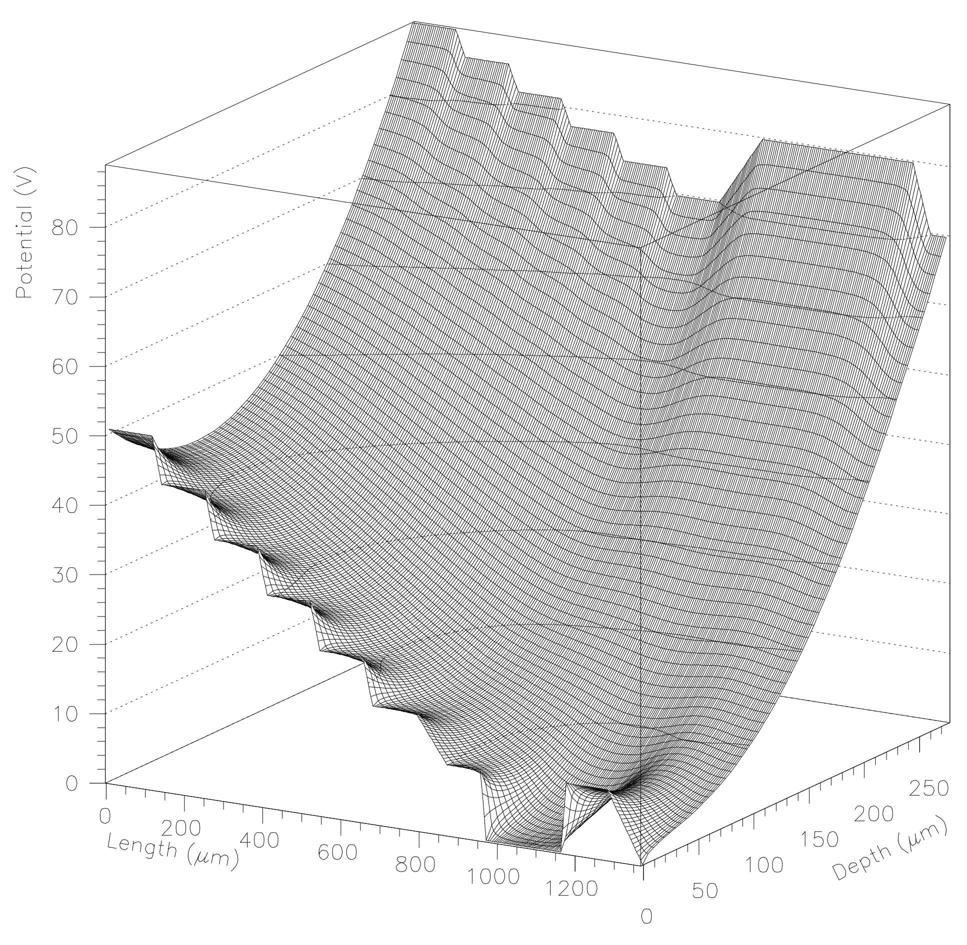

Figure 2.4: Map of the electric potential in a silicon drift detector

sage will drift “downhill” along the length of the detector, and then into the collection anode.

Chapter 3

Testing SDDs in a particle beam

In July, 1994, the STAR SVT collaboration used the B test beam area of the AGS at Brookhaven National Laboratory for in-beam studies of the operation and performance of prototype STAR-1 silicon drift detectors.

3.1 The experimental setup

3.1.1 Overview

The objective was to place a small number of SDDs in a particle beam of reasonably well-known composition in order to study signal and noise levels, detector efficiency, and position resolution. By placing three detectors in the beam, it is possible, in principle, to assemble hits in the detectors into tracks (which should be straight, since there was no strong magnetic field present in the experiment). These tracks can then be used to measure the distance between the track and each hit in the track; the resulting "residues" give a measure of the position resolution of the detectors.

Figure 3.1 shows a schematic layout of the experimental setup used during the SVT test run.

The AGS test beam contains several different particle species. Selection of particle charge and momentum is done via upstream bending magnets. An

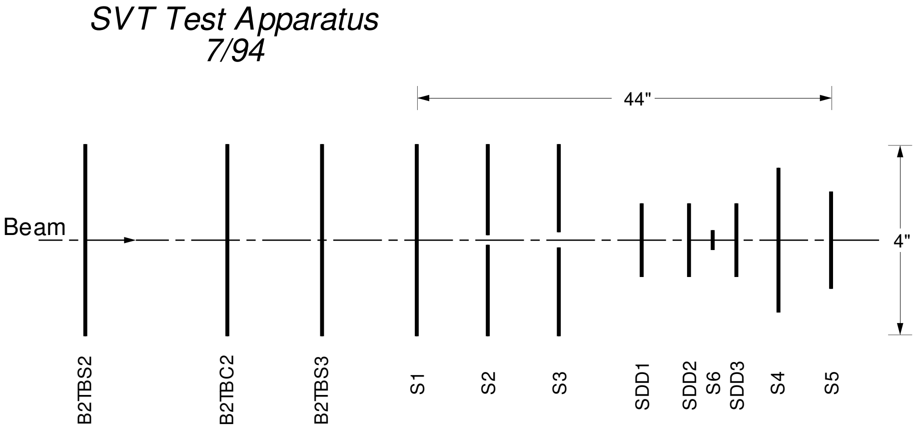

Trigger = (B2TBS2•B5TBS3)•($\overline{\text{B2TBC2}}$)•S1•$\overline{\text{S2}}$•$\overline{\text{S3}}$•S4•S5•S6

Figure 3.1: Schematic view of the SVT test run experimental setup

inline Cerenkov counter can tag electrons. As a result, the SVT collaboration was able to select particles with negative charge, and then veto electrons with the Cerenkov counter. The remaining particles were almost all pions; tuning the upstream bending magnets allowed the collaboration to select momenta of interest. Similarly, it was possible to select protons in the beam by switching to a positive-charge setting for the magnets.

3.1.2 Readout and data acquisition

Figure 3.2 shows a schematic of the readout chain and the data acquisition system. The first element in the chain is the pre-amp board. The pre-amp boards reside on the carrier board which holds the SDD itself. The pre-amps are placed as close to the anodes of the SDD as possible in an effort to minimize the total noise in the system. From the pre-amp boards the signal is carried via a ribbon cable to a shaping amplifier. The shaping amps reside in an external box attached to the experimental apparatus. From the shaping amps, the signal is transferred to a readout board.

The readout board contains a switched capacitor array (SCA) and an analog to digital converter (ADC). The SCA is a device for taking “snapshots” of the incoming analog signal. It consists of a large number of small capacitors, and a transistor switch network which changes state in response to a clock signal. At each instant, the SCA connects the input signal to one of the capacitors. When the next clock cycle arrives, the switch changes so that the input signal is connected to the “next” capacitor in the array. In this fashion, the input signal is sampled at discrete times while the SCA cycles. The RC time constant of the capacitor is roughly 100 picoseconds, so the capacitor is fully charged (to the same potential as the input signal) after only 1 ns or so. Further, the gate to the capacitor is open for only a few ns after the start of the clock cycle. Thus the charge stored on the capacitor does not represent an instantaneous sample of the input signal, but rather something like a time-average of the input signal over the few ns that the charging gate is open. However, for signals which do not change greatly over the time that the gate is open, the SCA comes close to providing an instantaneous sample.

The main purpose of the SCA is to sample the incoming analog signal and hold it long enough for the ADC to digitize the signal. Since the ADC takes up to a microsecond to digitize each sample, some method of holding analog signals until the ADC can digitize them is required. The SCA fills this role.

In the test-run readout system, the SCA is started in response to an

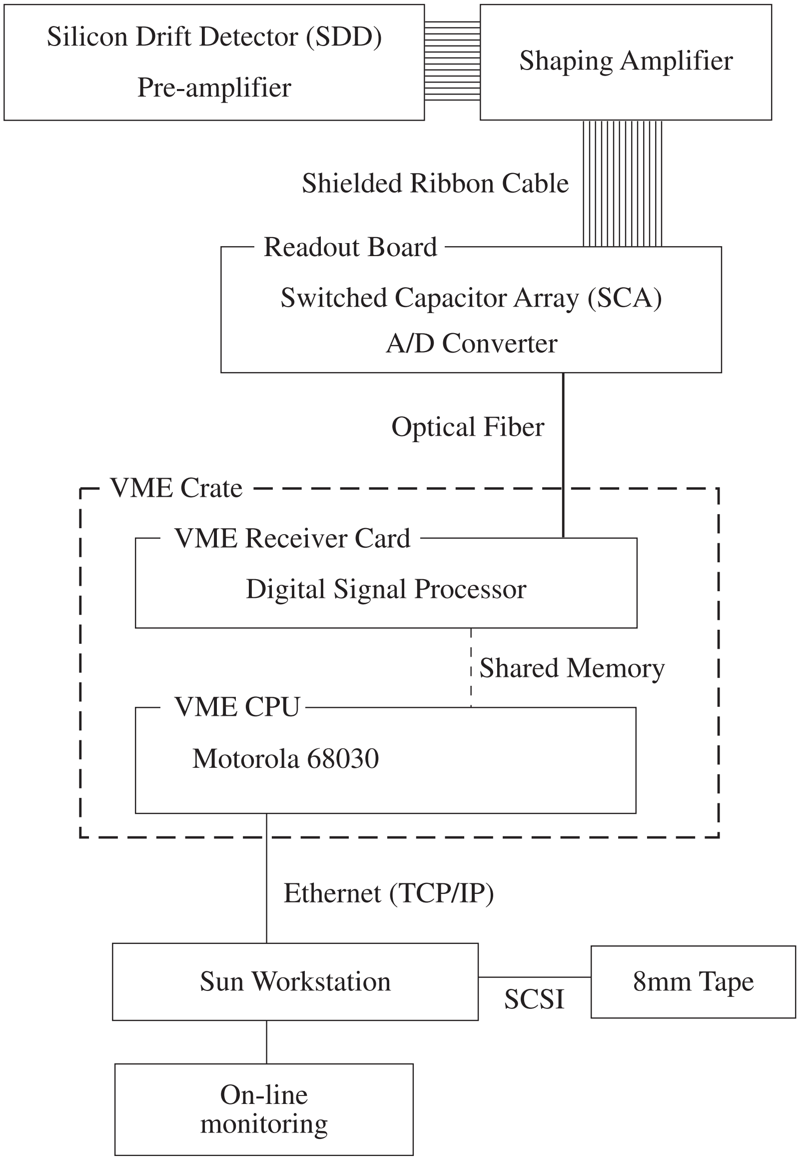

Figure 3.2: Schematic overview of the test run data acquisition system

event trigger. It takes a fixed number of clock cycles to start (1.2 μs in our case). Further, the detector events are asynchronous with respect to the SCA clock, which means that the actual delay between the event and the start of the SCA is 1.2 μs plus a random interval which could be as much as one clock cycle (25 ns). Thus the DAQ system has a "dead time" of interval $1.2\mu s < t <= 1.225\mu s$. Since the detector itself is operational during this time, the charge cloud produced by the passage of the particle through the detector is drifting during the dead time. This translates to a "dead area" in the detector, positioned adjacent to the readout anode. The dead area stretched from the anode to a distance equal to the dead time multiplied by the drift velocity. For our experiment, the dead area was approximately 1.2 μs $\times$ 6 mm/μs= 7.2 mm in length. As described at the end of this section, the actual dead time was measured for each event, and the resulting position offsets were calculated.

In the final STAR system, the SCAs will cycle continuously, operating in a "ring buffer" mode. When the end of the capacitor array is reached, the SCA will simply start over, writing at the beginning of the array. When a "stop acquisition" signal arrives (in response to an event trigger), the SCA will stop "spinning," and await a "readout" signal. Thus the STAR SCAs will provide a backward-looking series of samples, with no dead time.

Once the ADC has digitized the samples stored in the SCA, the data is passed on to a receiver card. The link between the readout card and the receiver card is via an optical fiber. The receiver card is a VME module, residing in a VME crate. The receiver card stores the incoming data in a buffer, where the data await processing by a digital signal processor (DSP) chip, a Motorola DSP

56000. In principle, the DSP could be used to perform pedestal subtraction or other processing. In this case, however, the DSP simply copied the data from the buffer memory into global VME memory, visible to other devices on the VME bus.

Once the data is written to VME memory, it is available for reading by any device on the VME bus. In the test run experiment, a Motorola MVME-167 VME controller, an MC 68040-based computer running the VxWorks operating system, was used to buffer the event data in order that the events could be written to tape.

A buffer manager, based on OS-9 code developed for CERN experiment NA35, created a ring buffer which could store up to 10 events. Clients of the buffer manager call the manager and request to be added to a list of clients which should see "any" event, or "all events." Then registered clients can call the buffer manager and request to be served an event. The buffer manager keeps track of which events have been seen by which clients, and serves the proper event to each client which requests an event. Once an event has been "seen" by all clients who requested "every event" service, the event is released (i.e. marked so that it may be over-written by new events).

In the case of the SVT test run, the only clients served by the buffer manager were a "taper" client, which wrote events to 8mm tape, and the *"Pad Monitor"* program, which was developed as a diagnostic tool for visualizing individual events. Both of these clients operated remotely, from a Sun SPARCstation LX workstation. The workstation was connected to the VME controller by an ethernet cable, using TCP/IP as the network protocol. The 8mm tape drive was connected extenally to the workstation's SCSI bus.

The entire detector readout chain, from the PASA and SCA readout boards to the receiver cards, consisted of equipment designed and built for NA35. A limited number of PASA chips were available, so only 8 anodes of each of three detectors were instrumented. As a result, only one readout board was required to process the 24 resulting signals, and only one receiver card slot was used in the VME crate. Thus, the output from the experiment came packed into a single receiver board's output. The signals from each SDD were packed into one 128-anode output array, separated by many "dead" or "blank" channels.

The test-run DAQ system also included several scalers, ADCs and TDCs. The scalers were used to count total beam spills, event triggers, etc. The ADC modules were used to measure integrated charge from the beam scintillators, and the TDC modules were used to obtain timing information for trggering.

One TDC channel in particular was important; the start signal for it was the event trigger, while the stop signal came from the SCA clock. In other words, this TDC measured the random delay mentioned in the description of the dead time inherent in our SCA unit. By adding the random delay measured by the TDC to the known 1.2 μs delay in the SCA startup, it was possible to obtain a correct measurement for the total drift time of each hit in the SDDs. This is of obvious interest if one wishes to obtain the best position resolution possible, since the measured position of the hit is dependent on the drift time.

3.2 Data analysis software

3.2.1 Overview

The software written to analyze the test-run data was designed to mimic the way final STAR SVT software will operate. Some liberties were taken because many details of the operation of STAR software were not available at the time the test-run software was written.

The goal of the SVT Cluster-finder, Analyzer, and Tracker (SCAT) software was to take the output from the DAQ system and analyze it. This analysis consisted of performing pedestal subtraction, finding sequences in the raw data stream which are above threshold ("zero suppression"), assembling these sequences into clusters which represent a single particle crossing ("cluster-finding"), analyzing these clusters, and performing basic track-finding and analysis when possible.

SCAT operates in a very simple loop:

1. Get an event (from tape or from the buffer manager)
2. Check if the event was a pedestal event: if it is not a pedestal event and no pedestals have been encountered yet, loop back to position 1.
3. If it is a pedestal event, store the event in the pedestal buffer, and loop back to position 1.
4. Processing cannot reach this point unless the event is not a pedestal event (i.e. it's a "data" event), and a pedestal already exists in the buffer. Subtract the stored pedestal from the current event.

5. Unpack the event into structures for each detector. This is required because there are data from three detectors packed onto the readout board.
6. Loop over detectors:
7. Loop over the active anodes, calling the zero suppressor for each anode. The result from this stage is a table of HIT structures for each detector.
8. Call the cluster-finder to find clusters in the HIT table. This cluster-finder also analyzes the clusters for centroids, etc.
9. When all detectors have been processed, call the tracking routine, which finds tracks, fits them with straight lines, and calculates residues.
10. Write out required information and loop back to position 1.

3.2.2 Detailed description of the routines in SCAT

The zero suppressor developed for SCAT finds “sequences” of ADC values above a user-supplied threshold, and uses them to fill a HIT structure. The HIT structure contains information such as the detector and anode on which the sequence was located, the start and end times of the sequence, and a copy of the actual ADC values in the sequence. This differs somewhat from the information which will be present in the STAR software arrangement, but the principle is similar.

Sequences are located using a simple single-threshold algorithm. This algorithm requires that a certain number of consecutive ADC values be above threshold. In the case of the SVT test run, the requirement was that three

ADC values must be above a threshold of three ADC counts (this corresponds to 30 mV).

Once a sequence has been located, the information described above is copied into a HIT structure, and the process continues. The “start” and “end” values are numbers which represent the number of time buckets since the start of acquisition. The values copied into the HIT structure are not actually the beginning and end of the sequence, but the beginning minus two time buckets, and the end plus two time buckets. This allows for some rise and fall time in the peak, and perhaps some minor examination of the background around the peak.

Once the zero suppressor is finished, all HIT entries are converted to SUBHIT entries. The intent was to allow for de-convolution of merged hits (i.e. de-convolution of multiple peaks in a single sequence). However, the test-run setup all but eliminated the chance that there would be more than one hit in a detector per event, so no software was developed to deconvolute any merged hits. In principle, one HIT might produce several SUBHITs; in reality there was a one-to-one correspondence.

The SUBHIT structure contained fitted parameters rather than raw data. Routines were developed to fit the ADC values with a gaussian, extracting an amplitude, centroid and width for each sequence. Also, the SUBHIT structure contained the value of the sum of all ADC counts in the peak, to be used for dE/dx studies.

In all of the cluster analysis performed by SCAT, gaussian peak shapes were assumed. When fitting gaussians to digitized data, there are several distinct situations which can arise.

If only one data point is available, no fitting can be performed. In this case, the centroid associated with this single point must be merely the abscissa of the point itself.

If two data points are available, it is possible to write down an equation for the centroid of the gaussian which fits those two points, but only if a width of that gaussian is known a priori. The equations which gives the centroid and the uncertainty in the centroid two data points are given below, with the derivation in appendix A.

If three data points are present, then the width can be calculated directly, along with an uncertainty. Also derived in appendix A the equations for the width, centroid, and uncertainties are given here.

Finally, if four or more data points are available, the most accurate position measurement is obtained by fitting a gaussian to the data points by a chi-square minimization process. The code developed for use in SCAT uses Levenberg- Marquardt minimization methods, based on sample code contained in Numerical Recipes In C [42].

The first step in the analysis of HIT entries into SUBHIT entries was to fit a gaussian to the time-series data representing a hit on a single anode. Since the zero-suppressor required a minimum of three time buckets above threshold, the time-domain data always contained at least three time bins. Therefore the two-bin and one-bin cases never arose.

The SUBHIT structures were what was used by the cluster finder. The cluster finder tried to arrange SUBHITs into clusters. In principle, each cluster would represent one particle crossing in the detector. The cluster finder (a relatively simple one) decided if a SUBHIT belonged to a cluster if the centroids of the cluster and the candidate SUBHIT were within a small tolerance of each other.

Once the SUBHIT structures were organized into clusters, the individual clusters were analyzed as units to obtain a two-dimensional position. The drift-centroids were used to compute an overall drift position for the cluster; each SUBHIT centroid was weighted by the sum of all bins in the SUBHIT, and the resulting weighted-mean drift centroid was computed.

In the transverse (anode) direction, it was not always true that three or more samples were present. Primarily, hits close to the anode deposit charge on fewer anodes, but also small-amplitude hits do not contain enough charge to be above threshold on any but the central anode. As a result, the two-anode clusters become an important part of the analysis process; if they are discarded, a significant portion of the total data set is lost. See section 3.3.2 for details of how the two-anode clusters were processed.

Finally, the clusters on each detector are converted to space points. SCAT used generalized rotation and translation matrices (provided as input by the user) to convert positions on the wafer faces into three-dimensional coordinates in the laboratory frame. Once the clusters were converted to lab coordinates, a tracking algorithm tried to group sets of three clusters into tracks.

Once tracks were constructed, a least-squares fit to a line was constructed, giving a three-dimensional straight line for the track. Subsequently, each point of the track was used to calculate a residue (distance between point and where the track line crossed the detector on which the point was located) in drift- and anode-directions.

3.3 Results from the analysis of test run data

3.3.1 Signal characteristics

The first thing done when analyzing the test-run data was to look at the signals from the detector with the *"Pad Monitor"* program. *Pad Monitor* shows a 2-dimensional overview of the data from a receiver card, coding the signal strength by color.

Pad Monitor also has the ability to show a single anode's data as a time series, with zoom capability. Figure 3.3 shows a pulse from a typical event. In this figure, the ADC pedestal has been shifted to allow the digitized signal to include negative analog values; the shifted pedestal is 20 ADC counts less than the normal pedestal. Hence, the zero-voltage level appears at 20 ADC counts in this plot.

Once average pulse characteristics were identified, it was possible to set the threshold and minimum length for the zero suppressor. The values selected were 3 ADC counts (30 mv) and 3 time buckets.

After setting the threshold and minimum length, SCAT was used to make a pass through all data tapes, writing only the zero-suppressed data to disk. This allowed for high-speed re-analysis of the data when bugs in SCAT were corrected or when new clustering or tracking algorithms were being tested. Since the SCSI tape drive used had a maximum transfer rate of 250 kB/sec, and events were 64 kB in size, the maximum processing rate from tape was 4 events

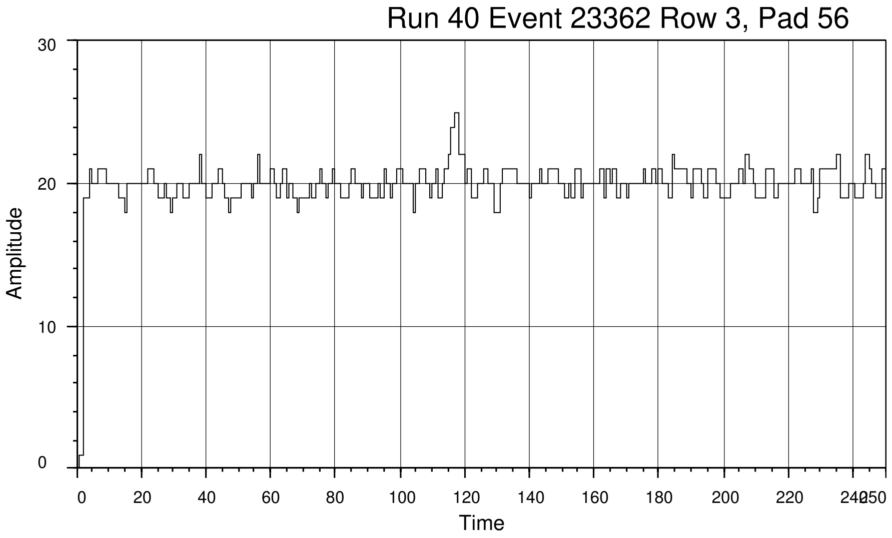

Figure 3.3: A single hit displayed by *Pad Monitor*

per second. When replaying zero-suppressed data from disk, the processing rate increased to 15 events per second.

During the analysis process, a program was developed to take the output from the hit-finding pass and create plots of the individual hits. This program, *"Plothits",* allowed for many command-line choices to control the number of plots created, allowing the user to specify cuts on peak height, drift distance, goodness-of-fit, etc. *Plothits* created PostScript files which were then printed.

Figure 3.4 shows typical output from *Plothits* (selecting a single cluster from a representative event). Clearly one can see the gaussian nature of the hit, both in the drift and anode directions. The reader should be aware that the scale is not a constant in fig. 3.4.

Once we were satisfied with the performance of the hit-finder and the

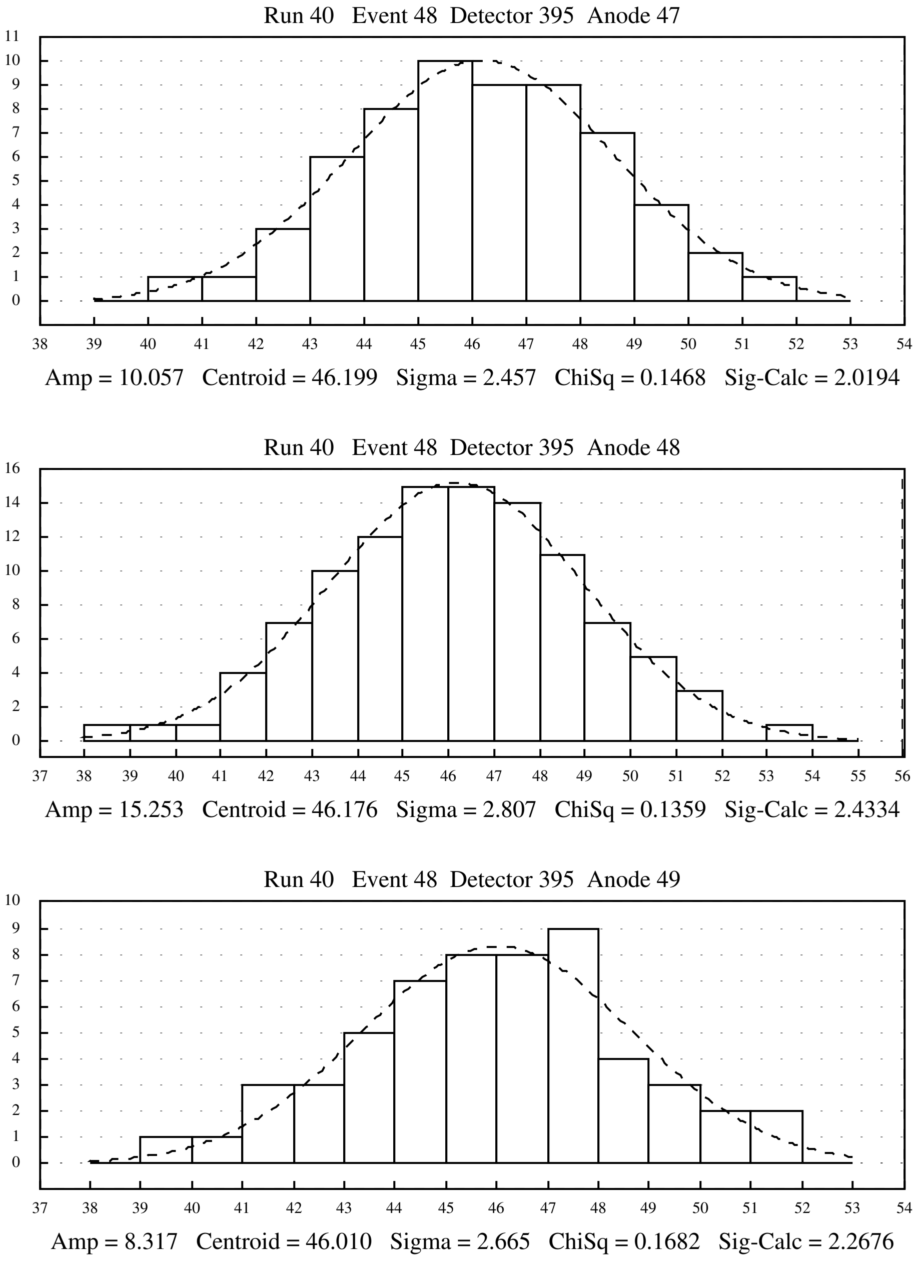

Figure 3.4: Sample output from the *Plothits* visualization program

fitting routines which were used to analyze individual hits, we continued with the analysis.

One of the first plots made during the analysis was an “occupancy plot;” the centroids of all clusters from a given run were plotted as a 2-dimensional histogram in order to assess the illumination of each detector and to check that all instrumented anodes were functioning at some minimum level. Figure 3.5 shows a representative plot, from the first detector in the beam line (SDD # 395). Both axes are in cm, after conversion from drift time and anode position. The drift time has been corrected for the SCA startup delay, as described earlier.

The first observation is the illumination appears not to be even. This is a somewhat erroneous conclusion, however, due to several features of the experiment. The first thing to consider is that only a few anodes were instrumented in each detector. Furthermore, a few anodes in each detector were faulty, and returned either no signal or spurious signals. Thus there are gaps in the active area of the detector which show up in the occupancy plot as regions with a lower hit density. The sharp edges on the illumination region are due to two effects. The lower edge (at drift distance ≈ 0.6 cm) is caused by the SCA dead time, as described earlier. The upper edge (at distance ≈ 2.2 cm) is caused by the fact that only 128 time bins were sampled. The few isolated hits beyond either limit are spurious points from erroneous results from the routines used to fit the charge distribution to obtain a centroid.

The solid lines present in the plot are from hits which contain only one anode. Since it is not possible to do anything with a single data point other than use its anode number to calculate a position, all single-anode clusters on a

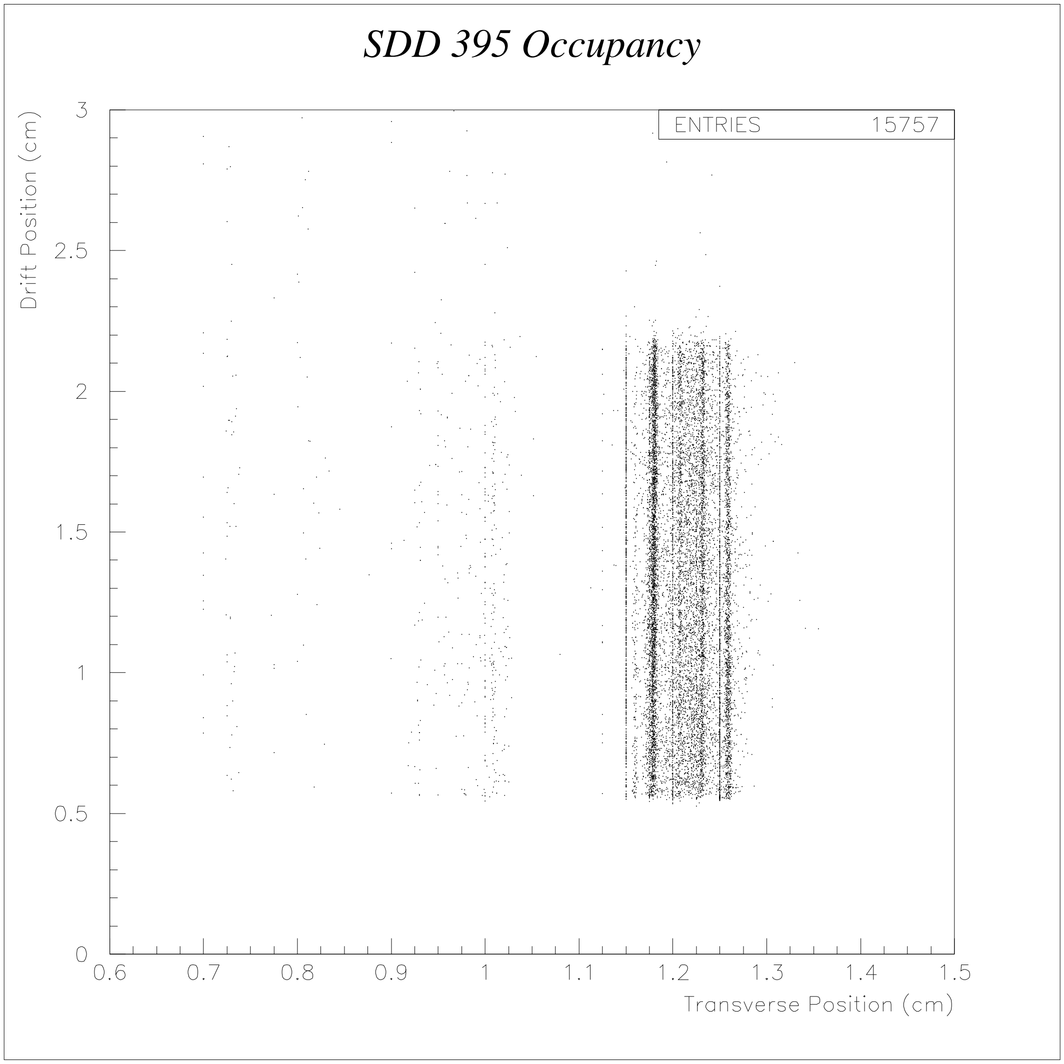

Figure 3.5: Occupancy of SDD 395 in the test beam

given anode will have the same anode-direction coordinate. Thus the straight lines appear.

Figures 3.6 and 3.7 show similar plots from the second and third SDDs in the beam (SDDs 405 and 416, respectively). SDD 405 behaves similarly to SDD 395, but SDD 416 does not. The reason for this is that SDD 416 was a last-minute substitution for a detector which was broken during shipment. Since no other "perfect" detectors were available at the time of the test run, the best of a number of defective detectors was selected. SDD 416 was a special case, however.

As an experiment, SDD 416 was irradiated at Ohio State University by a ^{60}Co source in an attempt to characterize radiation damage resistance for the SDDs [43]. The SDD was exposed to 9 krad of γ radiation, which is roughly equal to the dose expected during 10 years of STAR operation at RHIC. After the irradiation, SDD 416 was re-annealed in an effort to rehabilitate the unit.

However, the re-annealing was not an unqualified success, and the detector performed adequately on only a few anodes. This behavior can be observed in the occupancy plot in figure 3.7. Note that this plot clearly shows that only 4 anodes (anodes 48-51) are operational in this detector. This is not so critical for characterizing overall wafer performance numbers, but it makes tracking experiments (and hence position resolution measurements) somewhat more difficult.

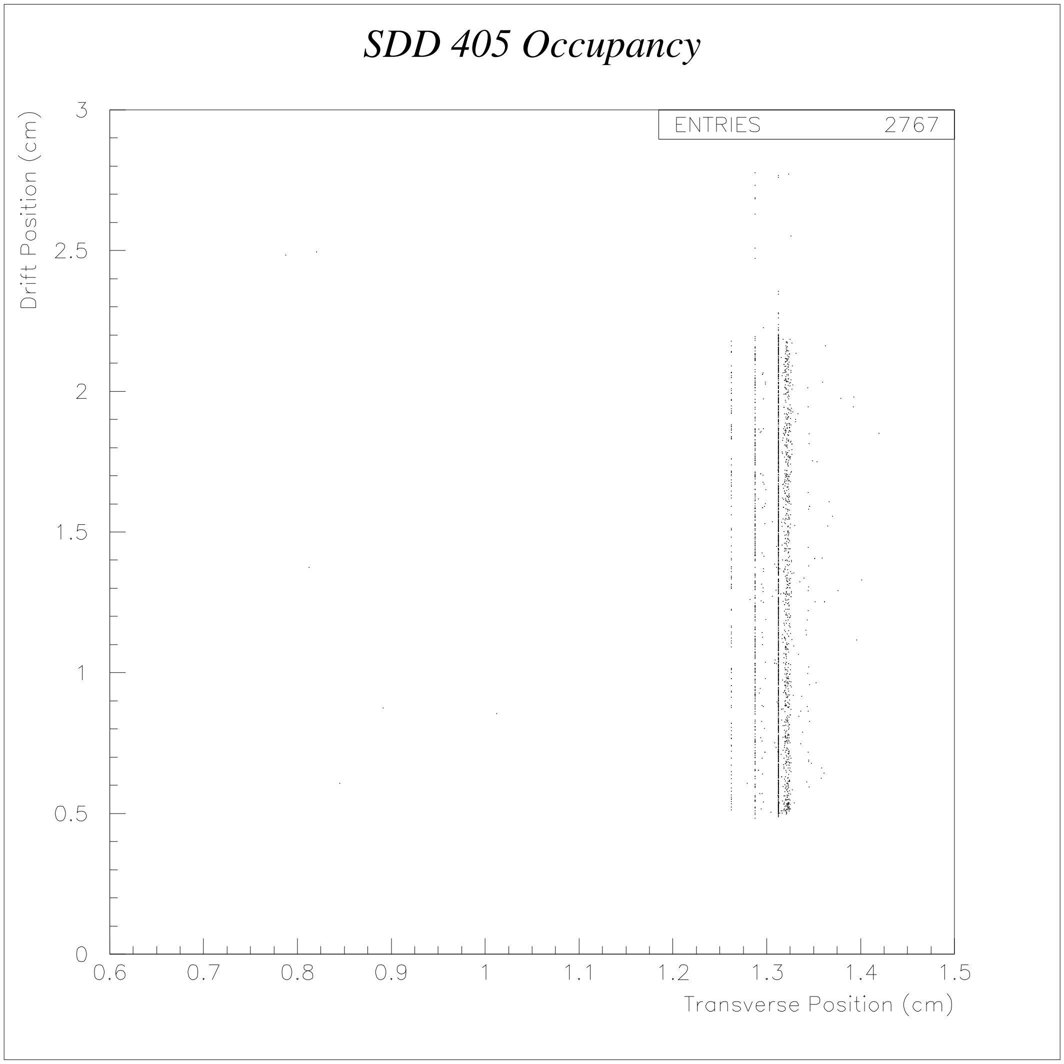

Figure 3.6: Occupancy of SDD 405 in the test beam

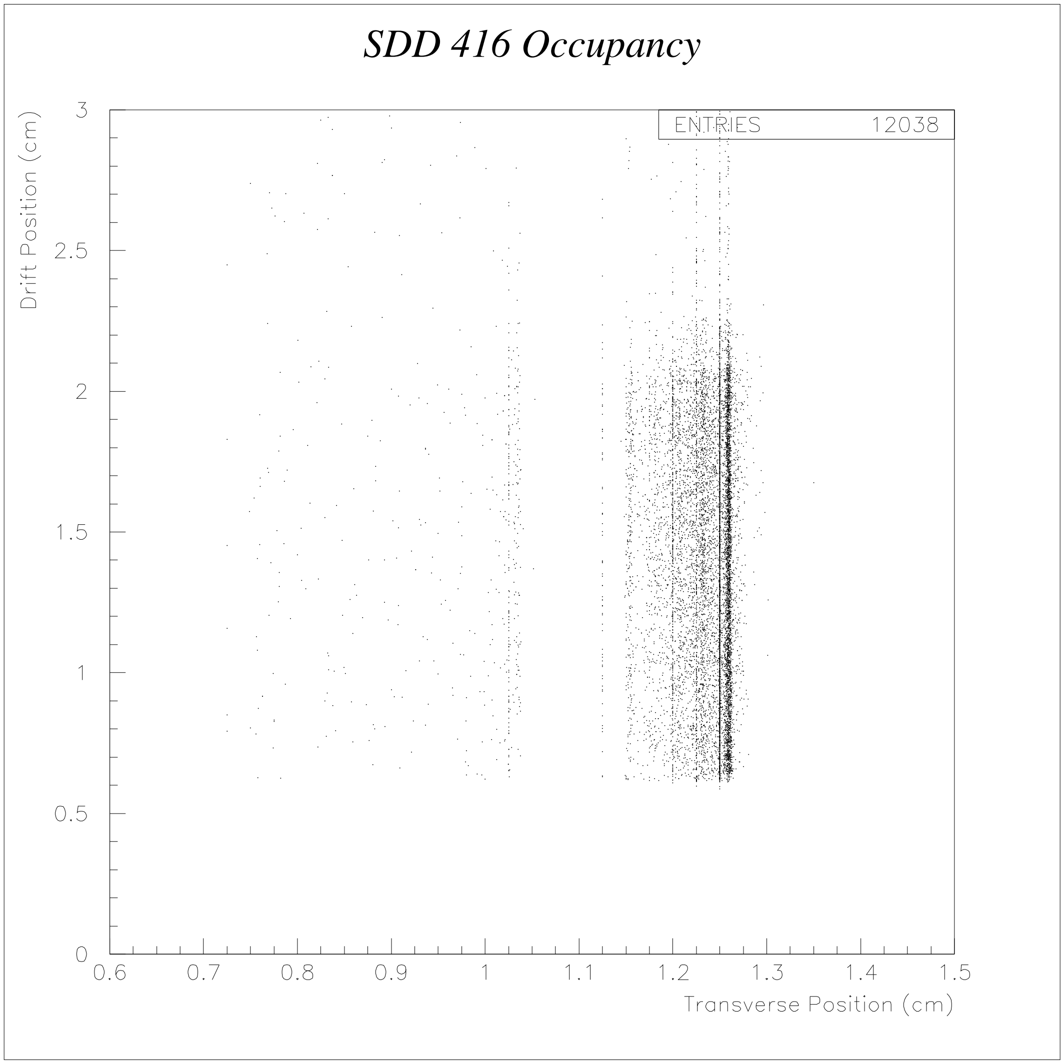

Figure 3.7: Occupancy of SDD 416 in the test beam

3.3.2 Determining the position of a hit from the cluster charge distribution

Clearly, one of the items of greatest interest in analyzing the test-run data is the ability to determine the position on the wafer surface of each cluster. In both the drift- and anode-directions, the process involves two steps. The first step is to determine the cluster centroid in units of time buckets (or anodes, as appropriate), and then to convert that position into physical units such as centimeters. For drift-direction centroids, the first step is easy, while the second step requires knowledge of the the drift velocity of the electron clouds, as well as the sampling rate of the SCA. These quantities will be discussed later.

The reason that the drift-direction centroids are easy to determine is that there were usually a large number of time bins available for fitting. The minimum length of any individual time sequence is guaranteed to be three bins (because the zero supressor requires three bins above threshold as a trigger condition), but often there were as many as 10 bins available for fitting.

In the anode direction, the process is slightly more troublesome; the conversion from anodes to centimeters is trivial, but the location of the anode-direction centroid can be a difficult process. The reason for this is that often clusters contained less than three anodes. Two- and even one-anode clusters were not uncommon in the data. Figure 3.8 shows a histogram of the number of anodes in clusters from a representative data acquisition run. Table 3.1 shows the same data, broken down into the three regions of interest for this discussion.

As described before, it is possible to fit a gaussian form to data and extract a centroid from the data if three or more data points are available. If

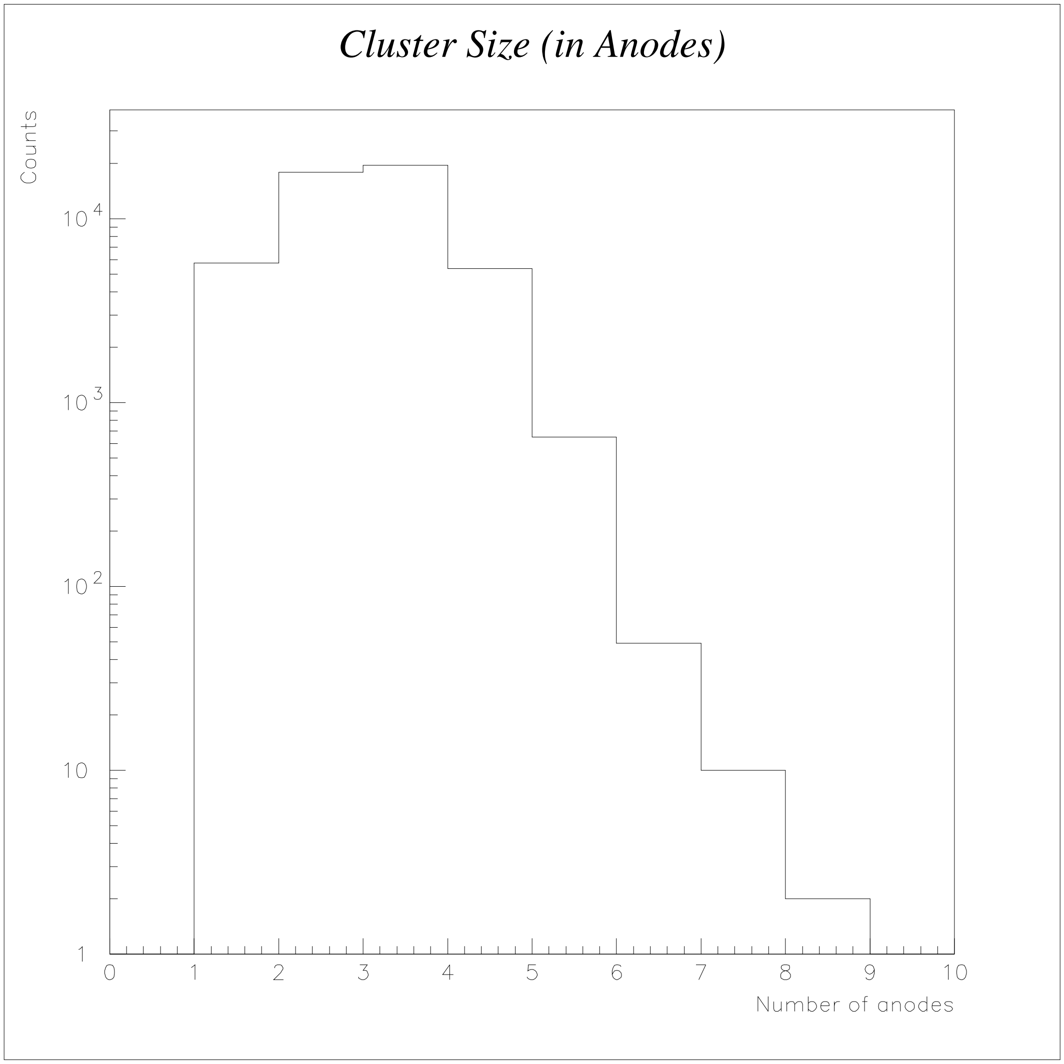

Figure 3.8: Histogram of number of anodes in a cluster

Anodes	Number	% of total
1	5751	11.7
2	17931	36.3
3+	25653	52.0
All	49335	100.0

Table 3.1: Table of cluster anode counts in SDD 395

two data points are available, then a fit can still be made, but only if the width of the peak is known in advance. Thus, the large number of two- and one-anode clusters present in the test-run data present a substantial problem. If we analyze only the clusters which contain three or more anodes, we lose appproximately 48% of our total data. However, analyzing the two-anode clusters, which would gain back 36%, requires knowledge of the cluster widths.

The first step is to plot the cluster widths in the drift and anode directions. This is of interest beyond the data analysis itself, because the cluster widths affect the total occupancy of the detectors in high multiplicity environments.

In theory, since the beam particles in the test run are normally incident on the SDDs, the resulting charge cloud should be isotropic. In other words, the cloud should be "circular," with identical widths in the drift and anode directions. Thus it should be possible to use drift-direction widths, calculated from fits of the time-sequence data, to determine the anode-direction widths so that two-anode clusters can be analyzed.

However, the drift-direction data include a gaussian width from the shaping amplifiers, which should add in quadrature with the width due to the shape of the charge cloud. Therefore, when determining the physical size of the cluster in the drift direction, one must subtract (in quadrature) the known width due to the shaping amplifiers. This width is 35 ns, or 1.4 time buckets.

Figure 3.9 is a plot showing the average cluster width as a function of drift distance, binned in 0.25 cm intervals. The filled triangles represent the drift-direction widths, de-convolved as described above, and converted to a physical size in μm. The open squares represent the anode-direction widths, converted to μm. The relationship between the two is clearly strong enough to use during data analysis.

So, armed with the knowledge that the cluster widths could be extracted from the drift-direction data, the clusters were re-analyzed. Figure 3.5 (previously presented) actually includes the two-anode and one-anode clusters.

Finally, an attempt was made to increase the average number of anodes per cluster. To do this, code was added to SCAT which looked on the anodes immediately adjacent to all clusters, calling the hit-finder again to search the proper drift regions using a slighly lower threshold. This re-processing of the data yielded a minor increase in the number of three-anode clusters.

3.3.3 Detector efficiency

Of obvious interest is the particle detection efficiency of the SDDs. To evaluate the efficiency of the center detector, we counted the number of events which contained hits in all three detectors, and divided this by the number of events which contained hits in only the first and second detectors. In order to try to avoid problems with small signals, a cut was placed so that we counted only events in which the peak ADC value for each clusters was above some threshold.

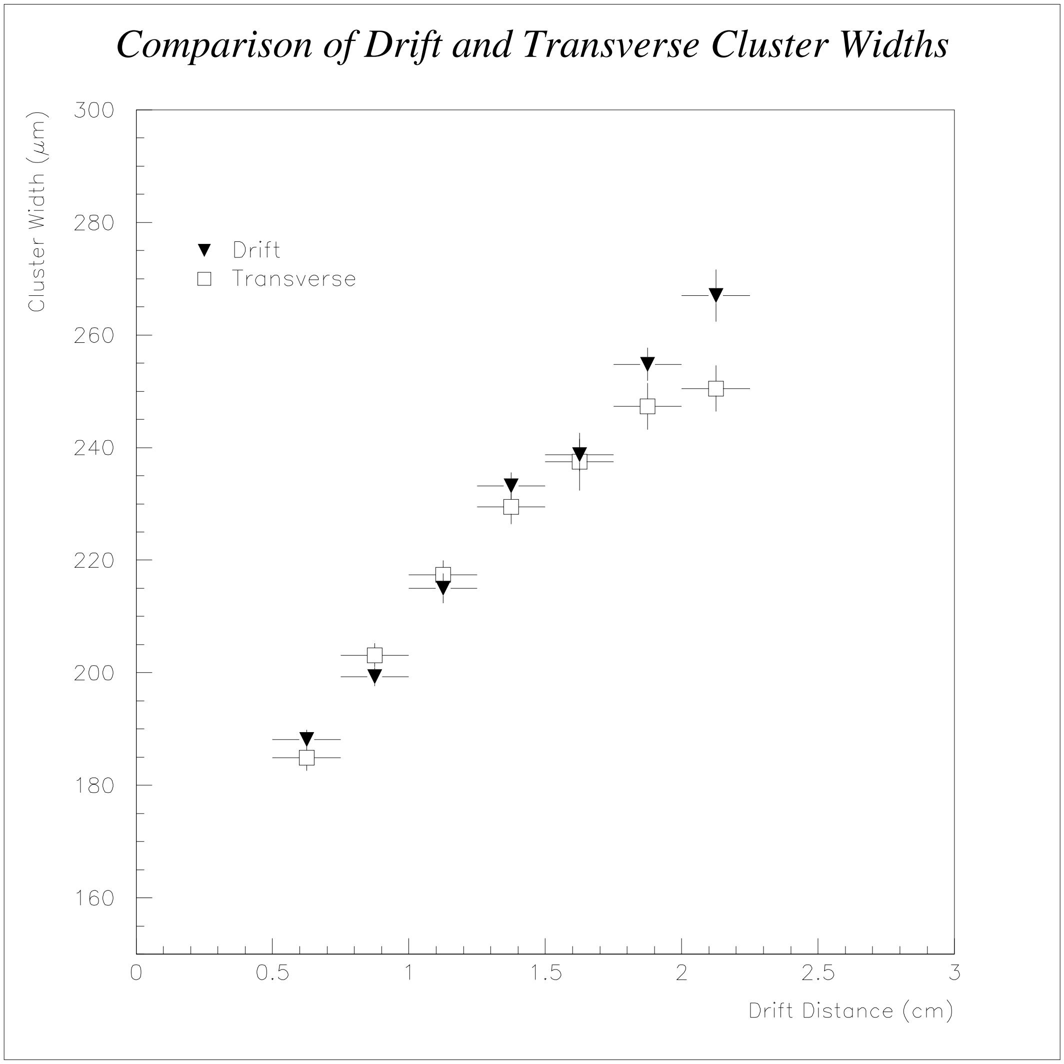

Figure 3.9: Drift and transverse cluster widths vs. drift distance

Further, an attempt to avoid the problems associated with the uninstrumented anodes was made, by placing a cut on the anode-direction centroid of the clusters.

When all cuts were made, there were 106 events with hits in SDDs 395 and 405, and 94 events with hits in all three detectors. Thus the efficiency of SDD 416 was measured. Clearly this is not 100%!. The 12 events in which a signal was not found in 416 were inspected by hand, using *Pad Monitor.* In all 12 events, a signal was discernable by eye, exactly in the expected location, based on the signal locations in the other two detectors. In each case, the signal was small enough to be below threshold for the hit-finder, with only one or two buckets present above threshold. Thus we conclude that the detection efficiency of SDD 416 is essentially 100%, but that fluctuations in the energy deposited were an important effect.

Similar measurements were made for SDDs 395 and 405, with similar results.

If the peak height cuts described above are removed, there are 5737 events with hits on all three detectors, and 6738 events with hits on just 395 and 405. This implies an efficiency for SDD 416 of 85 %, which is close to the 89 % efficiency measured with the cuts (without searching the troublesome events by hand).

In the final STAR SVT, the pre-amplifier/shaping amplifier units will have twice the gain, while the ADC units will have 20 times the sensitivity. Accordingly, the average ADC value of a MIP peak should be approximately 40 times what it was in the test run, and it seems unlikely that energy deposition fluctuations will have a significant impact on the efficiency of the SDDs.

3.3.4 Charge reconstruction

Clearly, the distribution of total charge in the reconstructed clusters is of interest. Figure 3.10 shows a two-dimensional scatter plot of the total charge in the clusters (uncalibrated) versus the drift distance. In theory, charge in the clusters should be conserved, except perhaps for small recombination effects during the drift time. If the drift time is kept small, this effect should be negligible. From examining the plot, however, it is clear that there is a gentle trend towards "losing" some charge at greater drift distances.

It is likely that this charge has not recombined during the drift, but that the charge has spread to adjacent anodes whose peak values fell below threshold for the hit-finder. Figure 3.11 shows a scatter plot of the average peak height vs. the drift distance. This figure is expected, since as the total charge is shared among anodes, the amount present on the central anode must decrease.

Knowing this, if one wants to make a sensible histogram of the charge present in the reconstructed clusters, the histogram must be made with cuts on both the drift distance and the number of anodes present in the cluster. Figure 3.12 shows just such a histogram. Clearly this histogram exhibits the characteristic "Landau curve" shape for deposited energy, similar to curves for other ionization-based detectors such as scintillators and wire chambers.

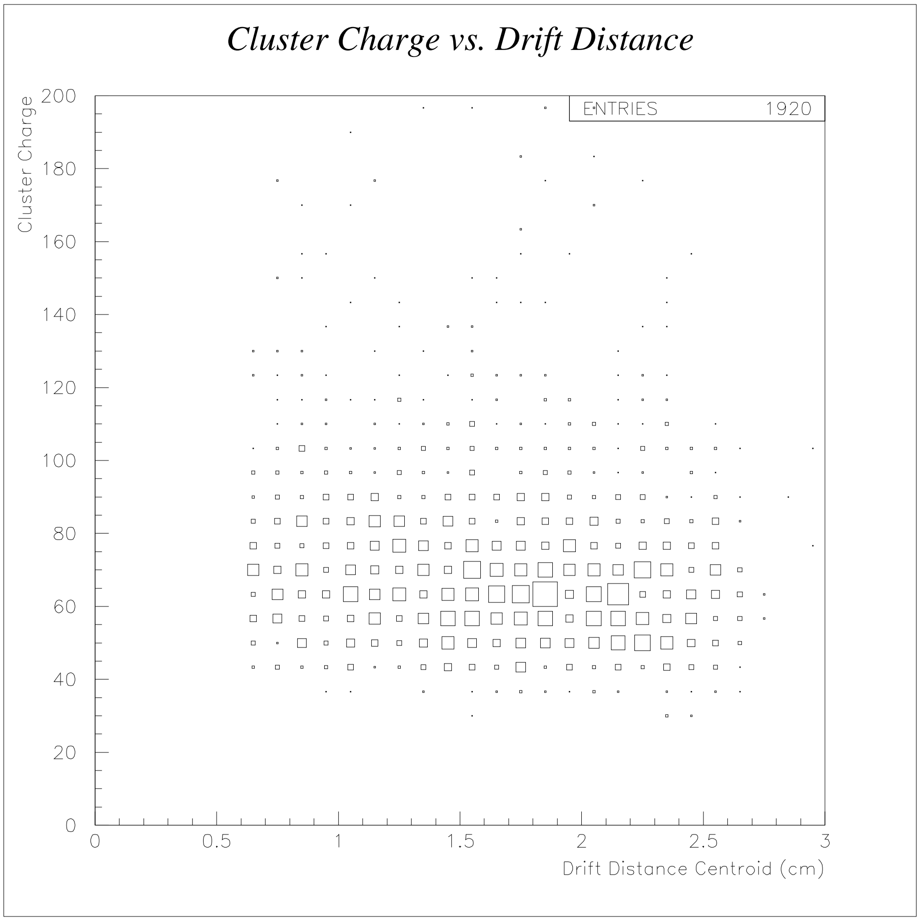

Figure 3.10: Cluster charge vs. drift distance

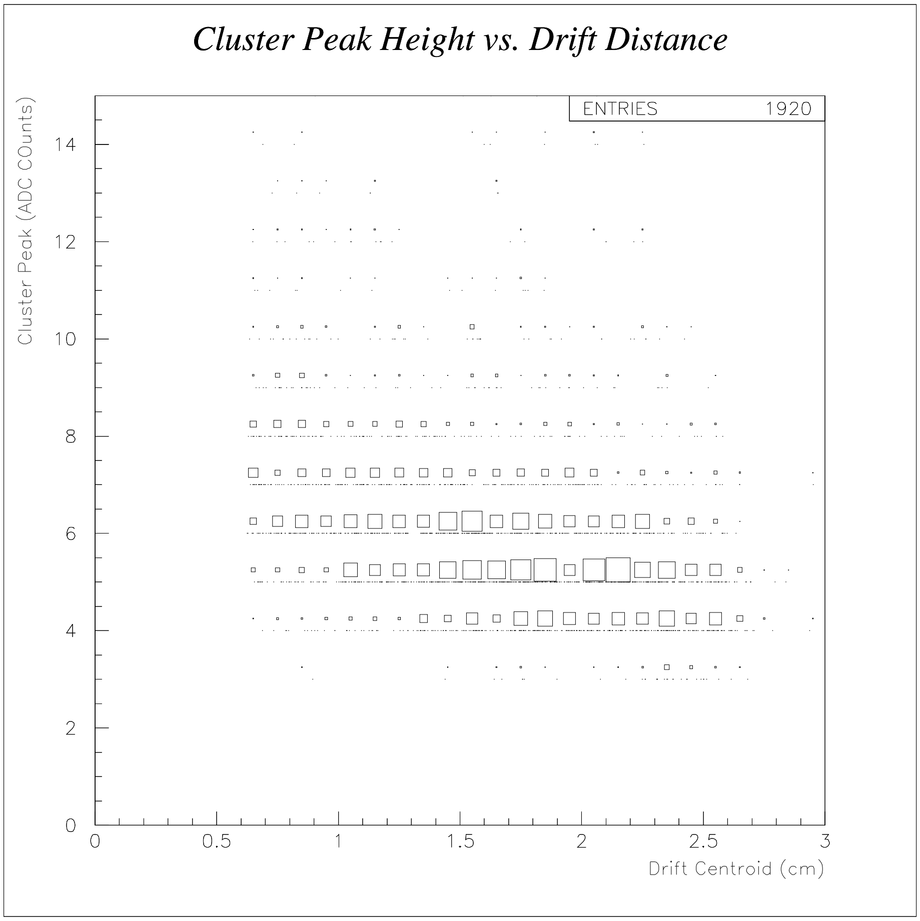

Figure 3.11: Peak height vs. drift distance

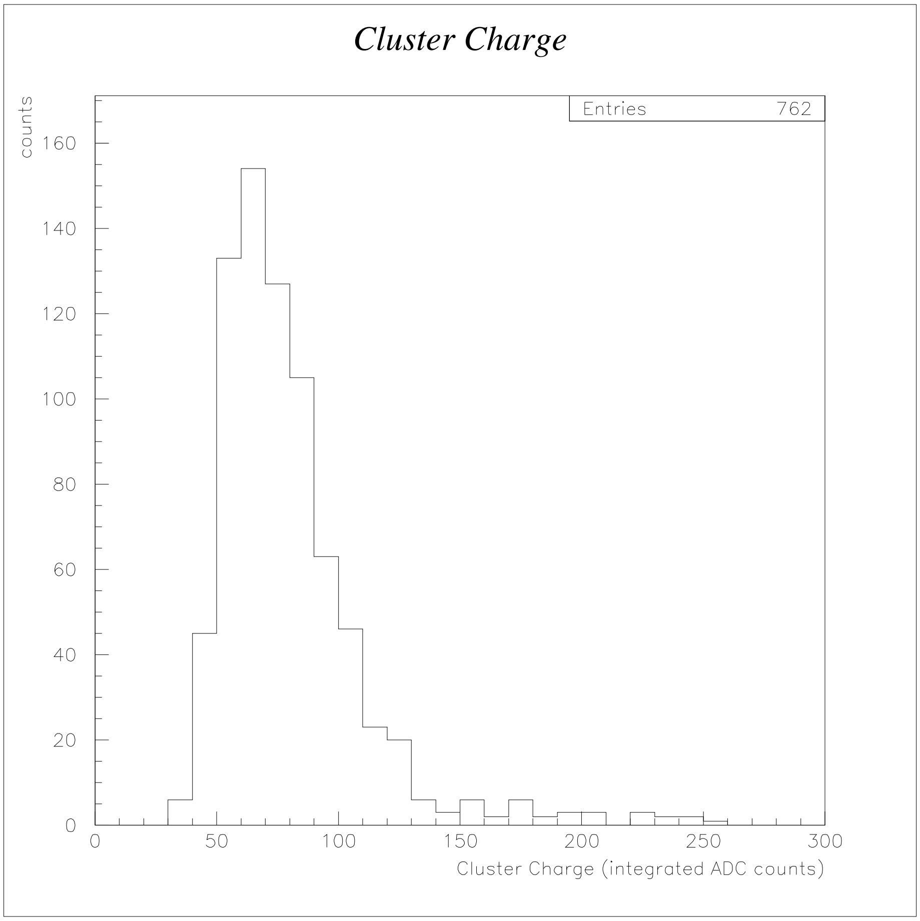

Figure 3.12: Total cluster charge

3.3.5 Position resolution

One of the principal goals of the test run was to determine the position resolution of the detectors for minimum ionizing particles. In laboratory bench tests of the detectors, position resolution can be measured easily since the detectors are triggered using an infrared laser striking at a known position. In the beam test, however, the exact position at which a particle strikes the detector is not known.

In order to measure position resolution, one needs three detectors. When hits are present in all three detectors, a track can be formed, and a straight line can be fit to the positions of the hits. Once a fitted line is available, the distance between each point and the line is easily calculated, and these residues can be used to calculate the average error in the calculated position.

Figure 3.13 and 3.14 show histograms of the residues from the accepted three- detector events from a single run. Cuts were placed in order to require the tracks to pass through the fully-instrumented regions of the detectors.

The number of interest comes not from the centroids of the peaks, but from the widths. The centroids should be roughly zero (which they are); the widths give a measure of the average excursion away from zero.

Note in figures 3.13 and 3.14 that the center detector's histogram shows twice the width of the other two detectors'. This is an artifact of the fact that three points were used to create (and thus fit) the straight line from which the residue was calculated. In such a fit, the best-fit line will always lie closer to the outer two points that the inner point, by a factor of two. This factor of two appears in the histogram of the residues from the center detector.

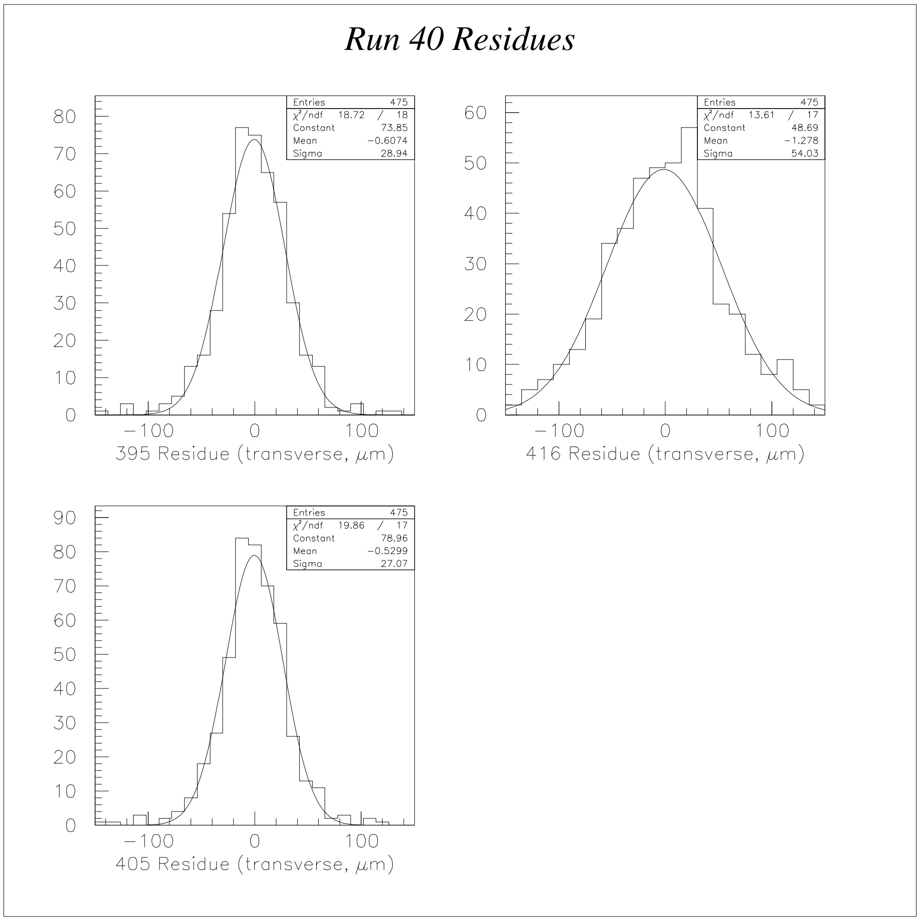

Figure 3.13: Transverse residues from good tracking events

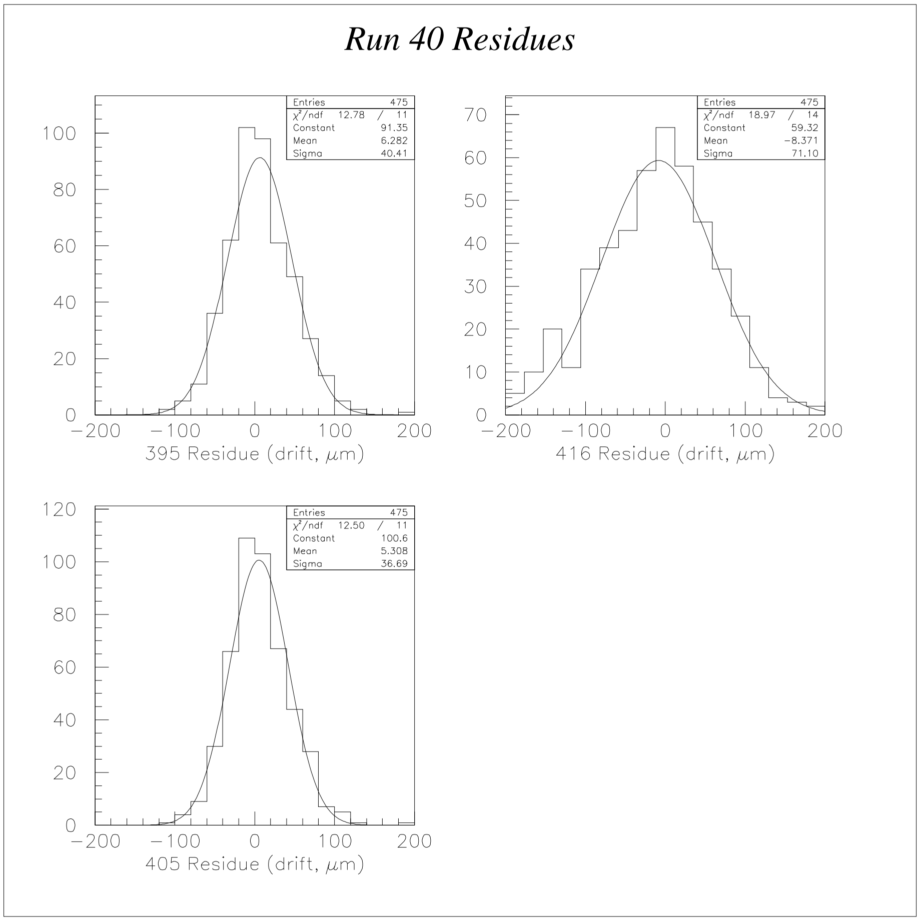

Figure 3.14: Drift residues from good tracking events

The width of the position error distribution can be calculated from the width of the residues distribution, assuming that all detectors have identical performance. Since the widths of the residue histograms differ by less than 10 %, this seems like a reasonable assumption to make. The relationship between the residues and the position resolution is given by [44]

$$\sigma_{position} = \sqrt{6}\ \sigma_{residue} \tag{3.1}$$

Thus from figures 3.13 and 3.14 we calculate that the position resolution for these detectors is 93 μm in the drift direction, and 68 μm in the anode direction. Note that this method does not give a way to measure the resolution of each detector individually, but gives an average resolution for the three detectors as an ensemble. Further, this method is blind to systematic errors in position calculation; if the same method is used to calulate the position in each detector, then all three detectors will show an offset in the same direction, and the fitted track will shift by the same amount. thus leaving the residue constant.

Finally, it should be mentioned that the position resolution should be dependent on the drift distance; there are several mechanisms which could affect the resolution. Foremost is the tendency to have fewer anodes in the clusters towards small drift distances, due to the small drift time (the charge cloud has not had enough time to spread to adjacent anodes). The test-run data do not show this effect, due to the 1.2 μs dead time from the SCA start-up, but it is an effect which may impact resolution in STAR. For large drift times, the decreased peak height can play a role in degrading the position resolution.

Chapter 4

Detector simulations and data analysis code

Detector simulations have two purposes. The first is to simulate detector signals in order to have input for testing data analysis code. The second, once the first has been done, is to evaulate the performance of the detector as a whole regarding its ability to measure physics observables of interest.

4.1 Simulating the silicon drift detector

The purpose of the SVT Signal Simulator (SSS) is to simulate the output of individual SDDs in response to particle crossings. The SSS package was originally written by V. Rykov, working at WSU. Later it was ported to the STAR "TAS" software enviroment by D. Read and C. Pruneau.

To understand SSS, it is useful to make comparisons to the two previous SDD simulation codes. The "fast" simulator, SRS (SVT Resolution Simulator) is merely a "position smearing" algorithm; it takes detector crossings as an input table, and copies them to the output table, while changing the coordinates slightly. The new coordinates are "smeared" from the old ones with a gaussian distribution, using a width which is estimated to be typical of SDD performance. This package was developed by K. Wilson.

The original “slow simulator,” also written by K. Wilson, is an electron-transport code. It simulates the actual physics of SDD operation, drifting the electron cloud in the applied electric field. It accounts for all of the major processes in the transport: Coulomb repulsion, diffusion, ion-pair recombination, field non-linearities, etc. This code can be used to fill an output table of ADC values, which is of interest because it allows development and testing of zero-suppression, cluster-finding, and space point extraction routines.

The advantage of SRS is that it is very fast; it takes at most a few milliseconds per point to complete its task. The disadvantage of SRS is that it is purely phenomenological, and has no way to allow for ”merged” hits (hits on the detector whose charge clouds overlap when they reach the anode). Also, it does not provide a mechanism for introducing noise.

The electron-transport code allows any imaginable situation to be simulated; it provides great flexibility by simulating the underlying physics processes involved in SDD operation. However, it is too slow to be used in STAR detector simulations; on STAR’s fastest computer, one hit takes approximately 10 minutes to simulate. Given that a typical 100 GeV-A Au-Au event in STAR may produce 7000 hits in the SVT, the electron-transport code is clearly not usable.

Thus we see the need for the SSS package. In retrospect, SSS should be considered a “medium-speed” simulator; at roughly 10 ms per hit, it is faster than the electron-transport code but slower than SRS. It allows for more flexible simulations than SRS, but might be blind to some processes that the electron-transport code could simulate.

SSS operates by integrating two coupled differential equations which represent the time-evolution of the size of the electron cloud. These equations were derived by V. Rykov and tuned to match previously-measured detector performance. SSS consists of three main routines. The first section integrates the differential equations to calculate a cloud size depending on the drift distance and total cloud charge (and SDD parameters such as electron mobility, drift voltage, etc.). This routine is called once per particle crossing. The second routine determines how much charge should be present on any given anode, and calculates the drift- direction centroid of the charge distribution on that anode. This routine is called once for each anode in the final distribution. The last routine in the SSS package calculates the output signal (in mV) for a single time bucket on one anode. It includes effects which simulate the performance of the SVT PASA units, including rise-time and undershoot effects.

The final output of the SSS package is a pixel array representing the raw output (in mV) of a single SDD. This array must be processed further by other routines or packages.

4.2 Analyzing silicon drift detector data

The analysis of SDD data can be broken down into five pieces. The organization of the SVT data-analysis code is intended to mimic the operation of the STAR data acquisition and analysis hardware. The five stages in this are as follows:

- Analog data from the PASAs is sampled and stored in a Switched Capacitor Array (SCA).

- The SCA is read out into an A/D converter.

- The digital data is then sent to a set of Application-Specific Integrated Circuits (ASICs) which perform a conversion from the 10-bit values returned from the ADC to 8-bit values used in the rest of the analysis. The ASICs also perform subtraction of the ADC pedestals. The ASICs also act as digital discriminators, passing on pointers to sequences of data which are above some user-defined threshold.
- Local processors (Intel i960 RISC microprocessors) organize these sequences into "clusters" which are intended to represent discrete regions of charge on the wafer face.
- The clusters are analyzed to convert them to space points, which can later be used in tracking.

Figure 4.1 shows a schematic view of the logic flow in the `svt_medium` routine, which is the master routine in the SVT simulation and analysis chain.

4.2.1 The SAD package

The SVT Analog to Digital (SAD) package mimics the operation of the A/D converters in the SVT DAQ system. Currently, it performs a linear conversion, with no noise effects. Future versions will include A/D non-linearities and random noise effects. Also, the SAD package performs a 10-bit to 8-bit conversion, something which is included in the SVT DAQ specifications.

SAD copies data values from the input table to the output table, converting mV to digital ADC counts for each pixel.

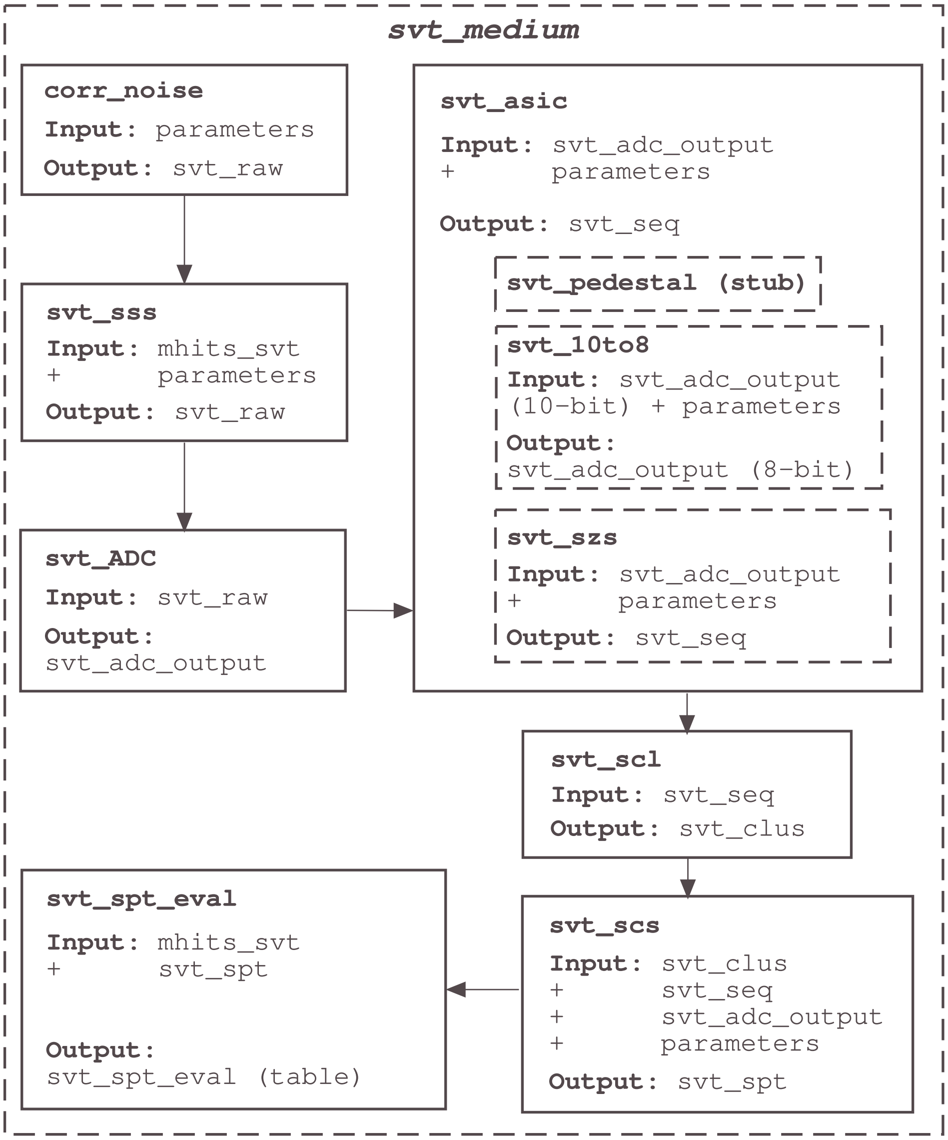

Figure 4.1: Logic flow in the `svt_medium` routine

4.2.2 The SAS package

The SVT ASIC Simulator (SAS) pacakges simulates the operation of the Application-Specific Integrated Circuits in the SVT DAQ. The first task performed by these chips is pedestal subtraction. Since pedestals do not currently exist in the simulated data, this function is merely a stub (i.e. a placeholder).

The next task performed by the ASICs is to convert the 10-bit digital ADC value to an 8-bit value using a lookup table. The lookup table is generated from a "tri-linear" function. The table is linear with a slope of 1.0 from input ADC value 0 to 127. From value 127 to 383 it is linear but with a slope of 0.25. Finally, in the region from 383 to 1023 the curve is linear but with a slope of 0.10.

The "standard" functional form for such conversion tables (used in previous heavy-ion experiments) is a logarithmic function. The tri-linear function described here is similar in shape. Either form may be used without any speed penalty, since the table is filled only once, and once filled, all lookups are completed with identical speed.

The effects of one alternate choice for the functional form of the 10-bit to 8-bit conversion table has been simulated. See section 4.3.5.

Finally, the ASICs perform zero suppression. The intent is to reduce the stream of raw data to include only regions of interest, removing the areas where the signal is zero. The input to the SVT Zero Suppression (SZS) routine is the output table from earlier in the SAS package, i.e. a table of 8-bit output values representing the pixel array from a single SDD. The output from SZS is a table of pointers back into the input array. These pointers point to the starting and ending positions of sequences of ADC values which are considered to be of interest. In essence, SZS is a digital discriminator.

The discrimination algorithm used is a simple "two-threshold" algorithm. I will call the first threshold the "acquisition threshold" and the second threshold the "trigger threshold." The acquisition threshold must be lower than the trigger threshold. SZS ignores all time bins which are below the acquisition threshold. The requirement for triggering is that at least `n_lo` consecutive bins must be above the acquisition threshold, and at least `n_hi` bins (not necessarily consecutive) must be above the trigger threshold.

Once SZS has been "triggered" by a sequence, it adds an entry to its output table. This entry contains the number of the anode on which the sequence was found, the start and end times (in time bins) of the sequence, and an index into the input pixel array which is the index of the first time bin in the sequence.

SZS searches the entire pixel array for sequences, as described above. It performs no analysis whatsoever of the sequences; it merely fills the output table with the information described above.

4.2.3 The SCL package

The SVT Cluster Locator (SCL) package tries to assemble into "clusters" the sequences which are the output of SZS. In the simplest case, a cluster contains only the sequences which belong to a single detector hit. In practice, however, clusters can contain sequences which belong to several hits. This will occur when the detector hits are close enough together that the charge clouds have merged before they reached the collection anodes. Such "merged hits" occur with relatively low frequency due to the low occupancy of the SDDs, yet they constitute a problem which must be handled.

The input to SCL consists of the table of sequences produced by SZS. Here is the algorithm used by SCL to find clusters in these sequences:

1. Mark all sequences as “unused.”
2. Loop over the list of sequences. Find an unused sequence; this becomes both the “first” sequence and the “last” sequence in the cluster. Clear the “cluster changed” flag.
3. Search the entire table of sequences from the beginning of the table.
4. If a sequence has been marked as “used,” continue the loop at step 3.
5. Is this sequence on the anode directly to the “left” of the first sequence in the cluster? If so, does it overlap the first sequence? If yes, this sequence belongs to the cluster; it becomes the “first” sequence in the cluster. Mark the sequence as “used,” set the “cluster changed” flag and loop back to step 3.
6. Is this sequence on the anode directly to to the “right” of the last sequence in the cluster? If so, does it overlap the last sequence? If yes, this sequence belongs to the cluster; it becomes the “last” sequence in the cluster. Mark the sequence as “used,” Set the “cluster changed” flag and loop back to position 3.
7. If the sequence fails both tests (5) and (6) above, leave it untouched and loop back to position 3. The “cluster changed” flag is left unchanged.

8. When the entire sequence list has been searched, if the "cluster changed" flag has been set, reset the flag and start again at position 3.

9. When all sequences have been used, SCL is done.

The "overlap" condition mentioned in steps (5) and (6) is a simple one. A sequence is defined to overlap an adjacent one if the start and stop times are consistent. Let b1 and e1 be the begin and end times of the first sequence, and let b2 and e2 be the begin and end times of the second sequence. Then these sequences overlap if $b1 \leq b2$ and $b2 \leq e1$ or if $b1 \geq b2$ and $b1 \geq e2$. Figure 4.2 demonstrates these two cases. Each rectangle in the figure is meant to represent a single sequence. This method for determining if two sequences overlap should handle most cases. Particles with very large crossing angles might produce completely disjoint clusters, but charge-spreading in the electron cloud should minimize any such effect.

SCL, in addition to filling its output table, makes some changes to the entries in the sequence table. Each entry in the sequence table contains one field which is not filled by SZS. This field is the "next sequence" field. Essentially, SCL fills these fields to construct a singly-linked list of sequences; each "cluster" points to the first sequence in the list, and any code wishing to find all sequences in the cluster must follow the linked-list until it ends. A "next sequence" value of 0 indicates the end of the cluster. When SCL starts, it sets the "next sequence" field of all sequences in the table to be -1, which indicates that the sequence is unused. If any value 0 or greater appears in this field, the sequence has been used in a cluster.

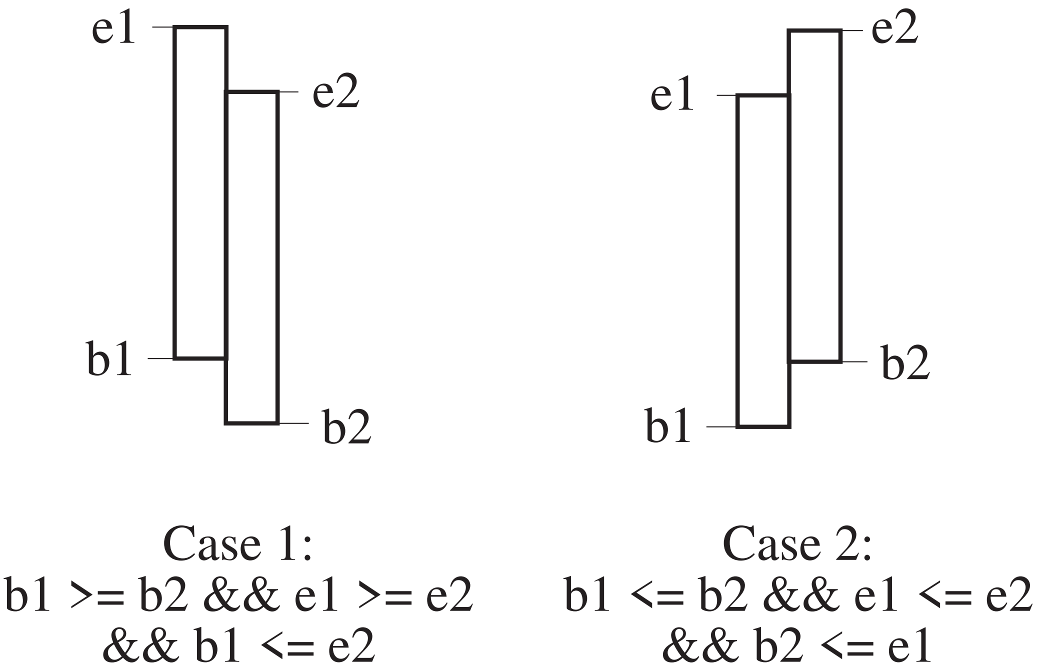

Figure 4.2: Two conditions for adjacent sequences to be considered to be overlapping

Note the use of the "cluster changed" flag in the algorithm above. This allows clusters to "grow" in either direction. In principle, the sequence list will be completely unordered, since the sequences will be found in parallel by several ASICs. This implies that the "seed" value found by SCL might not be the "first" sequence in the cluster. This, in turn, implies that SCL must search for new sequences which belong on "either side" of the cluster as it "grows." Any time a new sequence is added to the cluster, the possibility arises that another sequence is now adjacent to the cluster edge. Thus, the table of sequences must be searched repeatedly until the cluster has stopped growing. Once it has been verified that no unused sequence in the table belongs to the cluster, the cluster may be considered “finished,” and a search for new clusters may begin.

The algorithm above also guarantees that all sequences will be used as part of a cluster, even if some clusters will contain only one sequence.

One potentially troublesome situation which can arise is the “merged hit” case. This situation can manifest itself in several ways. One way in which it can appear is when the hits have been merged in the “drift” direction; the hits were separated along the drift axis rather than the anode axis. In this case, one or more of the sequences in the cluster will contain two peaks. However, no analysis has been performed upon the sequences at this point, so SCL has no way to know of this situation.

Alternatively, hits might be separated in the anode direction. In this case, there must exist at least one sequence which should belong to both clusters.

Finally, one can imagine many different geometries which can be assumed by multiple hits. Diagrams of a few possible shapes clusters can take for merged hits are below. The pixels from cluster A are represented by right-slanting hatches, while the pixels from cluster B are represented by left-slanting hatches, and the pixels which belong to both clusters are represented by cross-hatches.

The challenge for SCL is to construct clusters which contain <u>all</u> sequences which belong to either sub-cluster in the merged-hit case. This is desirable because a later routine (SCS) will extract the space point information from the clusters. If possible, it would be nice to signal SCS that a cluster

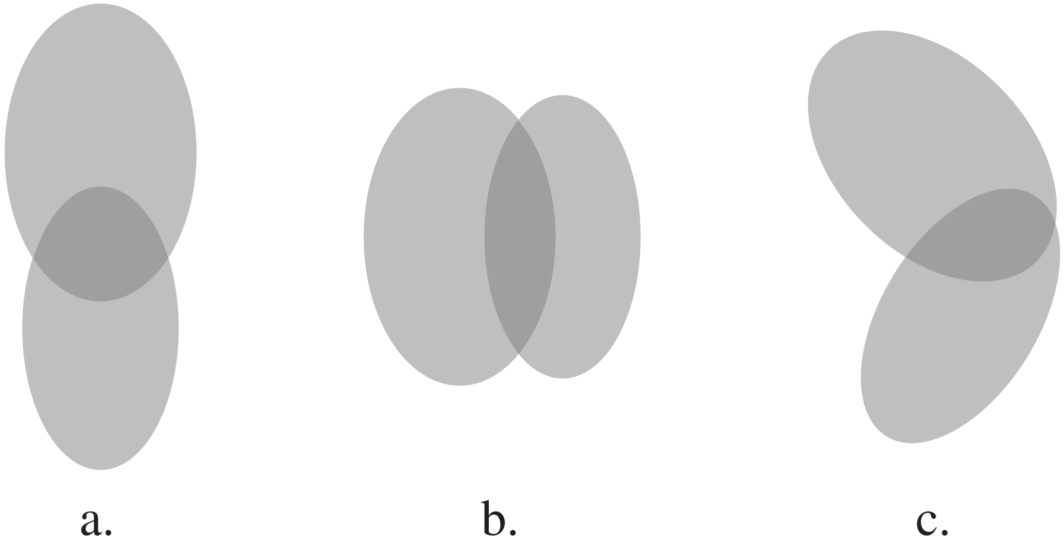

Figure 4.3: Three examples of possiblities for merged-hit cluster shapes

contains more than one hit.

In order to handle multi-hit clusters, a small modification must be made to the cluster-finding algorithm described above. Steps (5) and (6) must be replaced with a combined step. This step should take the candidate sequence and search for an overlap condition on either side of any sequence already in the cluster, not just on the "left of the first" sequence or on the "right of the last" sequence. In the case when an overlap is found with an interior sequence, the cluster should be marked as containing multiple hits so that SCS may know right away that special analysis must be made of that cluster. This will allow single clusters to manage shapes such as the "kidney bean" of case (c) above. In fact, it should allow clusters to manage any shape imaginable.

4.2.4 The SCS package

The SVT Cluster to Spacepoint (SCS) package converts clusters constructed by SCL into three-dimensional space points in the STAR coordinate system.

Clearly, the easiest case for SCS is when a cluster represents a single hit on a wafer. As mentioned previously, it is possible for two or more hits to become merged into one cluster, and it is the job of SCL to make sure that if any part of any hit becomes merged with another hit, that all sequences in both hits will be contained in a single cluster. Since the occupancy of even the inner layer of the SVT is fairly low, we expect that two-hit clusters will be relatively rare, and cluster containing more than two clusters will be virtually non-existant.

First I will describe the overall architecture of SCS.

The input to SCS takes the form of three tables: `svt_clus`, which is a list of clusters found by SCL (i.e. pointers to the heads of linked-lists in the `svt_seq table`), `svt_seq`, which contains pointers to “runs” or 'sequences' in `svt_adc_output`, and `svt_adc_output`, which should be considered to be the “raw” data. As its output, SCS fills the `svt_spt` table, which contains merely space points in global STAR coordinates, including the position of each space point, estimates for the errors on those positions, and the total energy deposited in the cluster from which the space point was extracted.

4.2.4.1 Processing isolated hits

For the simplest case, that of the cluster which represents only one hit on the wafer, the analysis is straight-forward. The first moment of the charge distribution (in both the drift and transverse directions) is used to calculate the position of the cluster on the wafer face. The zeroth moment of the charge distribution (i.e. the sum of the ADC values from all pixels in the cluster) is used as a measure of the total charge in the cluster.

4.2.4.2 Processing merged hits

Analyzing clusters which contain more than one hit is more difficult than analyzing clusters which contain only one hit. The first step in the process is to decide whether a given cluster contains one hit, or more than one hit.

In principle, one could just count the number of peaks in every cluster and decide what to do based on the number of peaks. However, peak-counting is a slow process (and one which can be fooled by noise), so it is useful to have some faster discrimination method. The one developed for use in the SVT software relies on the sizes of the clusters.

At the time that the centroids of a cluster are calculated (i.e. first moments), the second moments are also calculated. The second moments are related to the widths of the cluster. If a histogram is made of the second moment in the drift direction vs. the second moment in transverse direction, a clear trend appears. Figure 4.4 shows this histogram. The small points are the isolated hits, while the larger filled circles are clusters which contain two or more hits. The sharp vertical lines with transverse 2nd moments of 0.25 and 0.5 are from one- and two-anode clusters.

The grouping of the isolated hits shows just what one would expect, that the isolated hits have second moments which grow in relation to each other. This is true because the charge cloud spreads essentially isotropically as it drifts. By placing a cut on the second moments, it is possible to discriminate isolated hits from merged hits; the isolated hits all lie close to the dense diagonal group in the histogram. Also, the cut includes all clusters with both second moments below a cut-off value. However, the cut excludes clusters whose second moments are on the line described above if the second moments are larger than a maximum value of 1.6.

If a cluster is determined to have passed the cut for being an isolated hit, it is processed in the manner described in section 4.2.4.1. If the cluster fails this cut, it is processed in a different way.

The first step in processing a cluster which is suspected to contain more than one hit is to find out exactly how many hits are present. To do this, a peak-counting routine was written. First, the cluster, which consists of a linked-list of sequences, is converted to a pixel array which contains all of the pixels in the cluster plus a one pixel border around the edges of the array. The array is dimensioned to cover the full extent of the cluster (plus the border). Unused pixels in the array are initialized to zero.

The peak-counting routine requires a second pixel array, known as a “shadow” array. The shadow array has the same dimensions as the ADC pixel array, and is initialized so that all pixels in it are zero. The shadow array is used to track whether a pixel has already been used in a peak (see below)

The peak-counter creates a linear array of all the pixels in the cluster, and sorts them in descending order of their ADC values. Then the peak-counter loops through this list, starting with the pixel with the highest ADC value.

The goal is to find all pixels whose ADC values are larger than any of their neighbors. For each pixel in the ordered list, the same procedure is

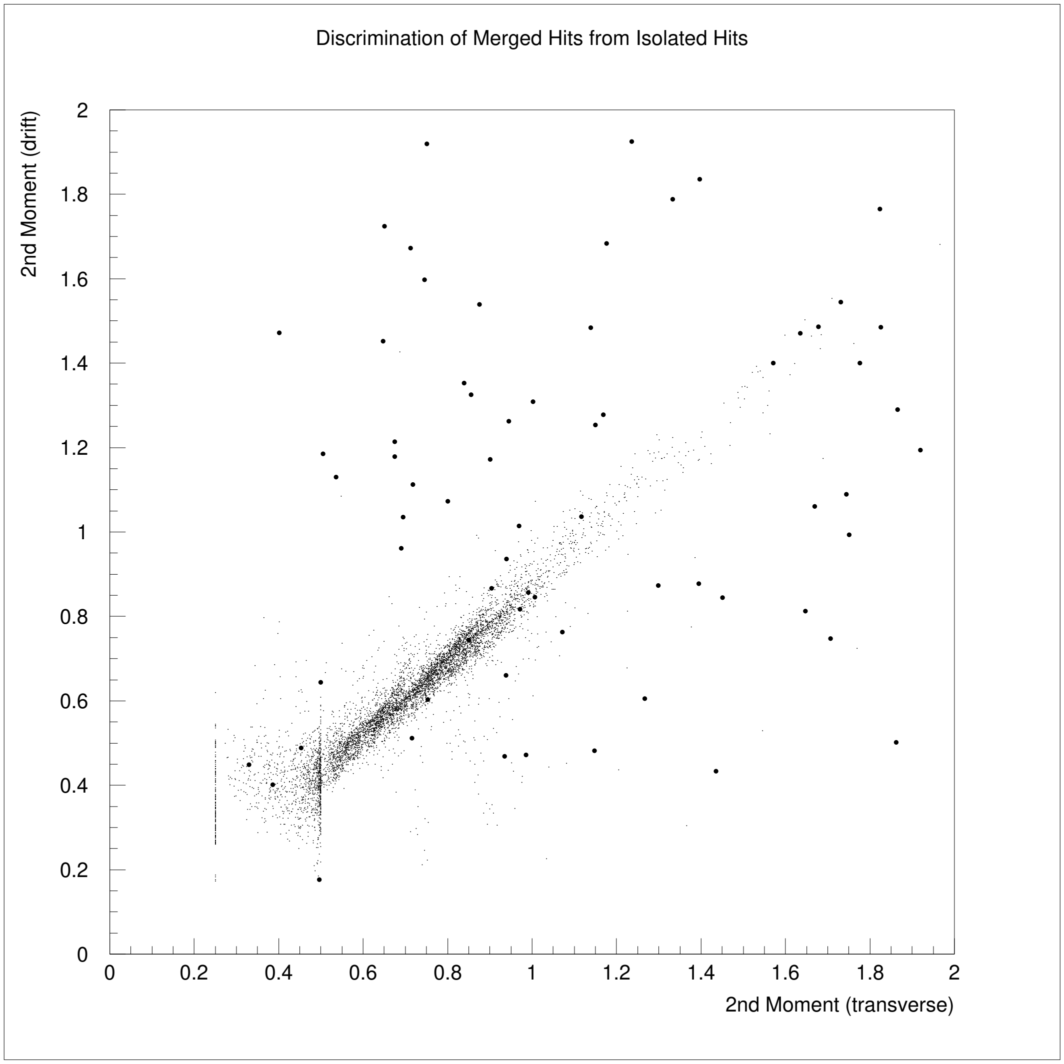

Figure 4.4: Distribution of cluster 2nd moments

followed: the shadow array is queried to see if the pixel in question has been used in a peak. If not, a new peak is added to the list of peaks, and the 3 x 3 grid in the shadow array centered on the pixel is marked as being used. If the pixel has been used already, the 3 x 3 grid in the shadow array is marked anyway. The effect of this is that existing peaks will "use" the pixels around them, and the effect will trickle down the peak unless a larger ADC value is encontered first. The processing of the pixels in the list continues until some lower cut-off value is reached.

Once all of the peaks have been found, a peak-to-valley cut must be applied, else noise might cause spurious peaks to be found. This mechanism is relatively straight-forward. For each peak in the list, a line is drawn between that peak and its nearest neighbor. The pixels along this line are searched to find the smallest ADC value. Once this has been found, the ratio of the ADC value of the peak to the ADC value of the valley is computed. If this ratio is below a cut-off value, the peak in question is removed from the peak list. This process continues until all peaks have been validated or removed.

When the peaks have been counted, a processing decision is made. If there is only one peak present in the cluster, the cluster is processed by the normal isolated-hit routines. If more than one peak is present, a different path is taken.

At this point, no deconvolution of the peaks is performed. The nine pixels centered on the peak are used to determine the centroid and widths of the peak (which allows determination of the total charge in the peak). In cases where the peaks are very close to each other, this method will certainly result in a larger average error in the determination of the peak centroid (and thus the position of the resultant space point), but it is not clear at this time that least-squares fitting of the peaks will result in a siginifcantly better result.

Finally, all clusters, whether isolated or merged, will result in at least one entry being added to the space point table. For each space point to be added, the centroids supplied are converted from (anode, time bin) coordinate space to (drift, transverse) coordinates on the wafer face, in cm. This conversion includes an empirically-determined (and user-specifiable) correction to the drift distance. This correction is necessary due to centroid shifts introduced by the response time of the PASA. The correction is specified as a 10th order polynomial, although presently only terms up to cubic are used.

Figure 4.5 shows the error in the centroid positions at various drift distances without the correction. Figure 4.6 shows a plot of the average errors as a function of the drift distance (determined from the results shown in figure 4.5), with a least-squares fit to a cubic function. Finally, figure 4.7 shows the errors in the drift-direction centroid positions with the correction applied.

4.3 Simulations of detector performance

With a signal simulator and an analysis chain, it is possible to simulate and study virtually all aspects of the performance of the SVT. Areas of interest include efficiency, position resolution, charge resolution (all of which can depend on the noise levels in the system), and more.

STAR uses the GEANT pacakge from CERN to simulate the propagation of particles from an initial distribution, including all relevant physical processes related to the passage of the particles through the various detector

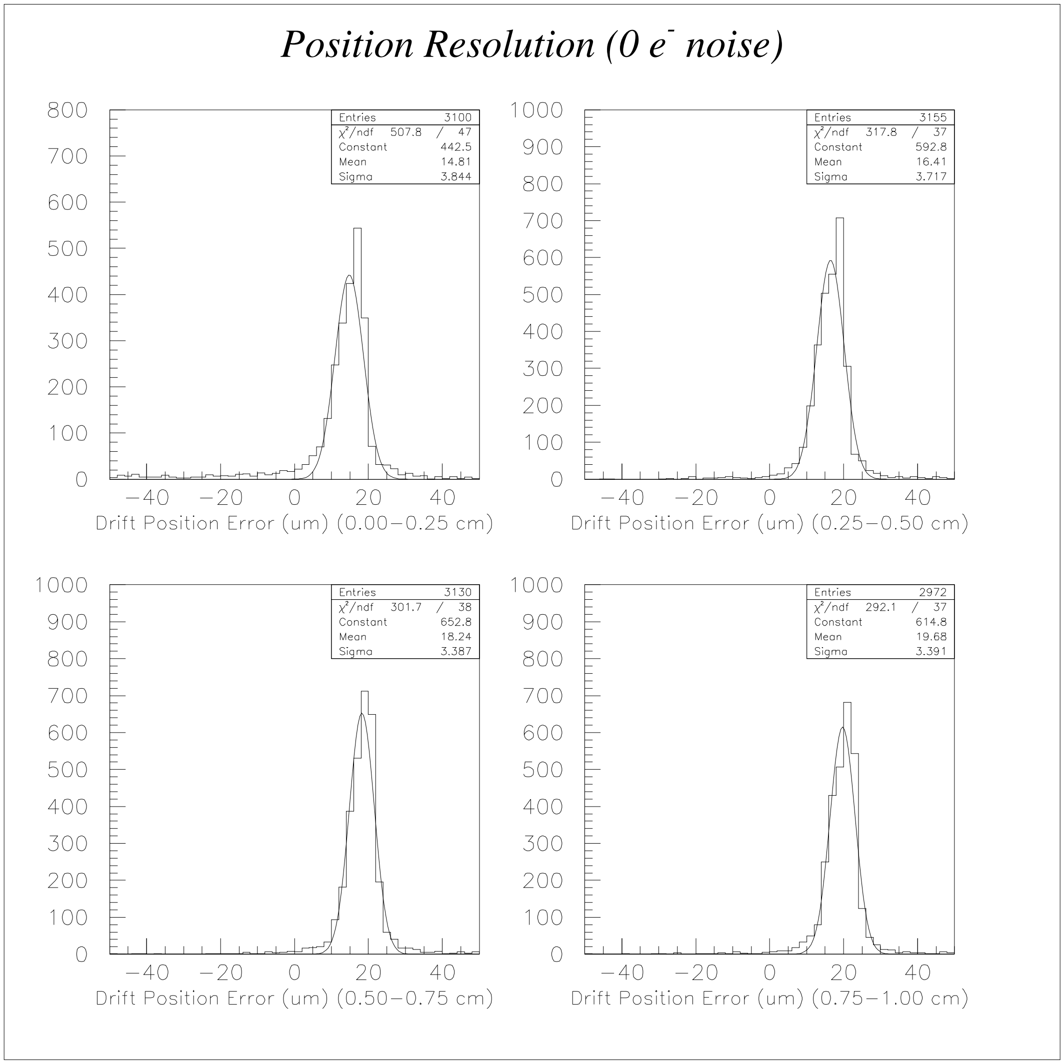

Figure 4.5: Uncorrected errors in drift-direction positions

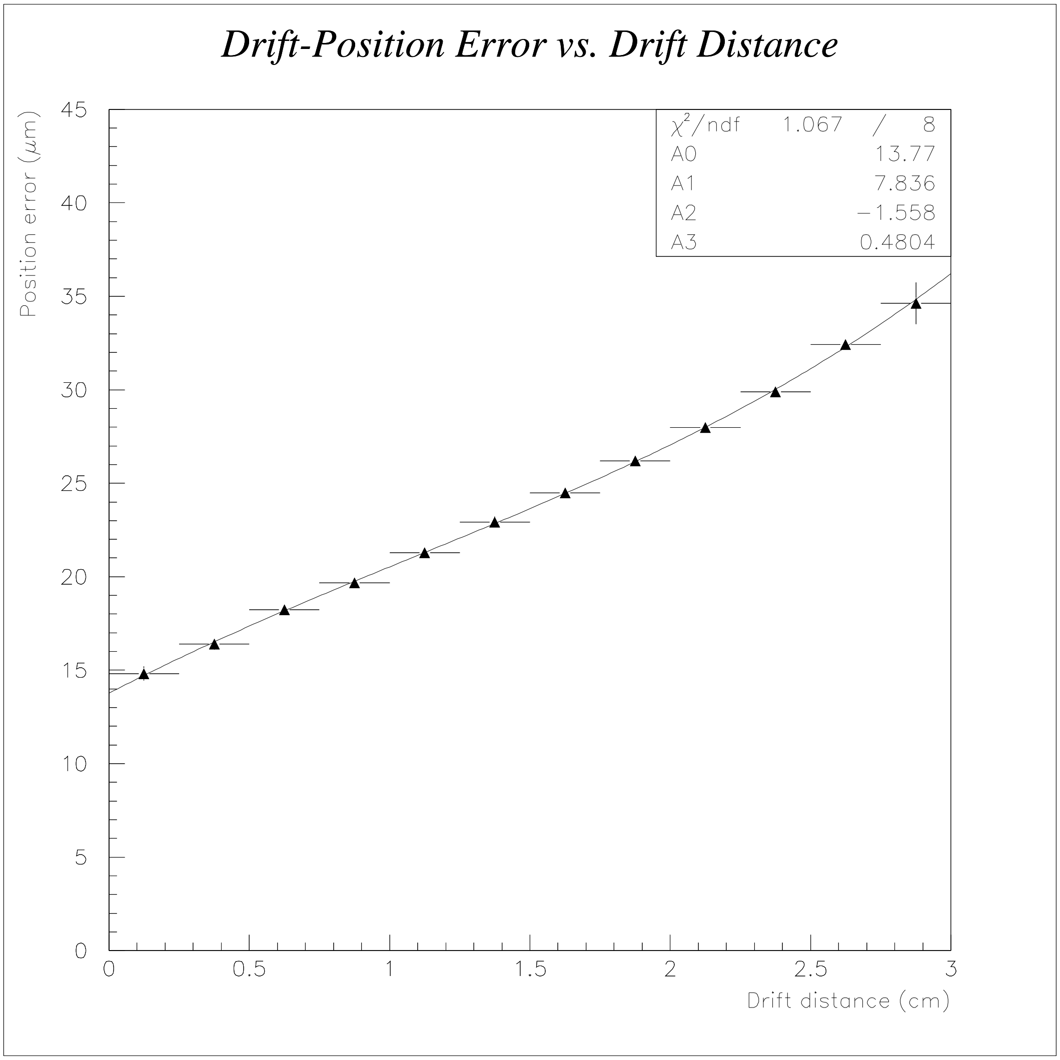

Figure 4.6: Fitting the drift-position errors

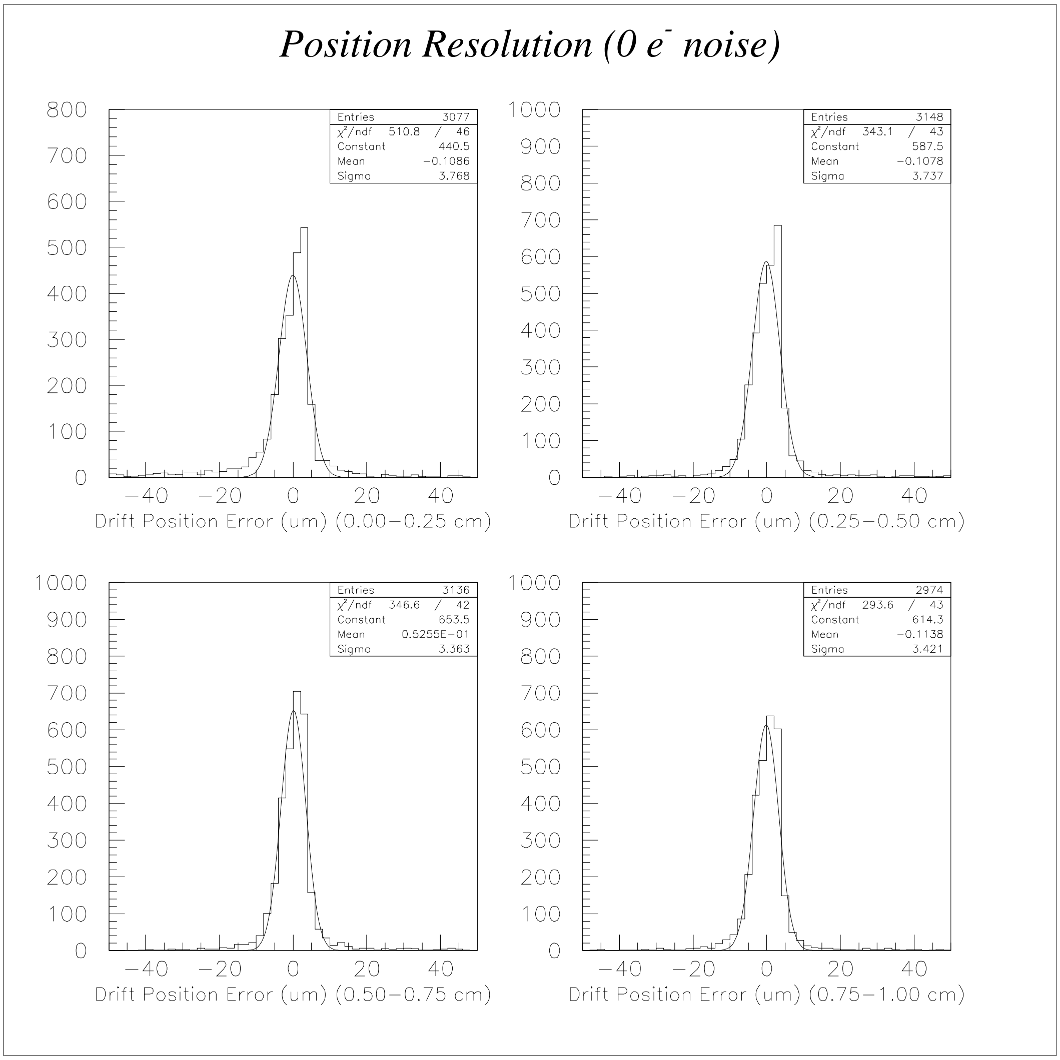

Figure 4.7: Errors in drift position after correction

elements. The initial particle distributions are generated by any of several "event generators" which try to produce realistic particule distributions based on a theory of interest to the author of the generator.

First the GEANT hits on individual SVT wafers are converted to positions on the wafer face. Then the input hits are passed through the simulation chain to produce a pixel array which represents the data as it would be returned from the data acquisition system. The next step is to analyze the pixel array in order to produce a table of reconstructed space points. Finally, one must match entries in the the input monte carlo hit table to entries in the space point table in order to evaluate the performance of the data analysis code. To this end, an evaluation module was written.

The matching mechanism used in the evaluator is a simple one. For each entry in the monte carlo table, all entries in the space point table are searched to find the set of space points which are on the same wafer as the monte carlo (MC) hit. Then the matcher finds the space point (from that set) which is closest to the MC hit. A maximum acceptance is set; space points which are farther than a maximum distance from the MC hit in question are not considered for matching. Currently this maximum distance is set at 500 μm.

Once an entire event is analyzed and all the matches are obtained, it is possible to calculate the quantities of interest.

4.3.1 Detector efficiency

The question of detector efficiency is this: what percentage of particle crossings result in an entry in the space point table? Geometric acceptance plays a large role in the answer to this question, but that feature is not of interest for this discussion.

In order to eliminate the issue of geometric acceptance from the discussion, the efficiency question must be re-phrased as: what percentage of all particles which cross a wafer *in the wafer's active area* result in an entry in the space point table?

The efficiency can be calculated as the ratio of the number of successfully reconstructed particle crossings to the number of input particle crossings. A reconstruction is considered to be successful if the space point can be matched to a monte carlo point on the same wafer, as described in the previous section.

In addition to efficiency, it is possible to calculate the "purity" of the simulation. Purity is the number of "ghost" or false hits generated by the analysis process, expressed as a percentage of the total number of input hits. A purity of 100% implies that no ghost hits were generated.

Clearly the efficiency can depend on several conditions. Foremost of these is the zero-suppression threshold setting, which must be set according to the noise level. Figure 4.8 shows a plot of the efficiency and purity, simulated from a single Hijing Au+Au event (100 Gev-A, central collision) with 400 RMS electrons noise (see below), as a function of the threshold setting.

This plot allows one to select the best set of thresholds for the given noise level. The intent is to choose the thresholds which maximize both the efficiency and the purity. Clearly there is a trade-off; one can attain higher efficiency by lowering the acquisition threshold, but this will cause lower purity, since noise pixels will trigger the hit-finder more often. Raising the thresholds will have the reverse effect; the efficiency will decrease while the purity increases.

The two-threshold algorithm adopted by STAR takes advantage of this effect; the efficiency is sensitive mostly to the lower threshold, while the purity is sensitive mainly to the high threshold.

Clearly, the noise levels will affect the efficiency, as well. Figure 4.8 shows the effects of noise on the efficiency by plotting the efficiency and purity as a function of threshold settings for two different noise levels, namely 800 and 400 RMS electrons.

"400 RMS electrons" means that the input voltage at the pre-amplifier stage of the PASA unit had a noise signal chosen randomly from a gaussian distribution with a width of 400 electrons. A noise signal was applied to every time sample of the pixel array. Since the gain of the PASA units is 9.5 μV $/e^-$, this translates into an RMS noise level of roughly 4 mV in each pixel. The noise level chosen for this plot is the same as that in the design specifications for the PASA units in the SVT data acquisition system.

Clearly the efficiency of the SDDs is not 100%. The major reason for this is that the zero-suppression routines work by requiring a minimum number of samples to be above threshold. This number is currently 3 samples. The practical implication of this is if the centroid of the charge distribution falls within the last three time bins on any anode, the entire sequence will be lost. Since each anode is sampled 128 times, this accounts for $3/128 = 2.34\%$ inefficiency. As noise levels (and hence threshold levels) rise, this inefficiency will rise also. Eventually, the centroid must be at least 4 time bins from the end, which translates into 3.1% inefficiency, and so on.

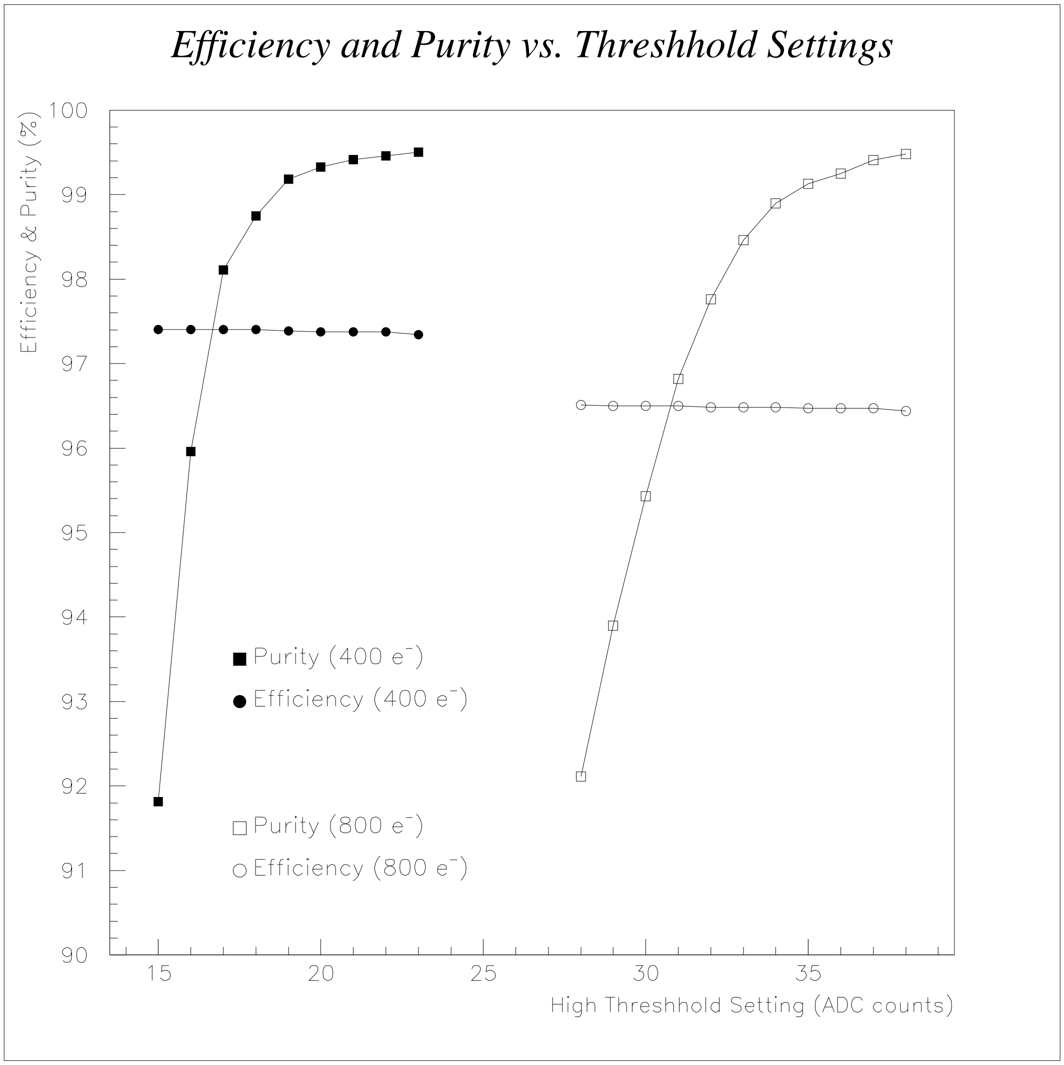

Figure 4.8: Efficiency and purity of the analysis process as a function of threshold settings

Tables 4.1 and 4.2 contain the complete results for detector efficiency. Each table contains data from simulations at several different acquisition thresholds. For each noise level/acquisition threshold setting, the detector efficiency was simulated at a range of trigger thresholds. In all cases, `n_lo = 3` and `n_hi = 1`. Note that for maximum efficiency, it is desirable to set the acquisition threshold as low as possible, but to attain greatest purity, high trigger thresholds are required. Figure 4.8 is a plot of the data points from 400 $e^-$with an acquisition threshold of 9 and 800 $e^-$and an acquisition threshold of 18.

4.3.2 Detector occupancy

One advantage of Silicon Drift Detectors is that they typically have very small electron cloud sizes, which leaves most of the surface area of the detector unoccupied during normal events. Simulating the total occupancy of SDDs is of interest because it provides a way to estimate how much data the Data Acquisition (DAQ) system will have to transport and analyze during normal operation. Further, the total occupancy of the SDDs is important because it indicates what fraction of all hits will suffer from "merging" with nearby hits.

Simulated detector occupancy can be measured by simply counting the number of pixels above threshold in a typical event. Since the three layers of the STAR SVT are at three different radii from the vertex, it is necessary to measure the occupancies in each layer separately.

Table 4.3 shows the occupancies in the three layers of the SVT. A single standard central collision event (from the Hijing event generator at 100 GeV-A) was processed, and the results were averaged over each layer (approximately

Noise (e^-)	Acq. Thresh	Trigger Thresh	MC hits	Space Points	good	bad	eff.	pur.
400	8	15	6968	7392	6787	605	97.4	91.8
400	8	16	6968	7073	6787	286	97.4	96.0
400	8	17	6968	6918	6787	131	97.4	98.1
400	8	18	6968	6873	6787	86	97.4	98.7
400	8	19	6968	6842	6786	56	97.4	99.2
400	8	20	6968	6831	6785	46	97.4	99.3
400	8	21	6968	6825	6785	40	97.4	99.4
400	8	22	6968	6822	6785	37	97.4	99.5
400	8	23	6968	6817	6783	34	97.3	99.5
400	9	14	6968	7306	6777	529	97.2	92.8
400	9	15	6968	7066	6777	289	97.2	95.9
400	9	16	6968	6917	6777	140	97.2	98.0
400	9	17	6968	6841	6777	64	97.2	99.1
400	9	18	6968	6817	6776	41	97.2	99.4
400	9	19	6968	6801	6775	26	97.2	99.6
400	9	20	6968	6795	6774	21	97.2	99.7
400	10	16	6968	6826	6772	54	97.2	99.2
400	10	17	6968	6798	6771	27	97.2	99.6
400	10	18	6968	6790	6770	20	97.2	99.7

Table 4.1: Efficiency and Purity at 400 e^- noise

Noise (e^-)	Acq. Thresh	Trigger Thresh	MC hits	Space Points	good	bad	eff.	pur.
800	16	30	6968	7413	6740	673	96.7	90.9
800	16	31	6968	7203	6740	463	96.7	93.6
800	16	32	6968	7054	6739	315	96.7	95.5
800	16	33	6968	6946	6739	207	96.7	97.0
800	16	34	6968	6882	6739	143	96.7	97.9
800	16	35	6968	6843	6739	104	96.7	98.5
800	16	36	6968	6833	6739	94	96.7	98.6
800	16	37	6968	6811	6739	72	96.7	98.9
800	16	38	6968	6798	6737	61	96.7	99.1
800	17	29	6968	7385	6737	648	96.7	91.2
800	17	30	6968	7202	6737	465	96.7	93.5
800	17	31	6968	7059	6737	322	96.7	95.4
800	17	32	6968	6961	6736	225	96.7	96.8
800	17	33	6968	6884	6736	148	96.7	97.9
800	17	34	6968	6838	6736	102	96.7	98.5
800	17	35	6968	6810	6736	74	96.7	98.8
800	17	36	6968	6801	6736	65	96.7	99.0
800	17	37	6968	6788	6736	52	96.7	99.2
800	17	38	6968	6778	6735	43	96.7	99.4
800	17	39	6968	6772	6735	37	96.7	99.5
800	18	29	6968	7161	6724	437	96.5	93.9
800	18	30	6968	7046	6724	322	96.5	95.4
800	18	31	6968	6945	6724	221	96.5	96.8
800	18	32	6968	6877	6723	154	96.5	97.8
800	18	33	6968	6828	6723	105	96.5	98.5
800	18	34	6968	6798	6723	65	96.5	98.9
800	18	35	6968	6781	6722	59	96.5	99.1
800	18	36	6968	6773	6722	51	96.5	99.2
800	18	37	6968	6762	6722	40	96.5	99.4
800	18	38	6968	6755	6720	35	96.4	99.5

Table 4.2: Efficiency and Purity at 800 e^- noise

Layer	Clusters per Detector	Sequences per Detector	Pixels per Detector	Pixel occupancy
Inner	80	275	1137	1.9 %
Middle	32	114	475	0.8 %
Outer	19	66	270	0.4 %
Whole SVT	32	113	468	0.8 %

Table 4.3: SVT detector occupancies by layer

2000 hits per layer) in order to obtain these numbers.

4.3.3 Position resolution

Of critical importance to the performance of the SVT is the position resolution of the individual Silicon Drift Detectors. Large position errors will result in negative impact on the physics capabilities of the detector, ranging from poor vertex position resolution to potentially large errors in the determination of tracking parameters (namely particle momenta).

Many things can affect the position resolution of the detector. Of primary importance is the effect of noise on the input signals. The effects of noise are several. As mentioned in section 4.3.1, the presence of noise requires the use of some sort of discriminator, which then requires the setting of thresholds. This can lead to lost hits, especially near the edges of the detector. Additionally, the discriminator can truncate the sequences which are passed to the cluster finder, potentially causing shifts in the centroids of the measured charge distribution.

Noise can also affect the centroids directly. Since noise appears with an essentially random distribution, it is possible (and likely) that noise will appear unevenly on any given peak. If the noise has greater amplitude on one side of the peak than on the other, a centroid shift results.

Figure 4.9 shows four histograms of the differences between the positions of the monte carlo hits and the matching reconstructed space points (in μm), binned into six different drift regions.

The centroids of the position-error distributions are all close to zero, as expected; the effects of noise should average to zero. The quantity of interest is the widths of these distributions. That number should be considered as the resolution of the detector; it gives a measurement for how much the position of the reconstructed space point can be expected to deviate from the position of the actual hit.

Figures 4.10 and 4.11 plot the resolution (in the drift and transverse directions), measured as described above, for an input noise of 400 RMS e^- at four different threshold settings. It is clear that the lowest threshold produces the best results. The lowest threshold used was 10 ADC counts, which corresponds to 5 mV. The 400 e^- noise level corresponds to 3.8 mV, so it would seem that a threshold set slightly above the noise level is the best one to use. When the threshold was set to 7 ADC counts (3.5 mV), the simulation aborted due to overflowing of the sequence tables.

Raising the threshold produces worse position accuracy because the higher thresholds can cause samples to be lost from the tails of the charge distribution.

Figures 4.12 and 4.13 are plots of the position resolution of the SVT, measured as described above, for various noise levels. As expected, lower noise

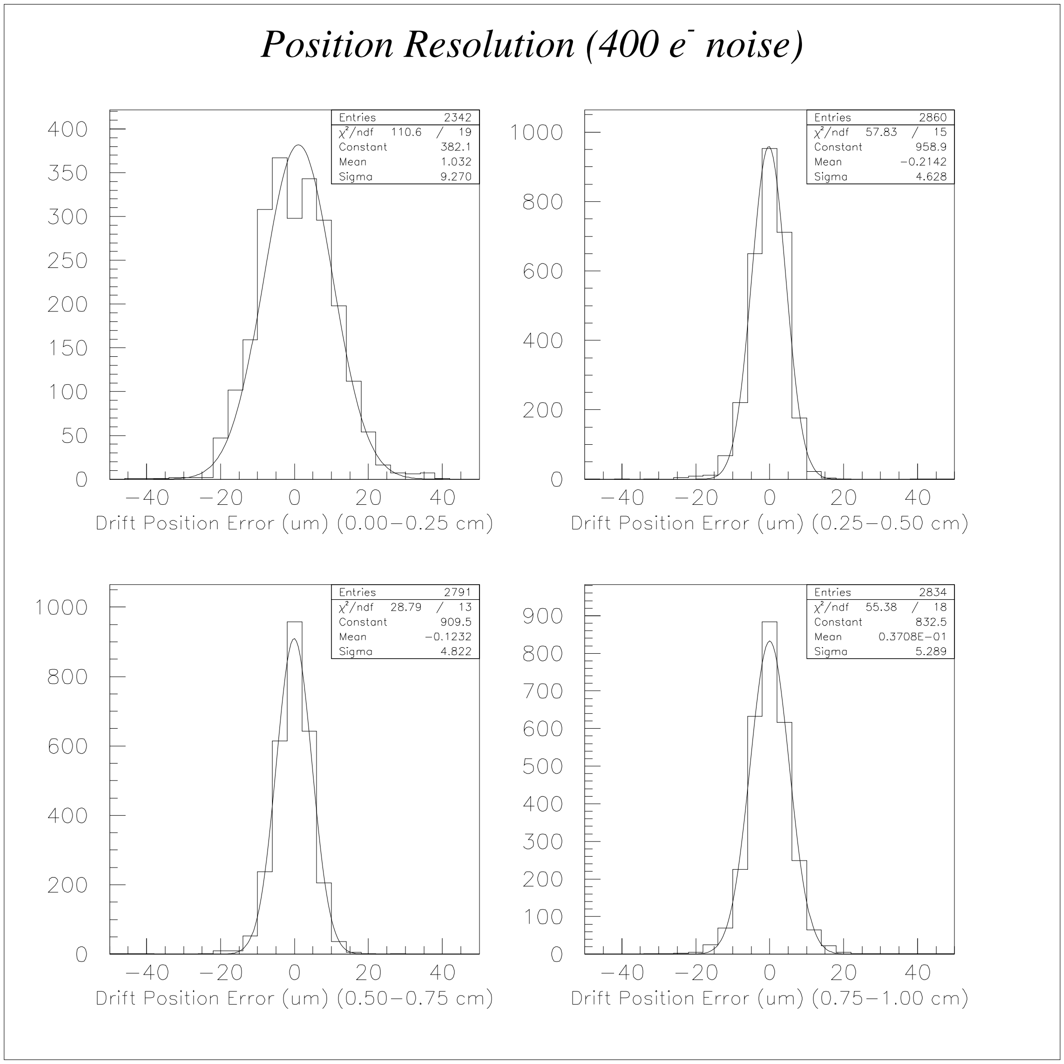

Figure 4.9: Errors in the reconstructed hit position at various drift distances

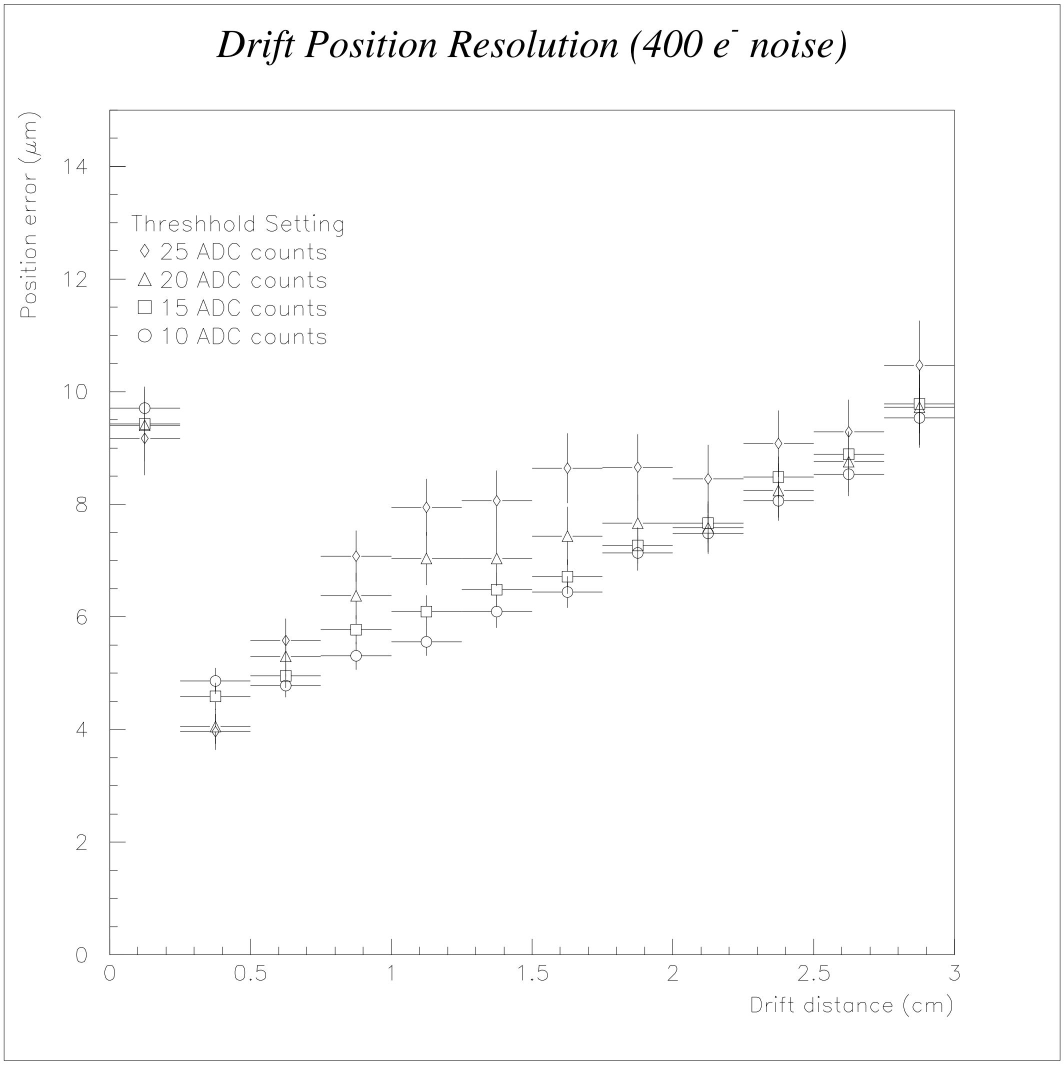

Figure 4.10: Drift position resolution as a function of drift distance and threshold levels

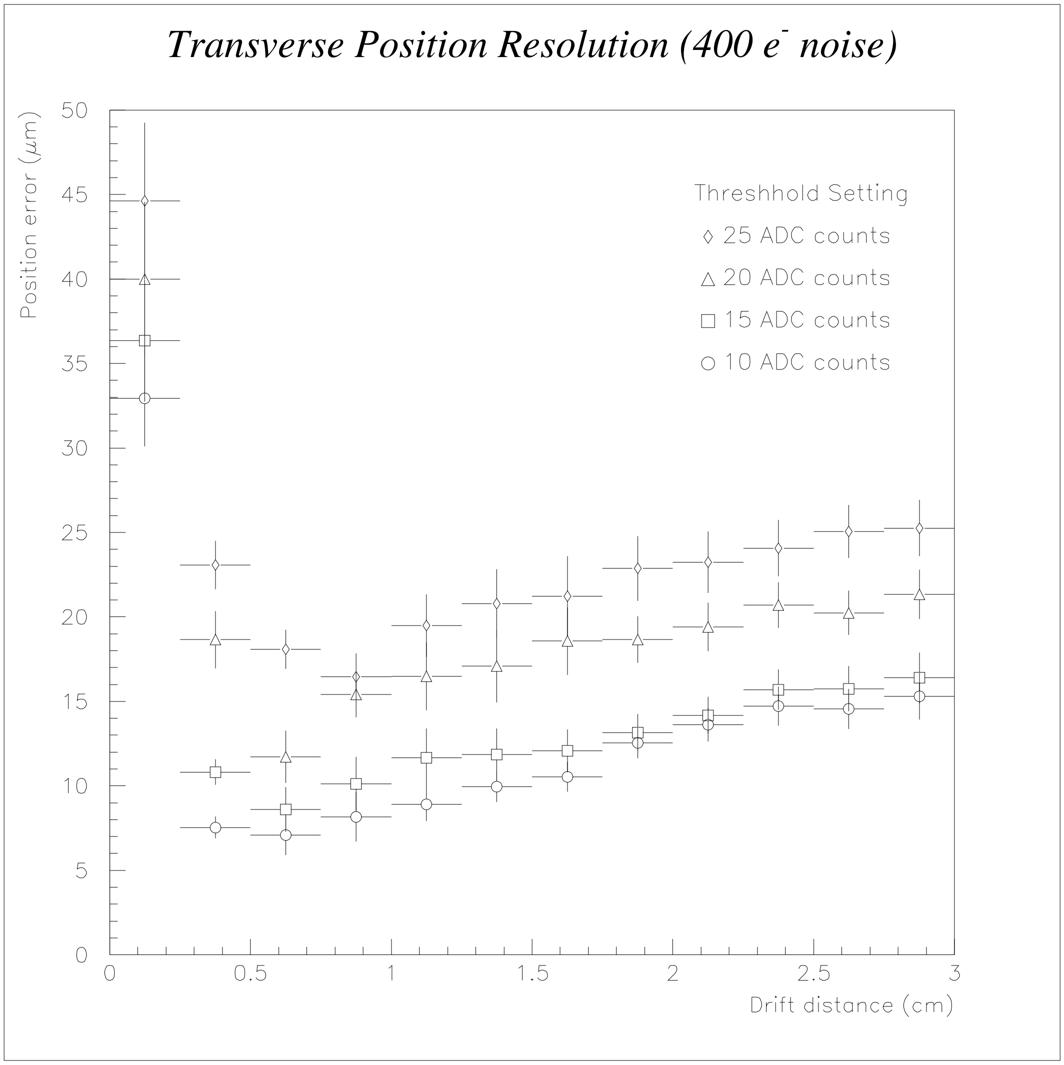

Figure 4.11: Transverse position resolution as a function of drift distance and threshold levels

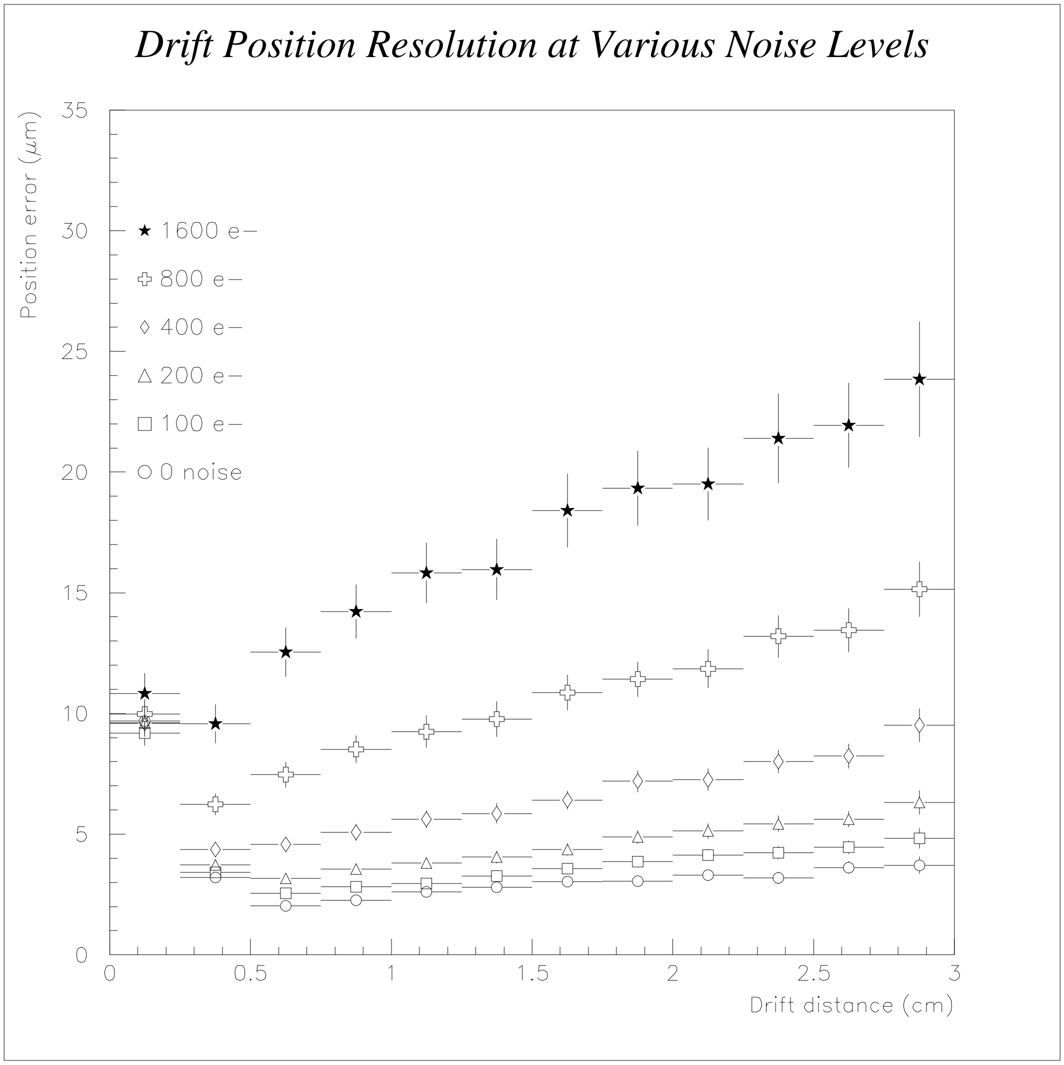

Figure 4.12: Drift position resolution as a function of drift distance and noise levels

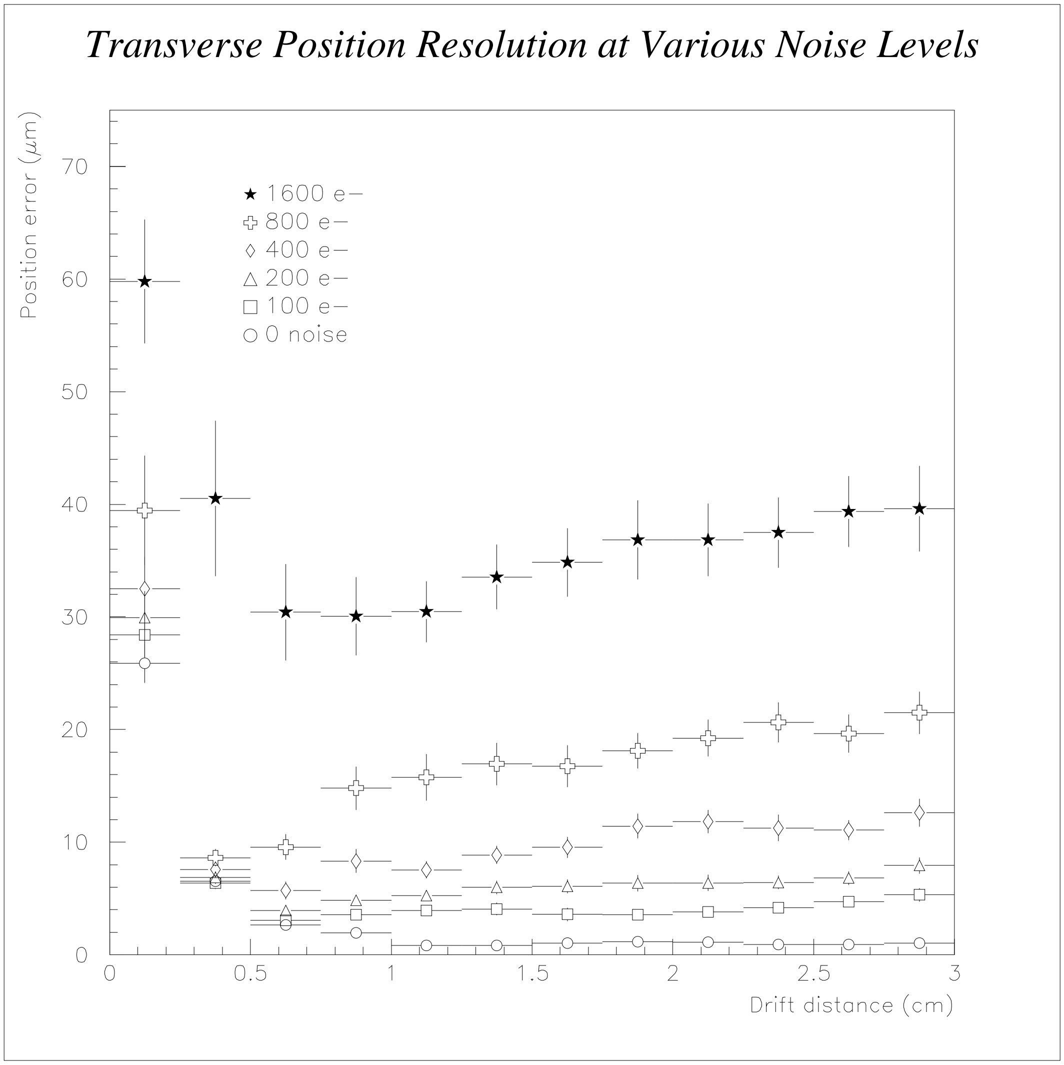

Figure 4.13: Transverse position resolution as a function of drift distance and noise levels

results in better position resolution.

4.3.4 Charge resolution

Among the capabilities of Silicon Drift Detectors is the ability to measure the total ionization charge (i.e. deposited energy). When used in conjunction with the particle momentum (determined from tracking), particle identification is possible.

In an ideal world, the SDD would collect (and thus measure) all of the negative ionization charge from each particle crossing. However, several factors can cause charge to be lost before collection. In the SVT, the largest of these factors is the fact that the charge cloud spreads as it drifts. Since there is noise in the system, some sort of discriminator must be used to differentiate signal from background.

The discriminator used in the SVT system is based on a threshold mechanism, which means that signals which fall below a cut-off value are discarded. The practical implication of this is that the tails of a charge distribution may get cut off by the zero-suppressor. In the drift direction, this effect may be small, as only one or two time buckets will fall below threshold.

In the anode direction, however, the effect can be much larger, as entire sequences on the tails of a given cluster might fall below threshold. This is especially true for hits which occur at large drift distances, as the charge cloud will become very large in extent, leaving very long tails. The signal on the central few anodes will suffer very little, but there can be several anodes on the fringes of the cluster where the signal fails to reach threshold.

Figure 4.14 shows a plot of the ratio of the energy deposited in the

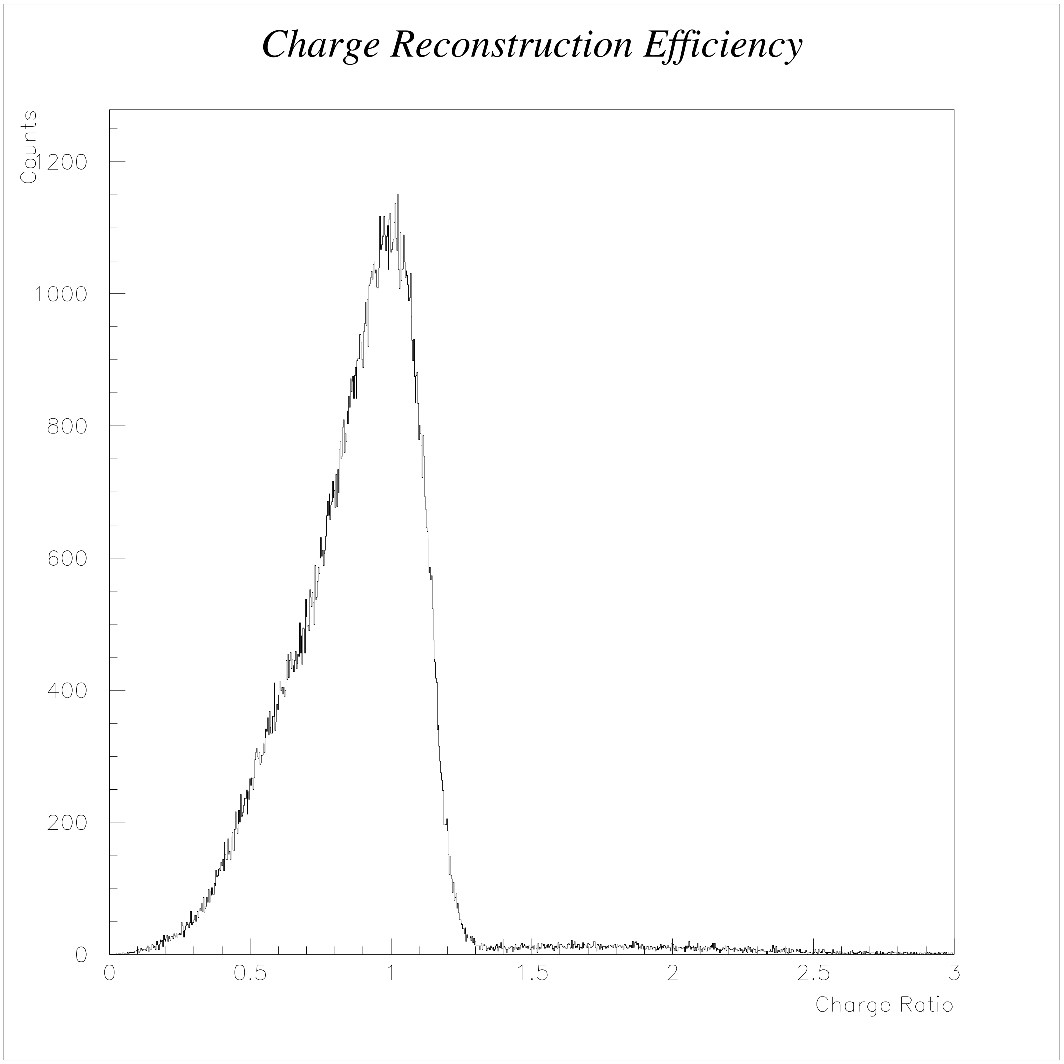

Figure 4.14: Charge reconstruction efficiency at 400 e^{-} noise

silicon (as reconstructed by the analysis software) to the monte carlo deposited energy. This plot represents 10 central collision Au+Au events, averaged over the entire SVT and all drift distances. Perfect charge reconstruction would appear in this plot as a very narrow peak centered around 1.0. Clearly this is not the case.

There are several reasons for the poor results seen in figure 4.14. The first is that the charge is reconstructed by simply adding together the ADC counts of all pixels in each cluster, forming an "integrated charge." This method is not a good one, because it is highly sensitive to noise. Furthermore, high thresholds can cause pixels to be lost, as described above. A more robust way of estimating the deposited energy would be to use the amplitude and widths calculated in the analysis process to calculate an integrated charge.

Another source of "lost charge" in the current analysis software is related to the 10-bit to 8-bit conversion of the ADC values performed by the data acquisition system. As described earlier, the DAQ system converts the 10-bit values from the ADC to 8-bit values using a non-linear lookup table. For the software used to make figure 4.14, those 8-bit values were used directly to calculate the quantities of interest. Clearly, a more accurate method would convert the 8-bit values back to 10-bit values, using the same lookup table used by the DAQ system. This method has been implemented, and the resulting charge-reconstruction histogram is shown below, in fig. 4.15. The width of this peak is noticeably smaller.

Figure 4.16 shows a plot of the charge-reconstruction efficiency of the simulation/analysis package, measured as described above. The quantity plotted is the FWHM of the plot shown in fig. 4.15, evaluated for several different

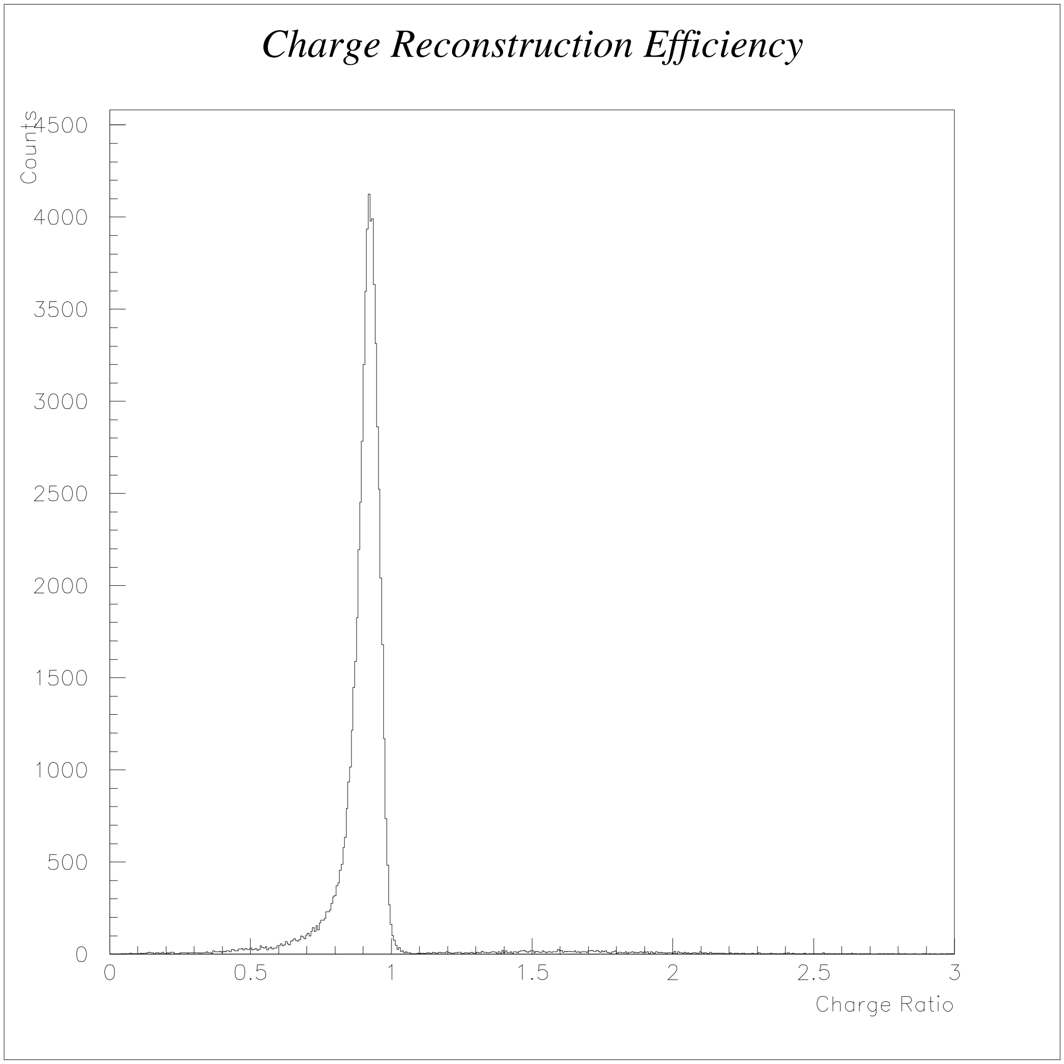

Figure 4.15: Improved charge reconstruction efficiency at 400 e^- noise

runs, with differing amounts of noise.

The first method (described above) for improving the charge resolution has also been implemented. The basis for implementing this method lies in assuming a gaussian form for the charge cloud distribution. Without loss of generality, we will take this distribution to be centered around $x_0 = 0$:

$$q(x) = Ae^{-x^2/2\sigma^2}. \tag{4.1}$$

Then we can write down an expression for the total charge in the cloud:

$$Q_{tot} = \int_{-\infty}^{\infty} q(x)dx. \tag{4.2}$$

Change variables to $u = x/\sqrt{2}\sigma$:

$$\begin{aligned} Q_{tot} &= A\int_{-\infty}^{\infty} e^{-u^2}\sqrt{2}\sigma du \\ &= \sqrt{2}\sigma A\int_{-\infty}^{\infty} e^{-u^2} du \\ &= \sqrt{2}\sigma A\sqrt{\pi} \\ &= \sqrt{2\pi}\sigma A \\ &\approx 2.5\sigma A. \end{aligned} \tag{4.3}$$

For the two-dimensional case, the charge distribution is a two-dimensional gaussian:

$$q(x,y) = Ae^{-\frac{x^2}{2\sigma_x^2}-\frac{y^2}{2\sigma_y^2}}. \tag{4.4}$$

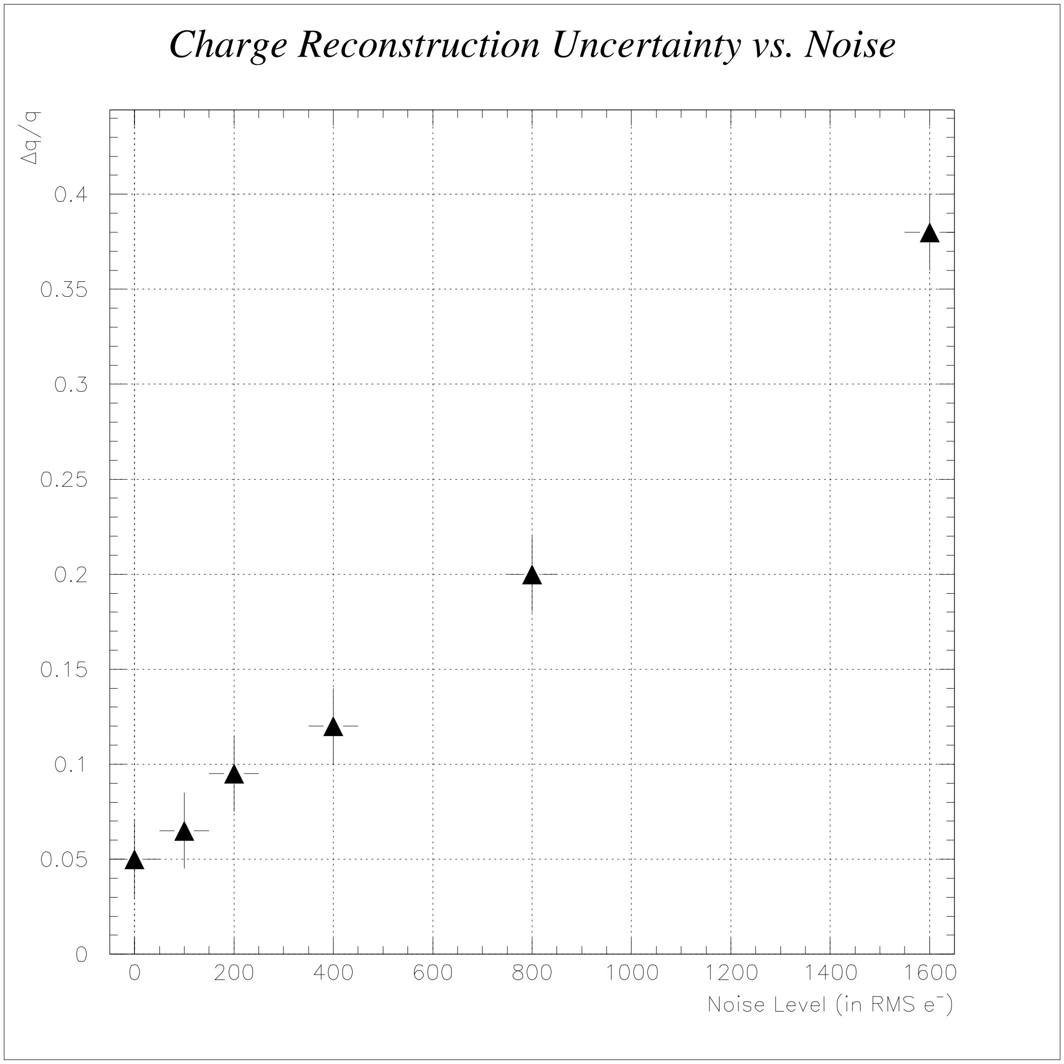

Figure 4.16: Charge reconstruction efficiency vs. noise levels

The solution, obtained in the same manner as eq. 4.3, is given by

$$Q_{tot,2D} = 2\pi\sigma_x\sigma_y A. \tag{4.5}$$

Figure 4.17 shows a histogram of the reconstruction effciency achieved using this method to estimate the charge in a cluster, and compares it to the results obtained by the simpler "summed pixels" method. The results are not good enough presently to justify using this method in actual analysis, but they are encouraging. It is likely that the fitted charge suffers from poor estimates of the amplitude and cluster sigmas. Perhaps if more accurate estimates were made, the fitted value for the charge would be more accurate. It is also possible that since the charge collected by each anode does not show a gaussian dependence, that the second moment in the anode direction is not a good estimate for the gaussian sigma of the physical charge cloud.

This method is of practical interest in obtaining an estimate for the charge contained in each peak in a cluster containing merged hits. Since it is not possible simply to sum the pixels in such a cluster, an estimate based on the central-most pixels must be made.

4.3.5 The effects of the 10-bit to 8-bit mapping table

As described in section 4.3.4, the 10-bit values from the ADCs are converted to 8-bit values using a lookup table. These values must be converted back to 10-bit values (using an inverse lookup table) before processing the data. The contents of the initial 10-to-8 conversion table can have a great effect on the capabilities of the SVT.

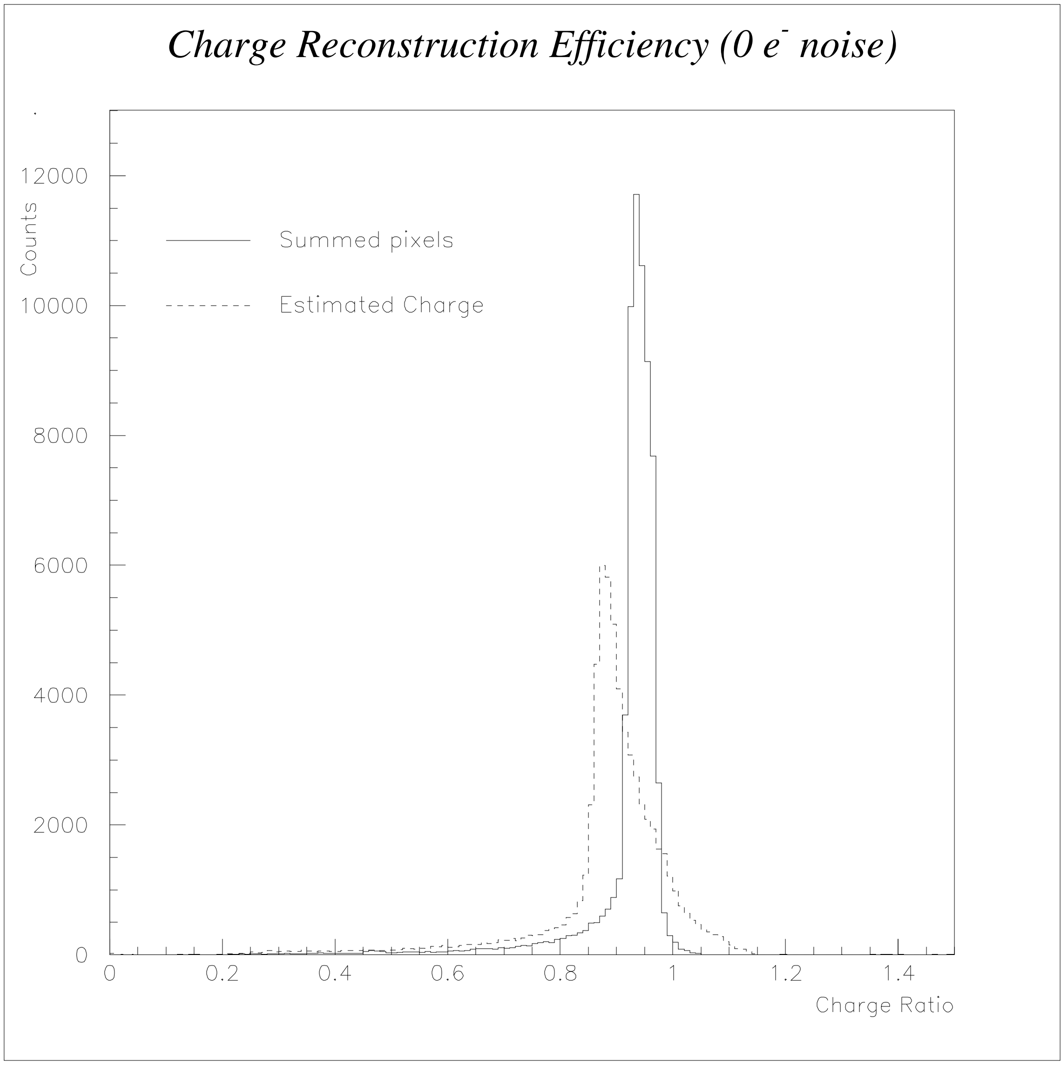

Figure 4.17: An alternate method of reconstructing the total deposited charge in a cluster

The table currently proposed is of the “tri-linear” form described in section 4.2.2. This table is intended to provide good resolutions (both spatial and charge) for minimum ionizing particles, but it remains to be seen how this table will perform for particles which deposit up to 10 times the “minimum ionizing” energy.

The SVT simulation and analysis software allows the effects of the 10-to-8 table to be simulated. As a first attempt at this, a simulation was performed using a purely linear table. This table converted 10-bit values (with a range of 0 to 1023) to 8-bit values (range 0 to 255) by dividing the 10-bit value by 4.

The results are shown below. The position resolution degraded significantly, and the charge resolution became markedly worse. This is due to the fact that the majority of the particles crossing the SVT are minimum ionizing particles; in the “old” table, the signals from these particles undergo no changes as a result of the conversion. In the “new” (linear) table, these signals undergo a transformation which downgrades their granularity by a factor of 4.

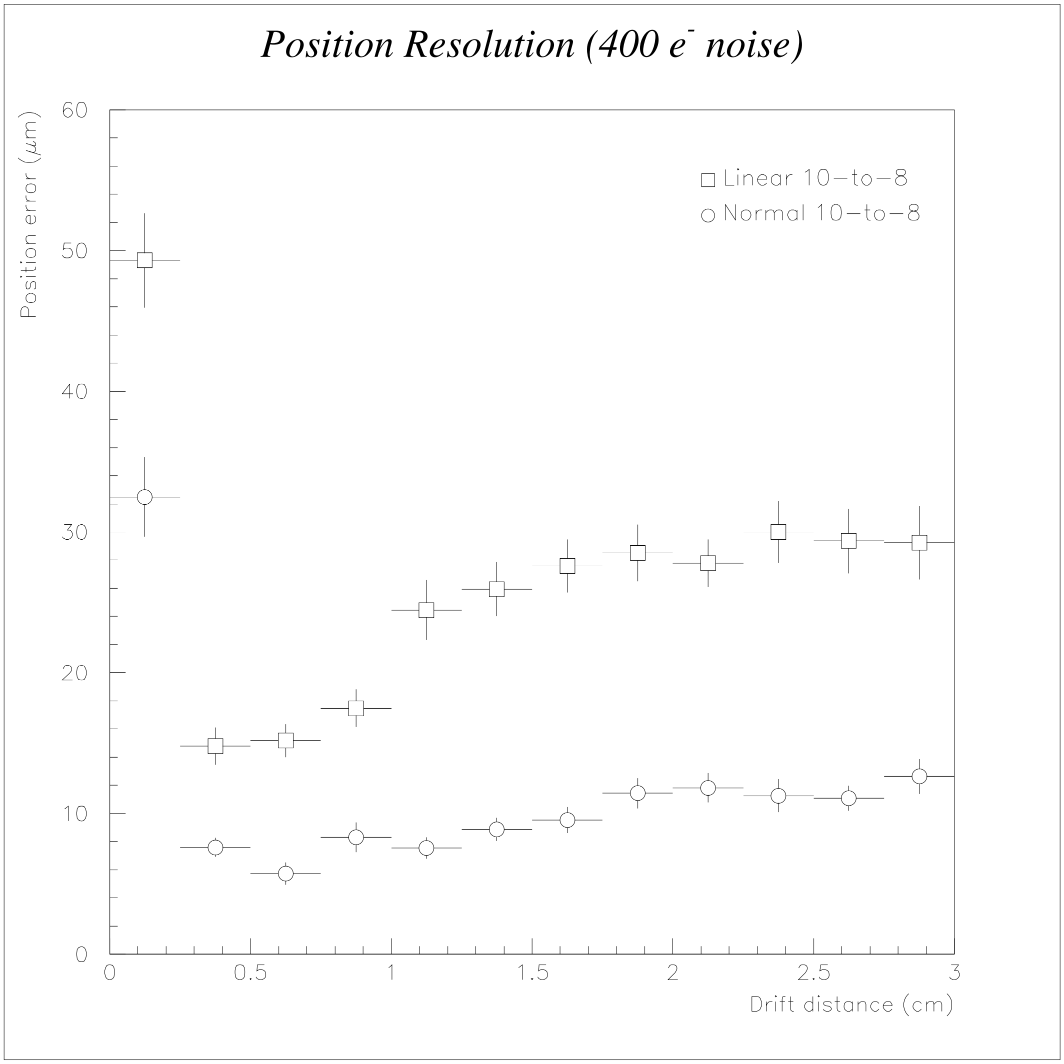

Figure 4.18: Position resolution with a linear 10-bit to 8-bit mapping table

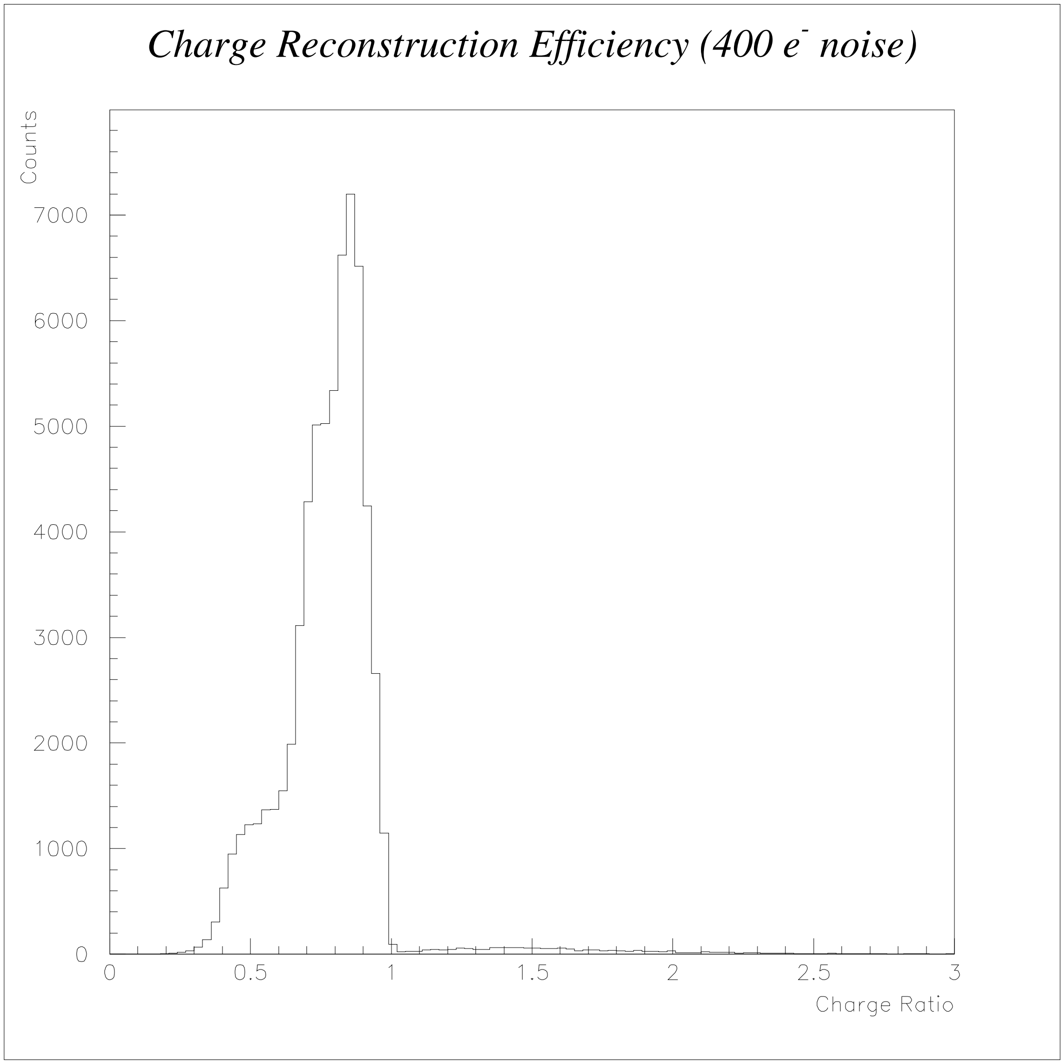

Figure 4.19: Charge reconstruction with a linear 10-bit to 8-bit mapping table

Chapter 5

SDDs in a Vertex Tracker

The Solenoidal Tracker at RHIC (STAR) Collaboration has designed a vertex tracker to be used in conjunction with STAR's main Time Projection Chamber (TPC). This device, the Silicon Vertex Tracker (SVT) is intended to allow STAR to observe particles with low transverse momentum, to allow STAR to make precise measurements of the position of the primary interaction vertex, and to provide STAR with the high spatial resolution required to observe secondary decays which occur close to the primary vertex.

5.1 Overview of the Silicon Vertex Tracker

STAR's Silicon Vertex Tracker consists of 216 silicon drift detectors, mounted in three concentric layers. Figure 5.1 shows a schematic drawing of the SVT layout. Each "super-layer" of the detector consists of two "sub-layers" which are tiled with respect to each other so as to minimize dead areas. Each superlayer consists of a number of "ladders" which carry SDDs. The ladders of the each superlayer contain a different number of SDDs than the other superlayers, because each superlayer covers the same range of pseudorapidity ($-1 < \eta < +1$). The inner layer's ladders carry 4 SDDs, the middle ladders carry 6, and the outer ladders carry 7. Additionally, each superlayer contains a different num-

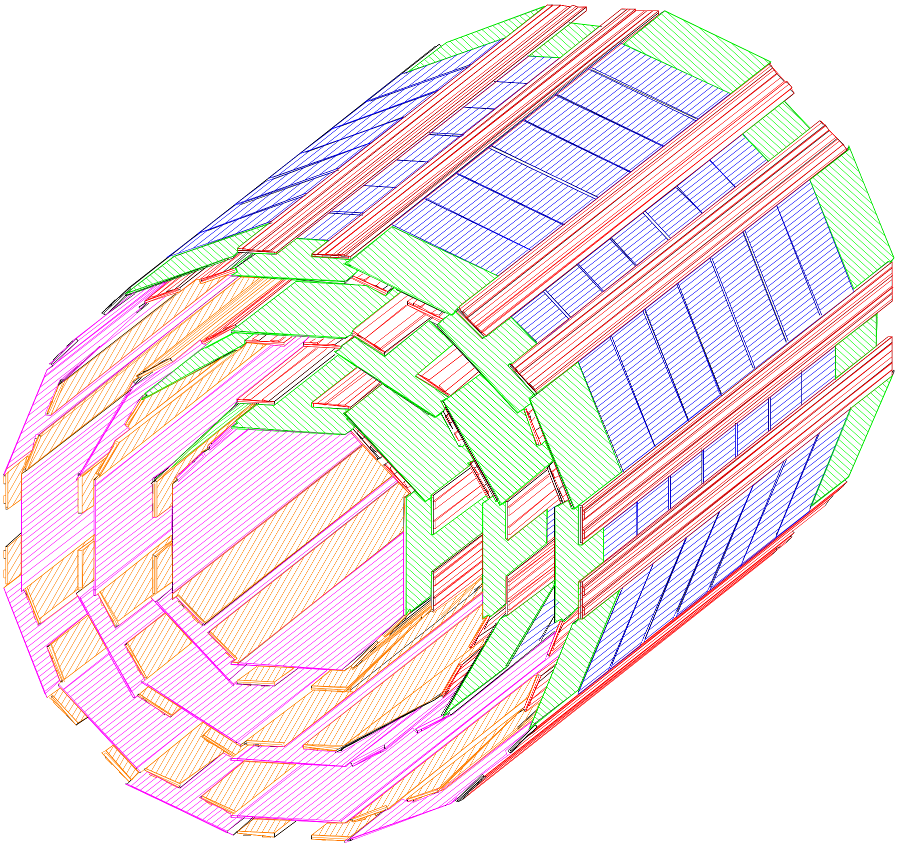

Figure 5.1: Schematic layout of the SVT

ber of ladders due to the differing radii; there are 8, 12 and 16 ladders on the inner, middle, and outer layers, respectively.

The ladders are attached to a Beryllium support frame. Each ladder connects to electronics carriers at both ends of the ladder. The electronics carries constitute the first stage of the front-end readout electronics. The ladders also contain channels for a water-cooling system.

length (active area)	42 cm (max.)
Inner Barrel (radial distance)	5.97 cm
Middle Barrel (radial distance)	10.16 cm
Outer Barrel (radial distance)	14.91 cm
Thickness of wafer (SDD)	300 μm
Total Thickness	900 μm $\sim$ 0.9% X_o
Number of wafers	216
Inner Barrel	32 (8×4)
Middle Barrel	72 (12×6)
Outer Barrel	112 (16×7)

Table 5.1: Specifications of the Silicon Vertex Tracker

The ladders and electronics carriers are supported at each end by two Beryllium end caps, configured as a clam shell to facilitate installation and removal of the detector assembly without breaking the RHIC beam pipe vacuum.

The two end caps are followed by two Aluminum water manifolds which distribute the cooling water to the respective ladder inlets. The detector assembly will be mounted around the beam pipe via two support cones attached to the STAR TPC end caps. Water cooling pipes, high voltage and signal cables are attached to the support cone.

5.2 Simulations of SVT Performance

Using the SSS code described in chapter 3, it is possible to simulate the performance of the SVT in various areas. Of primary interest are the SVT's tracking efficiency, momentum resolution, and particle-ID capabilities.

All of the simulations mentioned below consist of 25 central-collision Hijing Au+Au events at 100 GeV-A. Simulations were performed at two different

SDD noise levels (400 e^- and 800 e^-). For each noise level, an appropriate set of threshholds was chosen based on the results presented in sections 4.3.1 and 4.3.3.

5.2.1 SVT Tracking Efficiency

In order for a particle track to be reconstructed in the SVT, the particle must leave a detectable charge in each layer of the SVT; in other words, every track must have three hits in order to be reconstructed correctly.

This fact explains the two major ways in which tracks can fail to be reconstructed by the SVT. The first way is that on one or more of layers of the SVT, the track crosses a "dead area" of an SDD wafer. The other way is for an individual SDD to lose a hit due to noise or some other signal problem. As shown in section 4.3.1, the detection efficiency of the individual SDDs is expected to be very high, close to 97inefficiencies in the SVT to be the dead areas of the wafers.

Due to geometric considerations alone, the efficiency of the SVT is 95SDDs, this results in an effective efficiency of 92.15reconstruction requires a hit on all three layers of the SVT, the overall efficiency of SVT tracking is expected to be $\epsilon = 0.9215^3 = 78\%$.

Simulations agree with this. A full tracking SVT package has been written for STAR by C. Pruneau, and the SSS package has been successfully integrated with it.

In the simulations, 62067 tracks passed through the SVT fiducial volume. At an SDD noise level of 400 e^-. 46970 tracks were reconstructed successfully, yielding an efficiency of 75.6 %.

When the same simulation was performed for a noise level of 800 e^-, the number of reconstructed tracks dropped to 46584, resulting in an overall efficiency of 75.0 %.

From these numbers in can be concluded that the SVT tracking efficiency does not suffer dramatically from doubling the noise levels in the individual SDDs.

5.2.2 SVT Momentum Resolution

STAR determines momentum by measuring helix parameters of tracks, namely the helix radius R and the helix pitch angle λ. The relationship between these quantities and the particle momentum is given by [45]

$$p\ cos\lambda = 0.2998\ z\ B\ R \tag{5.1}$$

where p is the total momentum in GeV/c, ze is the charge of the particle, and B is the magnetic field strength.

Figure 5.2 shows a histogram of $\Delta p/p$ from the 400 e^- simulation, while figure 5.3 shows a histogram of $\Delta p/p$ from the 800 e^- simulation. In these plots, the quantity histogrammed is

$$\frac{\Delta p}{p} = \frac{p_{monte-carlo} - p_{reconstructed}}{p_{monte-carlo}} \tag{5.2}$$

It is instructive to note that the momentum resolution of the SVT does not change appreciably when the total noise is doubled. As discussed in section 4.3.3, the position resolution of the SDDs degrades by almost 50 % when the noise doubles from 400 e^- to 800 e^-. Clearly, since the momentum resolution

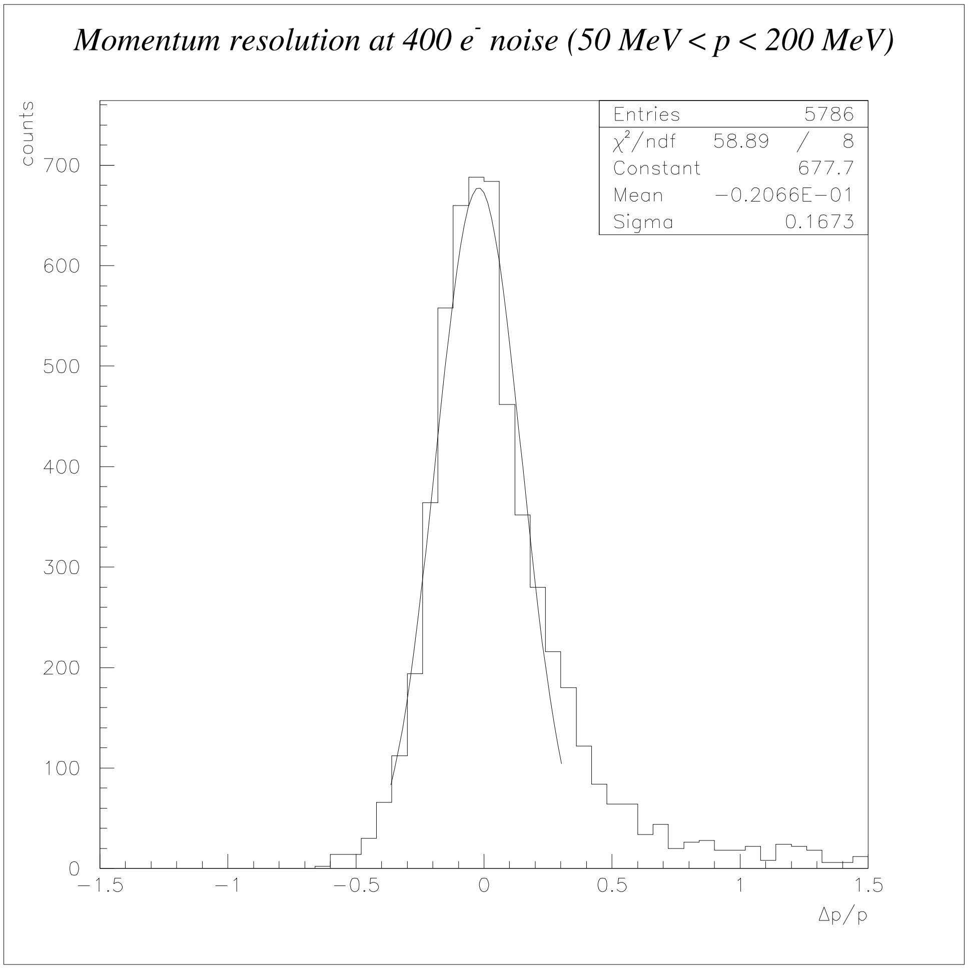

Figure 5.2: Momentum resolution in the SVT at 400 e^- noise

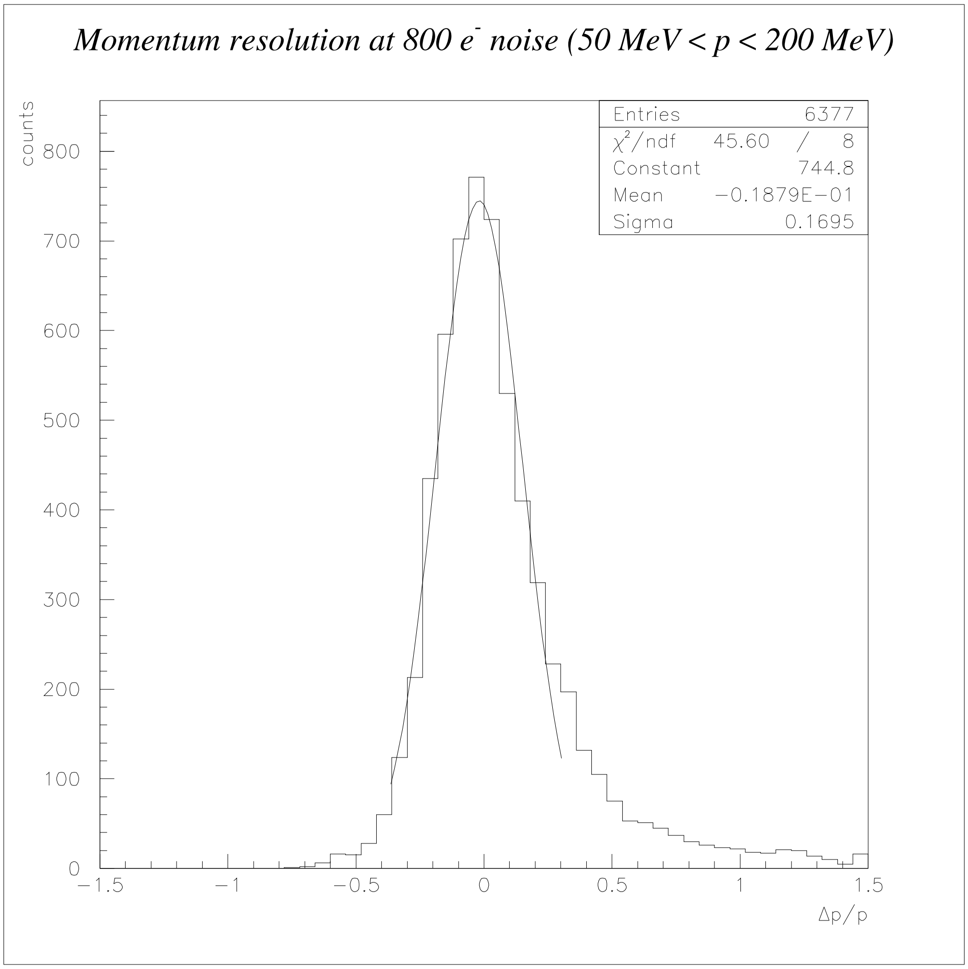

Figure 5.3: Momentum resolution in the SVT at 800 e^- noise

of the SVT did not change substantially, the momentum resolution must be due to some other effect.

Consider the simplest case: a particle traversing the SVT at zero rapidity. First the particle encounters the beam pipe, which is 800 μm of Beryllium.

The width of the angular distribution from multiple scattering is given by [46]

$$\theta_0 = \frac{13.6}{\beta p} z \sqrt{\frac{x}{X_0}} \left[1 + 0.038 \ln \frac{x}{X_0} \right] \tag{5.3}$$

where β is the particle's velocity divided by the speed of light, p is the particle's momentum (in MeV/c), z is the particle's charge number, x is the distance in the medium traversed by the particle, and X_0 is the radiation length of the medium.

For purposes of calculating the angular width of the scattering distribution, let us choose a particle and momentum typical of STAR interactions: a 100 MeV/c pion. Using the radiation length of Beryllium (35.3 cm), $x/X_0 = 0.227\%$. Then it is calculated that the angular width of the scattering distribution is 3.95×10^{-3} rad, or 0.11 deg.

The diameter of the beam pipe is 8 cm, so the radial distance the particle must travel between exiting the beam pipe and crossing an SDD is 1.97 cm. An angular deviation of 3.95×10^{-3} rad over a distance of 1.97 cm translates into a planar deviation of 38 μm on the wafer face; this is the average amount of "error" which multiple scattering in the beam pipe will introduce into the position measured by the first SDD the particle crosses. Already we see that the position resolution of the SDDs is better than the error introduced by multiple scattering in the beam pipe.

Now let us consider multiple scattering in the wafers of the SDD. Each wafer is 300 μm thick; the radiation length of silicon is 9.36 cm, which means that each SDD wafer represents 0.32 % of a radiation length. The same calculation as above, performed for the same particle traversing the SDD wafer at normal incidence, yields an angular distribution of 4.78×10^{-3} rad (0.14 deg). The radial distance between the first and second layers of the SVT is 4.19 cm, which implies an expected position deviation from multiple scattering in the first layer (measured on the second layer) of 102 μm.

The case presented above is the "best case" for multiple scattering. If particles traverse the SVT with non-zero pseudo-rapidities they will traverse more distance in the beam pipe and the wafers, and have longer to travel between the beam pipe and each layer of the SVT. Hence, angular widths of the multiple scattering distribution will be larger, and the position deviations due to multiple scattering will also be proportionally larger.

Clearly the position deviations connected to multiple scattering in the beam pipe and the first two layers of the SVT far outweigh the position resolution of the SDDs. It is for this reason that the momentum resolution of the SVT is not highly sensitive to the position resolution of the SDDs.

Appendix

Appendix A

Derivation of formulae for extracting parameters of a Gaussian distribution from few data points

A.1 Formulae for two data points

Let two data points of a gaussian distribution at ordinates x_0 and x_1 be represented by P_0 and P_1, respectively. Further, let $x_1 = x_0 + \Delta$, where Δ is a constant "mesh size." In this discussion, we will treat the ordinates as discrete integer variables, so that $\Delta = 1$ and $x_1 = x_0 + 1$.

Then the values of P_0 and P_1 can be written down:

$$\begin{aligned} P_0 &= Ae^{-\frac{(x_0-\bar{x})^2}{2\sigma^2}} \\ P_1 &= Ae^{-\frac{(x_1-\bar{x})^2}{2\sigma^2}} \end{aligned} \tag{A.1}$$

where A, $\bar{x}$ and σ are the amplitude, centroid, and width of the gaussian distribution. Then

$$\begin{aligned} \ln P_0 &= \ln A - \frac{(x_0 - \bar{x})^2}{2\sigma^2} \\ \ln P_1 &= \ln A - \frac{(x_1 - \bar{x})^2}{2\sigma^2} \end{aligned} \tag{A.2}$$

and

$$\ln P_0 - \ln P_1 = \ln \frac{P_0}{P_1} = -\frac{1}{2\sigma^2}\left[(x_0 - \bar{x})^2 - (x_1 - \bar{x})^2\right] \tag{A.3}$$

and

$$\frac{1}{2\sigma^2} \ln \frac{P_1}{P_0} = x_0^2 - 2x_0\bar{x} - x_1^2 + 2x_1\bar{x} \quad . \tag{A.4}$$

Now replace x_1 with $x_0 + 1$, yielding

$$\begin{aligned}
\frac{1}{2\sigma^2} \ln \frac{P_1}{P_0} \quad &= x_0^2 - 2x_0\bar{x} - (x_0+1)^2 + 2(x_0+1)\bar{x} \\
&= x_0^2 - 2x_0\bar{x} - x_0^2 - 2x_0 - 1 + 2x_0\bar{x} + 2\bar{x} && \text{(A.5)} \\
&= -2x_0 - 1 + 2\bar{x} && \text{(A.6)} \\
&= 2(\bar{x} - x_0) - 1 \quad . && \text{(A.7)}
\end{aligned}$$

Thus

$$\bar{x} = x_0 + \frac{1}{2} + \sigma^2 \ln \frac{P_1}{P_0} \tag{A.8}$$

Use of this formula requires an estimate for σ, which cannot be obtained by analyzing the two data points.

A.2 Formulae for three data points

Let two data points of a gaussian distribution at ordinates x_0 and x_1 be represented by P_0 and P_1, respectively. Further, let $x_1 = x_0 + \Delta$, where Δ is a constant "mesh size." In this discussion, we will treat the ordinates as discrete integer variables, so that $x_1 = x_0 + 1$.

Let three consecutive data points in a gaussian distribution at ordinates x_0, x_1 and x_2 be represented by P_0, P_1, and P_2, respectively. Further, let $x_0 = x_1 - \Delta$ and $x_2 = x_1 + \Delta = x_0 + 2\Delta$, where Δ is a constant "mesh size." In this discussion, we will treat the ordinates as discrete integer variables, so that $x_0 = x_1 - 1$ and $x_2 = x_1 + 1$.

The quantities P_0 and P_1 are identical to their definitions in section A.1, and P_2 has a similar definition. All parameters (A, $\bar{x}$, and σ) have the same definitions as in section A.1.

An expression for $\bar{x}$ can be derived in a manner similar to that of section A.1, yielding the following result:

$$\bar{x} = x_1 + \frac{\sigma^2}{2} \ln \frac{P_2}{P_0} \quad . \tag{A.9}$$

With three data points to analyze, however, an expression for the width of the gaussian distribution (σ) can be derived.

Start by calculating $P_1 \times P_1$:

$$P_1 \times P_1 = A^2 e^{-\frac{(x_1-\bar{x})^2}{\sigma^2}} \tag{A.10}$$

Then

$$\ln\left(P_1 \times P_1\right) = 2\ln A - \frac{(x_1 - \bar{x})^2}{\sigma^2} \quad . \tag{A.11}$$

Similarly,

$$\ln\left(P_0 \times P_2\right) = 2\ln A - \frac{-(x_0 - \bar{x})^2 - (x_2 - \bar{x})^2}{2\sigma^2} \quad . \tag{A.12}$$

Thus,

$$\begin{aligned} \ln\left(\frac{P_1 \times P_1}{P_0 \times P_2}\right) &= \ln\left(P_1 \times P_1\right) - \ln\left(P_0 \times P_2\right) \\ &= \frac{(x_0 - \bar{x})^2 - 2(x_1 - \bar{x})^2 + (x_2 - \bar{x})^2}{2\sigma^2} \quad . \end{aligned} \tag{A.13}$$

Expanding the binomials and collecting terms yields

$$\ln\left(\frac{P_1 \times P_1}{P_0 \times P_2}\right) = \left[-2x_1^2 + 4\bar{x}x_1 + x_0^2 - 2\bar{x}x_0 + x_2^2 - 2\bar{x}x_2\right]/2\sigma^2 \quad . \tag{A.14}$$

Now replace x_1 with $x_0 + 1$ and x_2 with $x_0 + 2$, and expand and collect terms to obtain

$$\ln\left(\frac{P_1 \times P_1}{P_0 \times P_2}\right) = \frac{1}{\sigma^2} \quad . \tag{A.15}$$

Thus,

$$\sigma = \left[\ln\left(\frac{P_1^2}{P_0 P_1}\right)\right]^{-\frac{1}{2}} \quad . \tag{A.16}$$

Bibliography

[1] *Proposal for a Silicon Vertex Tracker (SVT) as an upgrade for the STAR experiment at RHIC* Berkeley, CA. : The STAR SVT Collaboration, (Feb, 1995)

[2] *Conceptual Design Report for the Solenoidal Tracker At RHIC.* Berkeley, CA.: The STAR Collaboration, (June 15, 1992).

[3] B. Mueller, Hadronic Matter in Collisions, World Scientific (1989), p.379

[4] J. Rafelski, *AIP Proceedings* 176 (1988) 1068

[5] J. Rafelski, *Phys. Rep.* 88 (1982) 331

[6] J. Rafelski, *Nucl. Phys.* A544 (1992) 279c.

[7] S. Margetis et al., STAR Note #72 (1992)

[8] C.M. Ko, Plenary Session, QM93, to be published.

[9] D. Lissauer and E. Shuryak, *Phys. Lett.* B253 (1991) 15.

[10] E.V. Shuryak, *Phys. Rev. Lett.* 68 (1992) 3270.

[11] K. Geiger, *Phys. Rev.* C48 (1993) 4129.

[12] R. Bellwied, STAR Note #108 (1993).

[13] T. Akesson et al., *Z. Phys.* C46 (1990) 361.

[14] H. Stroebele et al., *Z. Phys.* C38 (1989) 89.

[15] B. Jacak, *Nucl. Phys.* A525 (1991) 77c.

[16] E.V. Shuryak, *Nucl. Phys.* A525 (1991) 3c.

[17] V. Koch, Stony Brook preprint SUNY-NTG 94-26 (1994).

[18] L. van Hove, *Ann. Phys.* 192 (1989) 66.

[19] J. Bjorken, *Act. Phys. Pol.* B23 (1992) 637.

[20] G.J. Alner, *Phys. Rep.* 154 (1987) 247.

[21] K. Rajagopal et al., *Nucl. Phys.* B404 (1993) 577.

[22] S. Pratt, *Phys. Lett.* B301 (1993) 153.

[23] R. Bellwied, STAR Note #149 (1994).

[24] S. Pratt, *Phys. Rev. Lett.* 53 (1984) 1219.

[25] G. Bertsch et al., *Phys. Lett.* C37 (1988) 1896.

[26] C. Greiner et al., *Phys. Lett.* B219 (1989) 199.

[27] M. Gyulassy et al., *Phys. Rev.* C41 (1989) R21.

[28] A.N. Makhlin et al., *Z. Phys.* C39 (1988) 69.

[29] M. Gyulassy, Preprint LBL-32051 (1992).

[30] U. Heinz et al., *Phys. Rev. Lett.* 58 (1987) 2292.

[31] D. Keane, STAR note #47 (1992).

[32] S.S. Padula et al., *Phys. Lett.* B217 (1989) 181.

[33] T. Humanic et al., *Nucl. Instr. and Meth.* A316 (1992) 283

[34] E. Gatti and P. Rehak, "Semiconductor Drift Chamber - an application of a novel charge transport scheme," *Nucl. Inst. Meth.*, **225** (1984) 608-614

[35] S. M. Sze, *Physics of Semiconductor Devices* Wiley, New York City, 1981

[36] G. Hall, "Silicon Drift Chambers," *Nucl. Inst. Meth.*, **A273** (1988) 559-564

[37] E. Gatti *et al.*, "Dynamics of electrons in drift detectors," *Nucl. Inst. Meth.*, **A253** (1987) 393-399

[38] S. Roe *et al.*, "Temperature dependence of the behaviour of silicon drift detectors," *Nucl. Inst. Meth.*, **A310** (1991) 192-196

[39] C. Jacoboni *et al.*, "A review of some charge transport properties of silicon," *Solid-State Electronics,* **20** (1977) 77-89

[40] P. Rehak *et al.*, "Progress in semiconductor drift detectors," *Nucl. Inst. Meth.*, **A248** (1986) 367-378

[41] S. M. Sze, *Semiconductor Devices: Physics and Technology*, Wiley, New York City, 1985

[42] Press, William H *et al.*, *Numerical Recipes In C*, Cambridge University Press, Cambridge, 1988

[43] S. U. Pandey *et al.*, "Studies of Ionizing Radiation Effects on Silicon Drift Detectors," internal memo, STAR Collaboration, 1994

[44] V. L. Rykov, private commincation (1994)

[45] R. L. Gluckstern, “Uncertainties in track momentum and direction,” *Nucl. Inst. Meth.*, **24** (1963) 381-389

[46] Particle Data Group, “Review Of Particle Properties” Physical Review D, **50** (1994) p. 1253

Vita

David McArthur Read was born in Atlanta, Georgia on November 19, 1966, the son of Martha Orr Read and John McArthur Read jr. After completing his work at Stuyvesant High School in New York City in 1984, he enrolled in the University of Texas at Ausin. He received the degree of Bachelor of Science from the University of Texas in May 1988. In August 1998 he entered the Graduate School of the University of Texas.

Permanent address: 4008 Avenue D
Austin, Texas 78751

This dissertation was typeset with LaTeX[‡] by the author.

[‡] LaTeX is a document preparation system developed by Leslie Lamport as a special version of Donald Knuth's TeX Program.