STAR Note #321
January 15, 1998
Version 1.1

The Performance of the SVT-SDDs in AGS Experiment 896

R. Bellwied[a], H. Caines[e], H. Dyke[e], Y. Guo[a], J. Hall[a], J. Hoffmann[c], T. Humanic[e], P. Jensen[c], I. Kotov[e], P. Kuczewski[b], B. Leonhardt[b], C. Liaw[b], G. LoCurto[e], D. Lynn[b], N. Mazeh[a], P. Middelkamp[b], B. Minor[d], S. Nehmeh[a], G. Ott[c], S. Pandey[a], C. Pruneau[a], V. Rykov[a], J. Schambach[c], J. Seldmeir[b], J. Sheen[a], B. Soja[b], E. Sugarbaker[e], J. Takahashi[a], K. Wilson[a]

The SDDA Collaboration

[a] *Wayne State University*
[b] *Brookhaven National Laboratory*
[c] *University of Texas, Austin*
[d] *Lawerence Berkeley Laboratory*
[e] *The Ohio State University*

The performance of silicon drift detectors (SDDs) in AGS Experiment 896 is discussed. STAR-SVT prototype wafers, in particular the latest design wafers (STAR 2.7 and STAR 2.9), are used in the setup. The front-end electronics consists of preamplifier/shaper die (PASA) and switched capacitor array die (SCA) mounted on a hybrid (or multichip module) circuit. In the SVT system test requirements document this constitutes STAGE 0 of the system test.

Table of Contents

1. Experimental Setup
1.1 E896 Environment and Support Structure
1.2 Wafers
1.3 Front End Electronics
1.4 Readout and Data Acquisition
1.5 Slow Controls and Monitoring

2. January 1997 Beam Time
2.1 Beam, Target, Trigger, and Timing
2.2 Running Conditions
2.3 Accumulated Statistics

3. Detector Performance
3.1 Basic Performance and Characteristics
3.1.1 Channel-by-Channel Performance
3.1.2 Front End Electronics
3.1.3 Slow Controls
3.2 Pedestal and Noise
3.2.1 Noise Characteristics
3.2.2 Pedestal Subtraction
3.2.3 Common Mode Noise Subtraction
3.3 Cluster Analysis
3.3.1 Zero Supression and Cluster Locater
3.3.2 Cluster Analysis and Fitting
3.4 Results

4. Problem and Future
4.1 Data Processing Problems
4.2 Mechanical Structure
4.3 Future Improvements
4.3.1 Online Software
4.3.2 Data Processing
4.3.3 New E896 boards
4.3.4 Slow Controls and Monitoring

5. References

1. Experimental Setup

1.1. E896 Environment and Support Structure

The E896 experiment is designed to search for the H_o di-baryon and other strange particles in nucleus-nucleus (AA) collisions at the AGS [1]. One aspect of this experiment is the Silicon Drift Detector Array (SDDA), a set of silicon drift detectors (SDDs) placed near the target on the negative rigidity side of the beamline. The array consists of ten to fifteen SDDs, each placed about 1.5 cm from the beamline, spaced 1.5 cm apart. Each SDD is also rotated by 21 degrees in the vertical plane so that a majority of the particles hit the detector at a right angle. The first detector is positioned roughly 8 cm from the target, giving the device a pseudo-rapidity coverage of η = 0.8-1.8 and an azimuthal coverage of 20°.

The SDDA is placed in a small superconducting dipole magnet. This magnet produces a 7 T field in a 15cm high, 30cm wide, x 100cm long volume. The individual detectors are mounted on G10 board with front end electronics and placed in a water-cooled aluminum support structure fabricated for this purpose. In addition, the detectors are kept dry using a total-loss nitrogen gas system, which also aids the cooling.

STAR-SVT prototype wafers, in particular the latest design versions (STAR 2.7 and STAR 2.9), are used in this setup. The E896 experiment can therefore be considered as the first stage of the SVT system test. Two heavy ion beam times were originally planned. The second one will have a full setup of 10-15 wafers, whereas the one reported here only utilized two detector layers to accumulate statistics about the detector performance itself. In the second run (scheduled for early 98) we hope to obtain physics results (hyperon yields and ratios) with a more complete setup. In between the two beam times we will perform additional stages of the SVT system test as outlined in the SVT system test requirements document. These tests will be setup in the electronics laboratory and will utilize additional SVT prototype components, like a Beryllium ladder for the detectors and the electronics, a prototype of the SVT cooling system and prototypes of the electronics interface boards. The data reported here show the performance of a SVT prototype detector connected to the proper SVT frontend and the SVT readout electronics all the way to the DAQ. In the SVT system test requirements document this constitutes STAGE 0 of the system test. In the remainder of this chapter we will describe in detail the mechanical setup used during the beam time, the wafers chosen for the test, the frontend electronics and the readout electronics. Chapter 2 describes the beam time conditions, Chapter 3 shows the detector performance and some physics results, Chapter 4 details the problems that occurred and will have to be solved for the next beam time.

The aluminum support structure was designed and manufactured at Wayne State University. The assembly of the structure includes a base plate, three support posts, and an upper plate which is essentially a mirror image of the base plate. Both the upper and lower plates have fifteen grooves, cut at 21° with respect to the beam line,

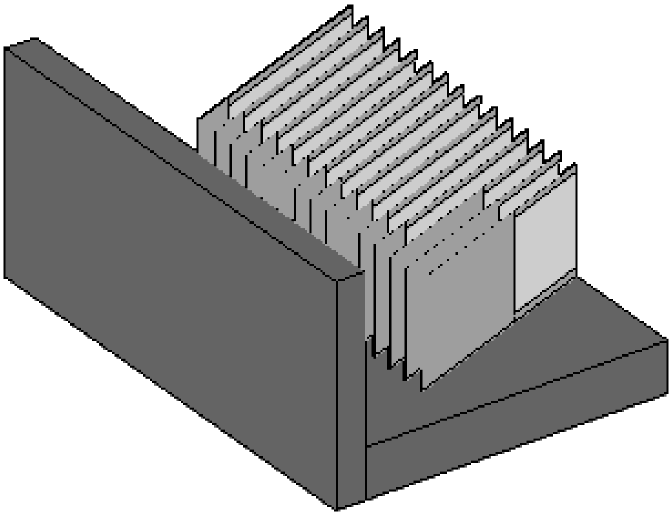

Fig. 1: Schematic of E896-SDDA Support Structure

designed to house up to two detectors each (as illustrated in Fig. 1). In both plates, a water channel winds its way between each of the grooves to provide cooling. The three support posts keep the two plates separated by 7cm. The structure as a whole measures 20cm across, 33cm long, and 14.5cm high. The dimensions of the structure were optimized to fit within the gap of the superconducting sweeper magnet in E896.

1.2. Wafers

The detectors used in the SDDA for the experiment E896 at the AGS were prototypes of the silicon drift detectors that will be used for the SVT in the STAR project. Two different prototypes were used in the experiment; STAR2.7 and STAR2.9. The detectors of the first type, STAR2.7 were produced by Sintef, a commercial company. The latter, STAR2.9, is the latest version designed and produced in the instrumentation division at Brookhaven National Laboratory.

The two different versions of the silicon drift detectors have the same functionality principles. The major differences between the two prototypes are in the charge collection region and in the HV guard line regions. Design changes in STAR 2.9 optimise charge collection through the anodes and a wider guard anode region minimises leakage current and possibilities of HV breakdowns. Further, the value of the implanted resistor between the cathodes was decreased from 1 $M\Omega$ to 0.5 $M\Omega$. The lower resistance value allows a stiffer resistor chain minimizing the effects of surface leakage currents on the voltage linearity. The two prototypes were used for comparison purposes.

Table 1 details the procedures used to test the detectors before mounting them on the experimental set-up. This same test procedure will be used to test the detectors for the SVT. In the following paragraphs, each test will be described in more detail. After these tests were completed, a few detectors were chosen to be mounted on PC boards and integrated with the corresponding electronics.

Type of test	test procedures
visual inspection	search for bridges causing electric shorts
	search for processing problems (e.g. mask misalignment)
Test Structures	resistor measurement
	diode current measurement
Probe Card	HV limit
	guard anode current
	anode leakage current
	voltage linearity

Table 1: Detector Test Procedures

The initial phase of testing is a visual inpsection of the wafers using a 50X optical microscope. Here it is possible to determine if the alignment of the processing masks had any major problems. One of the typical problems that can be found with the optical test is the existence of short bridges between cathodes. These could be caused by incomplete aluminum etching during the processing phase. With the recent improvements of the fabrication procedure we find that less than 2% of the produced detectors are eliminated in this phase of the test procedure.

The next phase involves characterizing the test structures that are built in around the detector on the same wafer. These test structures consist of small resistors series, diodes and P-N junctions. The resistors are designed to be identical to the implanted resistors of the detector. Using probes to electrically contact the test structures, it is possible to measure I-V curves and determine the resistor values. With this measurement it is possible to evaluate the doping implantation used during the fabrication process. In a similar way, the diode measurements can provide information on the P-N junction of the cathodes. The current measured through these diodes when reversed bias with 100V should be lower than $100pA$. A high current could indicate the existance of impurities on the junction surfaces. At this point it is possible to verify that there were no major problems during the fabrication process of the silicon wafers. In general, the detectors from the same batch of fabrication showed similar results.

The final step before mounting the detectors on to the PC-boards is the probe card testing. On the probe card station built at BNL, it is possible to apply voltages to the cathodes of the detector in the same voltage configuration that would be used in the experimental set-up. With this test it is possible to check if the detector can withstand the experimental operating conditions. Also a voltage linearity measurement is performed to detect possible drift non-linearity points and existence of surface leakage currents. Guard anode currents and HV current are measured as a function of the applied HV. In addition to these tests, it is possible to measure the anode leakage currents using another probe card. The anode leakage current is the final test that determines whether a detector is acceptable or not. The probe card testing was performed in a dark clean room as the silicon wafers are sensitive to light. Results from the probe card testing shows that the latest detector version, STAR2.9, has improved voltage linearity and lower guard anode currents. Also, the wider spacing between HV guard lines reduce the probability of break-down in the operational voltage range to a minimum. Table 2 summarizes the required values for each detector test sequence. The wafers chosen for the system test passed all the requirements.

Type of test		requirement
visual inspection	nr. of bridges	none
Test Structures	resistors	$0.5M\Omega$ +- 10
	diode current	$< 500pA$
Probe Card	voltage linearity	better than +-5V
	guard anode curr.	$< 50\mu A$
	HV currents	$< 5mA$
	anode leakage curr.	$< 150nA$

Table 2: Summary of Required Values for Detector Testing

A total of 10 silicon wafers including STAR2.7 and STAR2.9 prototypes were tested for use in experiment E896. All tested wafers passed the requirements summarized in Table 2. Two of each kind were selected to be mounted on PC-boards for the experimental set-up.

In parallel to the detector testing, the corresponding PC-boards were also tested for voltage linearity. An external resistor chain is used to bias the detector to maintain a linear voltage distribution. A single HV line is used to bias the resistor chain that will apply all the voltages necessary to power the cathodes and the focusing voltages.

1.3. Front End Electronics

The front-end electronics consists of preamplifier/shaper die (PASA) and switched capacitor array die (SCA) mounted on a hybrid (or multichip module) circuit. These were designed specifically for the STAR Silicon Vertex Tracker (SVT). The PASA is a 16 input/output channel charge amplifying and shaping integrated circuit designed in the BNL Instrumentation Group by D. DiMassimo. Each channel is intended to be (wire)bonded to a detector anode so that an incoming current pulse from the detector is amplified and shaped for both charge and time measurements. The shaped pulses are driven via the 16 output channels of the device. For an input delta function current, the PASA outputs a shaped pulse with a peaking time of approximately 100ns and maximum height of 7.2μV/e-. The width of this signal (from 10% of maximum to 100%) is 50ns. The response functions of the produced PASAs were explicitly measured to verify they meet the design requirements. All PASAs tested were found to have a similar response illustrated in Fig. 2.

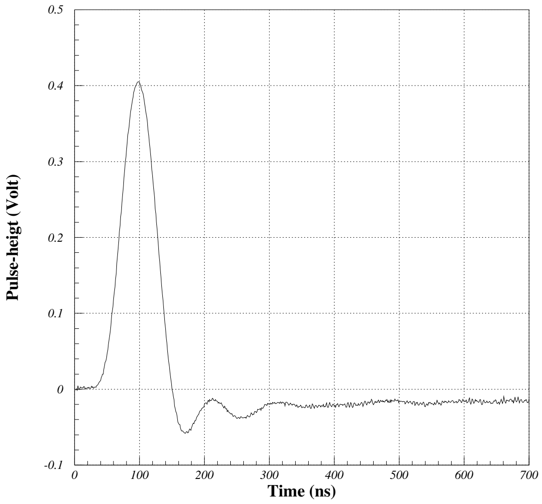

Fig. 2: PASA response function (measured on an oscilloscope).

The SCA is a 16 channel analog memory storage device designed at Lawrence Berkeley National Lab (LBNL) by Stuart Kleinfelder to accept signals from the PASA. Each SCA channel contains 128 sample and hold capacitors. At a typical sampling frequency of 25 MHz the output of each PASA channel is sequentially stored on the capacitors during “Acquire Mode,” when the SCA is sensitive to input signals. Upon receipt of a suitably delayed trigger the SCA is operated in “Readout Mode,” where the voltage on each capacitor is multiplexed onto a common SCA analog output line for digitization by the readout system.

The hybrid is a thick film circuit fabricated on a beryllia ceramic substrate. It was designed at BNL and fabricated by Emtron, a hybrid manufacturing company located on Long Island. Each hybrid contains 15 PASAs and 15 SCAs, capable of reading the 240 anodes of a half drift detector. Each group of five SCAs multiplex its analog output to a common analog buffer (Analog Devices Buf04) having higher drive capability. Thus each hybrid contains three analog buffers. Additionally, each hybrid contains a pseudo-ECL (PECL) level driver to fan out the differential clock signals, and a CMOS buffer to drive the four control signals required by the SCAs. Thick film resistors for biasing and surface mount capacitors for decoupling are also contained on the hybrid. Fig. 3 shows a photograph of a hybrid.

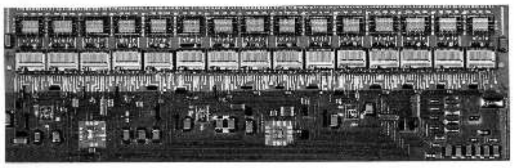

Fig. 3: Photograph of a fully assembled hybrid.

Two hybrids from the first prototype batch fabricated for the SVT were selected for use in E896. Each had known defects due to the manufacturing process. The first, FH4 (full hybrid #4) had a malfunctioning PASA (channels 81-96) and a dead SCA (channels 224-240). This is seen in Fig. 4 where the average pedestals and channel noise are plotted for FH4. The second (FH2) had three dead PASAs as seen in Fig. 5. At the time of the experiment only those two hybrids were available. Instead of instrumenting both sides (up and down) of a detector, it was chosen to instrument two detector halves. For the experiment we decided to mount FH2 to the STAR 2.7 prototype wafer, FH4 was mounted to the STAR 2.9 wafer.

Before each hybrid was bonded to the detector, we first studied their noise performance. As seen in Figs. 4 and 5, typical noise on the order of 4.5 mV (or about 625 electrons) is achieved. This is just above the design specifications of 500 electrons.

Fig. 6 shows the readout of one daisy chain (5 SCAs) of the hybrid. Charge is injected into a group of hybrid input pads by placing a conductor near the pads and applying a voltage pulse. The capacitance of conductor to several adjacent input pads results in the pulse spread in the channel direction seen in Fig. 6. Through this method we were able to verify whether each channel was functioning and measure the relative gain of each channel.

1.4. Hybrid-Detector Integration and tests

Once the detector and the hybrid were tested, both were integrated on a E896 PC board (SMD). Fig. 7 shows a photograph of the detector and the hybrid mounted to the SMD.

Six analog differential driver (ad826) stages and one digital differential line receiver (DS90C032) are surface mounted to the board. Input and output signals were

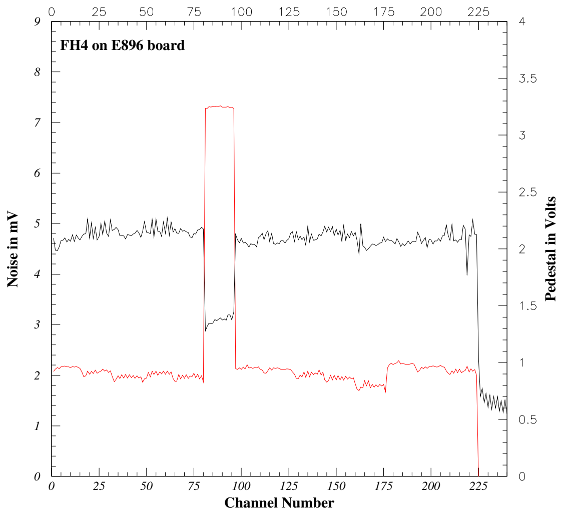

Fig. 4: Average pedestals(lower line) and noise (upper line) for each of the 240 channels of FH4.

transmitted through the 80-pin card edge connector. The high voltage needed to drift the charge in the detector was fanned out on the board to the different cathodes from a resistor network to obtain a constant electric field in the detector. The hybrid is fixed on the board via a heat conductive epoxy to facilitate the cooling. The detector cathodes were connected to the voltage divider on the board through wire-bonding. Ultrasonic wedge bonding technique was used with a 25 μm aluminum wire. The anodes are connected to the hybrid also through wire-bonding. Signal and power connections between the hybrid and the PC board were made with 34awg magnet wire.

The SMD permitted the first test of the readout of a SDD via the hybrid. The voltage linearity of the detector was remeasured using the resistor chain from the PC-board. Results were consistent with the measurements done previously using the probe-station.

The hybrid electronics was re-tested to verify that there was no damage due to the integration step. All results were consistent with the previous measurements, which indicates that the integration step is a relatively safe process.

Some capacitors (aproximately 100 μF) had to be added to the power line of the SCA on the PC-board to provide a stable bias voltage for the SCA. After the experiment, further optimizations of the SCA readout sequence reduced the need of capacitors on this power line.

The noise level of each channel was remeasured after the hybrid was bonded to the detector. Of particular concern was whether the additional input capacitance to

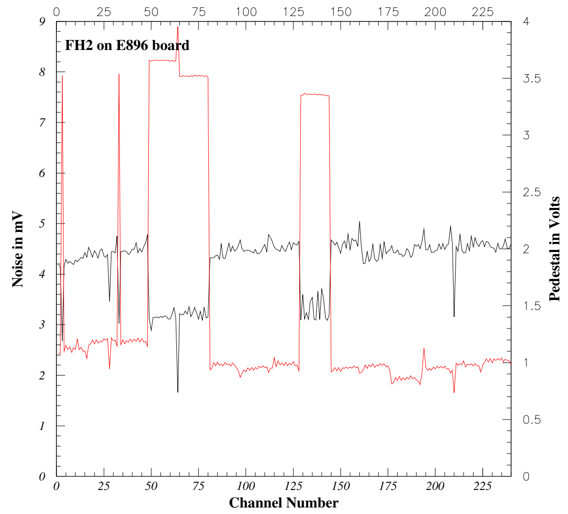

Fig. 5: Average pedestals (lower line) and noise (upper line) for each of the 240 channels of FH2.

the PASA from the detector anode and bond wire would result in a unnacceptable increase in noise. It was found that once that detector was biased with high voltage and a suitable settling time allowed (around 10 min.), the noise settled around 4.0mV to 4.5mV, indicating an increase of only 1.5mV due to coupling with the detector.

The final test performed before mounting the boards to the experimental set-up was a complete measurement of the signal generated on the detector through charge injection using the calibration lines and laser ionization. Using the hybrid testing readout system, it was possible to readout simultaneously the signal from 60 anodes at a time. Fig. 8 shows a two dimensional pixel plot of the signal generated by charge injection through the dot injection lines. The time axis shows the time distribution divided in 127 pixels of 39 ns. The drift time of the injection lines will provide the calibration of the drift velocity. Similar measurements were performed with charge injection using a infrared laser.

After finalizing the tests with the hybrid testing readout system, the detectors were connected to the readout and DAQ set-up of the experiment. In this case, it was possible to readout the complete hybrid with 240 anodes to look for dead anodes and/or complete dead PASA chips via laser injection. A more analysis of these results of these tests is presented in section 3.1.1.

1.5. Readout and Data Acquisition

The SDDA readout and data acquisition used in E896 is similar to that for the SVT in STAR [2]. This allowed us to test the components of the SVT DAQ in

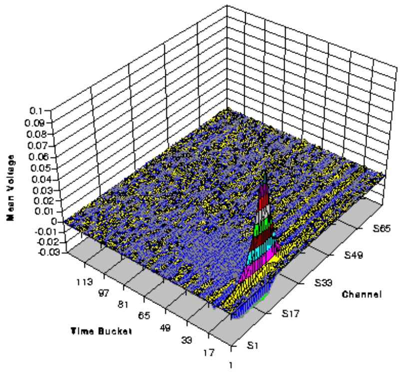

Fig. 6: Two dimensional mapping of output of 5 SCAs(one daisy chain) on the hybrid. The pulse seen on channels 15-40 is the result of capacitive charge injection to the input pads of the hybrid.

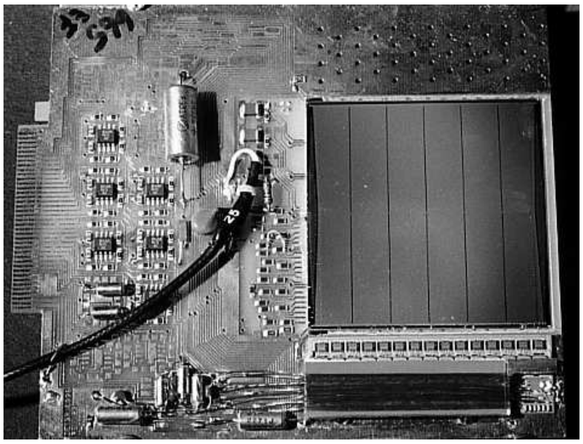

Fig. 7: Photograph of the detector mounted on the E896 PC-board integrated with the hybrid.

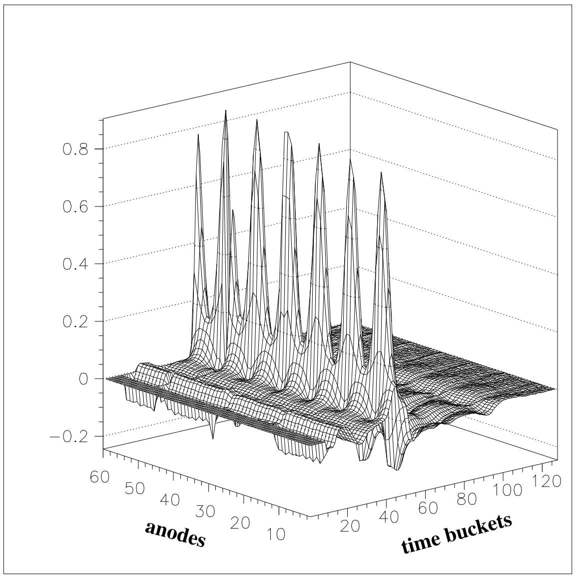

Fig. 8: Signal amplitude as a function of drift time and anode number. Signal measured corresponding to charge injection through dot implanted injector line.

E896, since both use the same detectors (SDDs). The basic SVT readout system accomodates 9 SDDs, and thus 18 hybrids. It consists of a readout board (RDO) and a VME receiver, both of which are STAR prototypes. To interface with the E896 trigger system, a trigger and clock distribution board was also developed.

The RDO is connected through a optical fiber to an optical receiver card, the prototype "ROSIE" card developed for STAR. Each ROSIE resdies on a Cyclone CVME964 card, which contains an Intel 80960 processor that assumes the role of one of the processors of the STAR receiver cards. In E896, these processors can perform zero suppression of the data. However, during the Winter 1997 beam time, this option was not used and all the data were stored. The layout of the readout and DAQ systems is summarized in Fig. 9.

During the E896 beam time, this DAQ was used to read out 2 hybrids on 2 SDDs through their associated electronics (see **Front End Electronics** above).

The data readout sequence jumps from SCA to SCA reading out the first time bucket before moving on to the second and so on. Each readout theoretically serves 18 hybrids (9 SDD's), therefore the data sequence in E896 for this run consists of 480 data words from time bucket 1 (from two hybrids), followed by 3840 empty words (from hybrids 3-18), followed by 480 data words from time bucket 2 and so on. The STAR-SVT ASIC which would allow zero-suppression on the processor level was not available yet. Therefore all data were recorded to tape. The data header was used to store the number of the first sample and hold capacitor that starts the data sequence

Fig. 9: Schematic of E896-SDDA Readout and DAQ

1.6. Slow Controls and Monitoring

During the E896 run in January 1997, the primary task of the SDDA slow controls was to monitor and control the high voltage power supply, a Le Croy 1454-HP mainframe, that was used to generate the drift field in the detectors. The computer used to accomplish this purpose was a Macintosh IIci with software written using Labview. These selections insured uniformity with the slow controls of the other components of the experiment.

The SDDA slow controls program enabled the user to change the drift field in the detectors by ramping up and down the detector's voltage at a safe rate. The voltages supplied and the currents drawn were monitored for each of the two detectors. Moreover, the user could monitor how each of these quantities evolved with time. Currents and voltages were measured every minute and were written on the Macintosh disk every 5 minutes.

The Macintosh was connected to the high voltage power supply through an RS-232 line running from the serial port of the computer to the serial port of the power supply. The distance between the two was about 100 feet. Although the ANSI/EIA RS-232 standard, which was established about 20 years ago, claims a maximum cable length of 50 feet at a data rate of 20 kbyte/s, it is currently safe to use RS-232 over larger distances. This can be attributed to subsequent improvements, both in the cables and in the drivers, such that RS-232 lines are now reported to work well over a distance of 200 feet at a data rate of 120 kbyte/s. However, our data transmission, over 100 feet of RS-232 line at a data rate of only 9600 byte/s, was carefully tested before it was adopted in the SDDA slow controls system.

2. January 1997 Beam Time

2.1. Beam, Target, Trigger, and Timing

In the January 1997 beam time, the experiment E896 used a gold beam with an energy of $11.3 GeV/A$. Two different gold targets were used with a thickness of 5% and 10% interaction length. Empty target events were taken from time to time for background subtraction purposes. The shape of the beam spot on the target was close to that of an ellipse with axis lengths of $1.5cm$ and $1cm$. Each target has a 2 cm diameter and the beam spot was centered in a way that no interactions with the target holder were likely to occur. The target wheel was mounted to a Geneva drive, allowing remote controlled target changes during the beam time.

The experiment used a central trigger (CT), a minimum bias trigger and a pulser trigger. However, the SDDA did have its own trigger, which in this run was chosen to be identical to the central trigger of the experiment. This trigger selected 20% of the most central events and was implemented as the coincidence of two beam counters and a multiplicity detector put in anticoincidence with a charge detector. In addition pulser data were taken with spill on and spill off for calibration purposes.

Cable delays and discriminator delays were taken into account to make the SDDA data acquisition start acquiring data $5\mu s$ after the primary interaction took place in the target. This timing would place the shortest drift hits in the first time bucket. The delay was varied during the beam time to observe the actual start time (t0) in the data and to verify that the t0 was indeed shifting across the detector. The delay time analysis is shown in the next paragraph.

2.2. Running Conditions

Throughout the E896 beam time the superconducting magnet was run at a current yielding a 4.7T field. The magnet is designed to provide up to 7T, but technical problems prevented us from running at full field. In a previous NIM paper [3] the SVT collaboration showed that the SDDs can be reliably operated at any field setting of the E896 sweeper. It was also shown that a 4.7T field leads to an effective decrease in the mobility of the drifting electron cloud which corresponds to an actual maximum drift voltage of 1070V instead of the nominal 1200V. The magnetic field was mapped after the beam time and shows the expected homogeneity.

During the beam time the detectors were cooled via the water cooling channels that had been grooved into the support structure. The SMD (E896 PC board) exhibits considerable warming if no cooling is provided. The cooled structure keeps the boards, the electronics, and the detectors themselves around room temperature. To control humidity and to ensure light tightness we enclosed the detector within a gas bag, which was continously flushed with nitrogen.

A Lecroy 1454 power supply was used to bias the detectors. Each SMD hosting one detector is connected through a single line to an output channel of the power supply. This power supply has the capability to be controlled remotely via RS-232 interface.

Due to the short time period available for testing the SMDs and the detectors and to eliminate any risk of voltage breakdown, the detectors were operated at a bias of 1200V instead of the 1500V for the most of the run. Under this condition, the active area of the detector reduces to 80% (2.4 cm maximum drift distance instead of 3 cm) due to the slower drift velocity. As a proof of principle the detectors were biased up to 1500V for 6 hours at the end of the data acquisition period.

The focusing voltages for the detector were set by the resistors of the external voltage divider. The exact voltage values were chosen based on extensive studies of the focusing region performed previously at WSU and OSU [4]. To keep the option of changing those voltages, it was possible to replace a few electric jumpers on the board and bias the focusing voltages independently. The voltage values used during the experimental run are listed in Table 3 below. Guard anode currents were measured

	STAR 2.7	STAR 2.9
E1 cathode	11.4V	10.9 V
W1 cathode	46.0 V	44.0 V
N0 cathode	30.5 V	28.8 V
P0 cathode	50.6 V	52.9 V
HV cathode	1200 V	1200 V

Table 3: Voltage Values used during January 1997 Beam Time

by the voltage drop on a 2 kΩ resistor mounted between the guard anodes and a ground line. Both detectors had guard anode currents higher than the acceptable range (below 150 μA), but still operational. The STAR 2.9 had a much lower guard anode current, below 10 μA, due to its improved design.

The HV current, supplied by the power supply was monitored during the whole experiment. The STAR 2.7 detector had a higher leakage current which also increased slightly (10%) through the time of the experiment. This increase was not observed on the new STAR 2.9 design.

For both detectors the sampling rate of the SCA was set to 25 MHz. To further investigate the accuracy of the sampling rate and to independently determine the drift velocity, we changed the delay time of the trigger to the detector (t0) and measured the effect on the STAR 2.7 wafer. Fig. 10 shows the response of the detector at three different delay times, t0 = 3.2, 4.2, and 6.2 μs. Almost all data were taken at t0 = 4.2 μs. Here the delay causes only the first four time buckets to be unusable. By comparing data taken with t0 = 3.2 and 4.2 μs we conclude that a time bucket corresponds to 41 $\pm$ 1 ns, very close to the expected value of 40 ns.

The steady fall-off in efficiency from a certain time bucket on when t0 = 4.2 or 6.2 μs is used can be explained as a loss in efficiency due to increasing signal width at long drifts. The far edge of the detector is not visible at any delay time due to

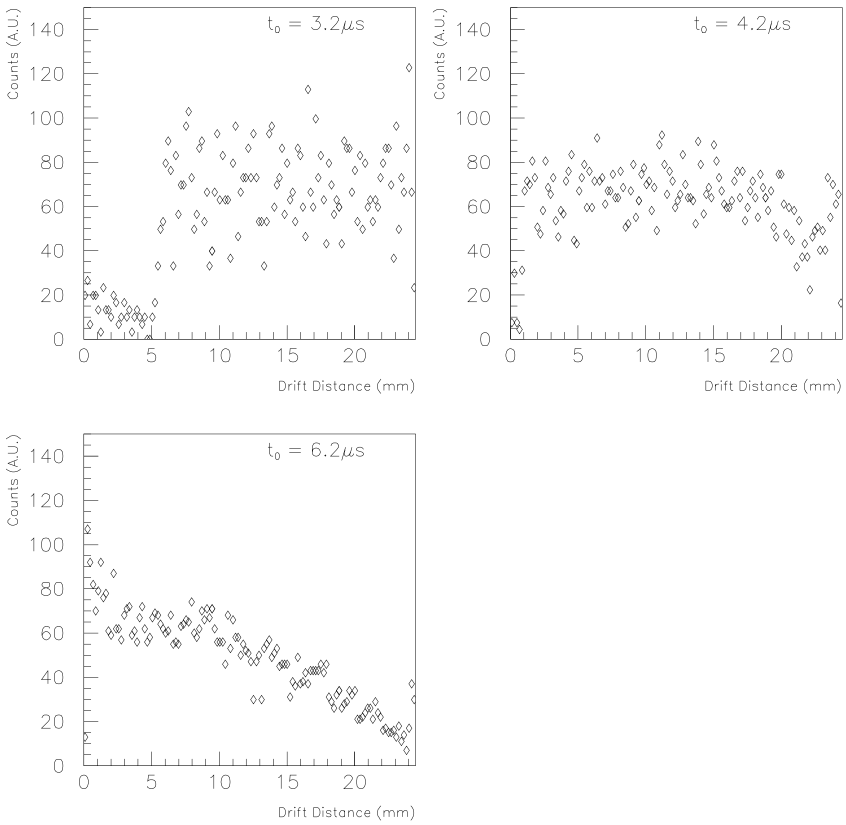

Fig. 10: Distribution of particle occupancies across the detector as a function of delay time.

the low maximum drift voltage. At 1070V instead of 1500V, the maximum drift (3 cm) corresponds to about time bucket 150. The 128 time buckets cover therefore only about 2.4 cm of the wafer (see above). Without measuring the far edge of the detector a direct measurement of the drift velocity is not possible. By treating the last few time buckets at the longest delay time as a far edge we obtain a drift velocity of about 4.8 μm/ns, which is in good agreement with simulations based on the maximum drift voltage of 1070V.

2.3. Accumulated Statistics

During the January 1997 beam time, we recorded over 65,000 events in 2 days. Due to readout electronics problems (see **4.2 Data Processing Problems**), only a third of this data is usable. Table 4 shows the accumulated statistics of this run by trigger and target type, as well as the delay time ("t0"). These three criteria were previously described above in Section 2.1.

Although the data acquisition systems for the SDDA and the remainder of E896 are independent, the detectors were connected via the shared trigger system. The trigger mix was set by the main E896 detectors.

Event Type	**Usable Events**
10% Central, t0=4543 ns	5168
5% Central, t0=4543 ns	962
Empty Target, t0=4543 ns	42
10% Central, t0=4200 ns	9936
5% Central, t0=4200 ns	359
Empty Target, t0=4200 ns	125
10% Central, t0=3200 ns	254
10% Central, t0=6200 ns	884
Pulser during spill-on	469
Pulser during spill-off	481
Laser pulse, varying fields	3542

Table 4: Accumulated Statistics

3. Detector Performance

3.1. Basic Performance and Characteristics

3.1.1. Channel-by-Channel Performance

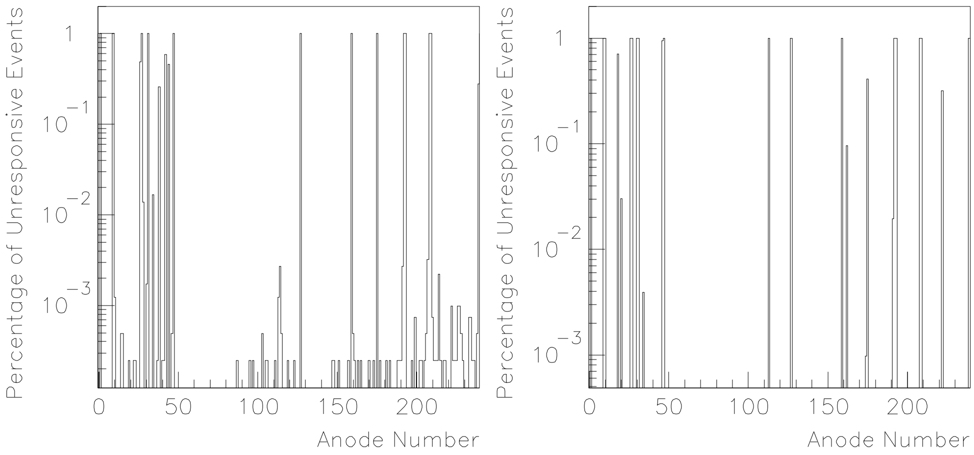

Fig. 11: STAR 2.7 anode distribution for unresponsive anodes. The left panel shows the unresponsive anodes in the events used for central-triggered data analysis, while the right panel shows the unresponsive anodes in the pulser-triggered data analysis. The unresponsive anodes may move from run to run. Any anode that is found to be unresponsive at least 50% of the time is removed from the analysis by setting the ADC values of the anode to zero.

Of the two detectors placed in the beam during the Winter 1997 beam time, one performed reasonably well. All of the results presented in this and the following section are from the analysis of the STAR 2.7 detector. The STAR 2.9 detector showed an uncharacteristic amount of noise, and was not used in this analysis (see **Problems and Future** below). The STAR 2.7 detector's single hybrid had three dead PASA chips (covering 48 channels) and two additional dead anodes. The three PASA chips were determined to be nonfunctioning before the experiment began. Dead channels are characterized by saturated ADC values. In addition, twenty of the anodes were characterized as *unresponsive*. An anode is characterized as unresponsive in an event where the ADC values recorded in 10 consecutive time buckets are identical. The percentage of events where an anode was found to be unresponsive as a function of anode number is shown in Fig. 11. Anodes that are unresponsive more than 50% of the time are considered dead and all of their time bucket ADC values are set to zero for the analysis. Unresponsive anodes change from run to run.

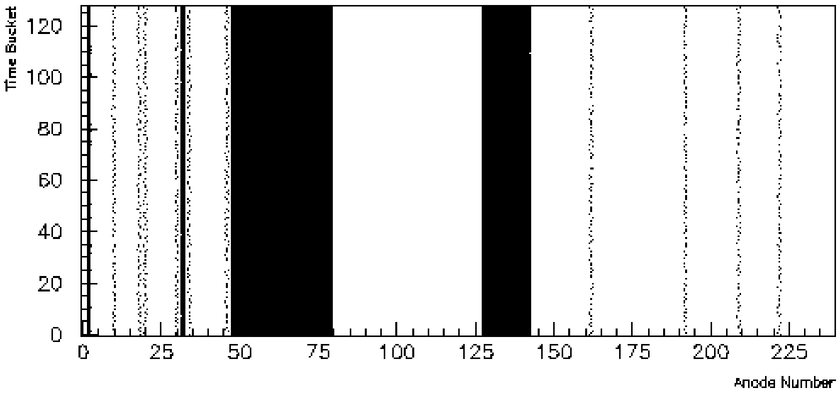

Fig. 12: STAR 2.7 anode distribution. Darkened areas represent dead or unresponsive channels. The beam is closest to anode number 239.

The distribution of dead and unresponsive channels across the detector is shown in Fig. 12. Darkened areas represent the channels for which all of the time buckets were set to zero for the analysis. About 5% of the channels besides the non-functional PASAs are unresponsive. This number is above the SVT requirement of 2%.

3.1.2. Front End Electronics

The front end electronics performed well during the experiment. In particular, the newly developed bipolar PASA chip performed as expected. The Laplace transform of its theoretical current impulse response function is:

$$F_{pasa}(s) \propto \frac{s}{(s+1/\tau_s)^5 \cdot (s+1/\tau_l)}; \tag{1}$$

where τ_l and τ_s are two time constants that depend on the details of the PASA. In the present version of the SVT PASA their approximate values are: $\tau_l \approx 500ns$ and $\tau_s \approx 11.5ns$. The signal from the PASA will appear as the convolution of the input gaussian signal with the PASA response function. Figure 13 shows a typical fit of a convoluted fit function to a signal shape obtained during the beam time. This cluster represents a charge of about 10 times minimum ionizing, producing a pronounced undershoot. This signal is not characteristic for the majority of clusters in the detector, which are mostly due to minimum ionizing particles. The signal was chosen because the clean shape allows a good comparison with the PASA response function. There is good agreement between the theoretical shape and the signal.

3.1.3. Slow Controls

The slow controls performed well during the time it was connected (about 6 hours). It was disconnected after a crash of the high voltage power supply apparently caused by a communication problem (this question will be discussed further in section 4.5.4).

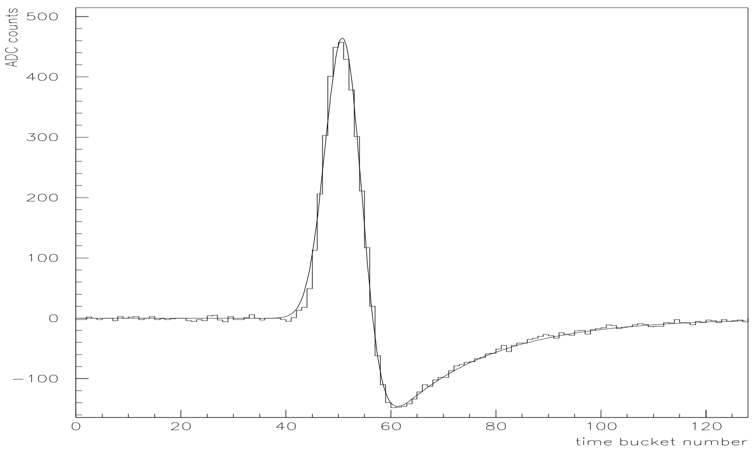

Fig. 13: Fit of the signal shape from one anode with the convolution of the PASA response function and a gaussian. The histogram represents the data and the line represents the fit.

While connected it recorded reliable information of the condition of the detector drift field and current.

Figure 14 shows data obtained by the slow controls program. It shows how the current drawn by one of the detectors varies as a function of the magnetic field. There was an overall increase of about 20% in the current while the field changed from $0T$ to $4.7T$.

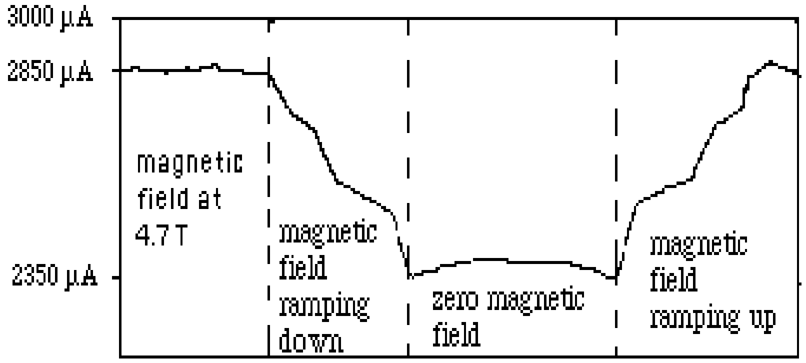

Fig. 14: The current drawn from the high voltage power supply as a function of the magnetic field

3.2. Pedestal and Noise

An important aspect of detector peformance is the analysis of its pedestal and noise values. More advanced analysis, like finding clusters from charged particle hits, cannot be attempted until these values are adequately characterized and subtracted.

Here, the methods of pedestal calculation and subtraction and common mode noise suppression are studied.

3.2.1. Pedestal Subtraction

The STAR 2.7 detector exhibited an RMS noise of less than 5 mV on the bench (2.5 ADC channels). In the experimental environment the raw signals picked up additional noise which was subsequently removed in the analysis. In this section we will describe the pedestal subtraction/noise suppression procedure.

Since each anode is read out by its own PASA, it is necessary to calculate a separate pedestal for at least each anode. We decided to go one step further and calculate a separate pedestal for each "pixel" (a single time bucket in the readout of a single anode). There are $240\times128=30720$ pixels in the data output from a single hybrid. Breaking the pedestal down into individual pixels allowed us to search for any variations in the pedestal over the course of the readout cycle.

Note that, since the SCA is treated as a circular buffer, the relationship between time bucket number (time since trigger) and SCA capacitor number (position on the SCA) changes from event to event. The pedestals reported here are only sensitive to the time bucket number; variations between individual SCA capacitors average out over many events. The preliminary version of the DAQ used for this run did not supply information on the alignment of the time bucket number and the capacitor number. In future studies this information will be available and we will examine the pedestal as a function of SCA capacitor number as well.

The time period in which a pedestal result is assumed to be valid depends on the time scale of pedestal changes. In Fig. 15 we show the ADC value of a single pixel for each event in a run. This indicates that the ADC value for a single pixel is stable over the course of a run. A run corresponds to about 2 hours of beam time. Therefore, we chose to perform one pedestal evaluation per run and will refer to these pedestals as "run pedestals." The pedestal is free to vary over longer time periods in response to temperature or other environmental changes.

The run pedestal for a specific pixel was taken as the average of that pixel's ADC values over the course of the run. A pixel may receive a minimum ionizing particle hit during the run, but this occurs infrequently and the ADC values are too small to bias the average. Very large pixel ADC values resulting from hits by heavily ionizing particles are eliminated from the sample before calculating the average. The resulting pedestal values for the pixels of one run are shown in the top panel of Fig. 16. Typically, the run pedestal values are near 450 channels (700 mV). This offset value is rather high and we will lower the offset in the next run by changing the FEE configuration. The peaks at 0 and 1024 channels are due to dead PASAs or bad individual channels.

The bottom panel of Fig. 16 shows the RMS deviation of each pixel from its pedestal (average), calculated over the whole run. This is a rough measure of the noise of the detector. The most probable noise value was about 5 channels (10 mV),

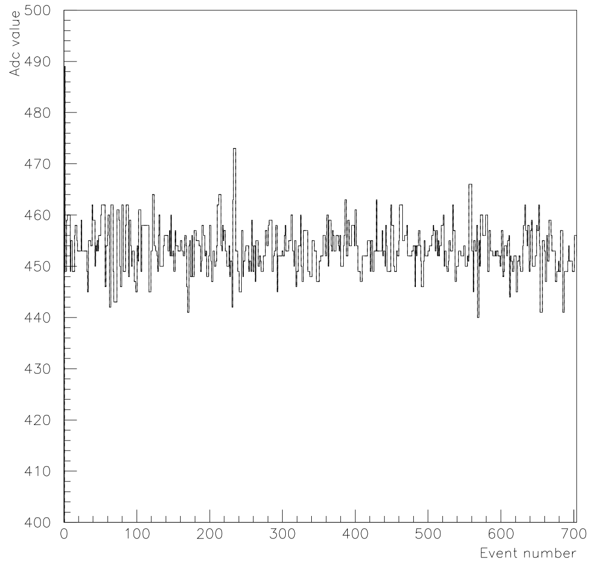

Fig. 15: The ADC values of a single pixel are shown for each event in a run.

a factor of two larger than the noise observed in the bench tests.

As expected, the run pedestal levels varied from PASA to PASA. The pedestal levels were also found to vary in an odd/even fashion with anode number across the anodes within a PASA. This pattern is apparent in Fig. 17 which shows the run pedestal variation across the anodes for one time bucket. Relatively large pedestal variations at the edges of the 16 channel PASAs are also apparent. (The edges of a PASA are the first and last channels.) None of these pedestal characteristics degrade the performance of the detector since they are all removed during the run pedestal subtraction.

The run pedestals exhibit a damped oscillation over the course of the readout cycle. This is apparent in Fig. 18 which shows the pedestals for several anodes as a function of time bucket. The magnitude of the damped oscillation is too small to be noticed in the data of a single event, but it becomes apparent when averaging over hundreds of events. Since it shows up in the run pedestal, the phase of the damped oscillation must be synchronized in time with the trigger. We believe that the damped oscillation is due to a transient created during the triggering of the SCA. It could be eliminated by increasing the delay between the filling and readout of the SCA from 1 μs to 10 μs. Since it is part of the pedestal values it is removed from the data and does not degrade the performance of the detector.

3.2.2. Common Mode Noise Suppression

The run pedestal for the appropriate run is subtracted from the raw data early in the analysis of each event. After this subtraction it was observed that an additional

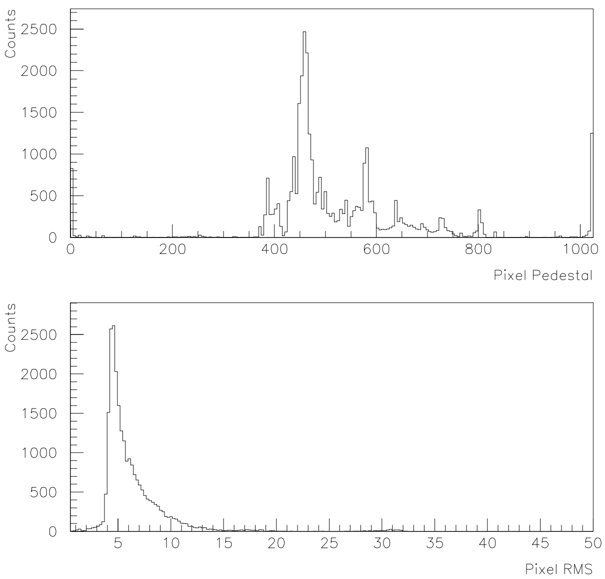

Fig. 16: The top panel shows the pedestal values in channels for each pixel in a run. The bottom panel shows the distribution of RMS deviations of each pixel from its run averaged value.

oscillation was present. This oscillation is characterized by approximately the same phase and amplitude on every anode in an event. The top panel in Fig. 19 shows a scatter plot of the ADC values against drift time (y axis) and anode number (x axis). The large irregular horizontal bars across the anodes indicate the presence of correlated noise pick-up common to many anodes.

Since it is common to all the anodes we refer to it as common mode noise. It does not appear in the run pedestals because the phase of the noise with respect to the SDDA trigger changes from event to event. This behavior suggests that the SDDA is picking up these oscillations from its environment.

The magnitude of the noise is indicated by the data shown in Fig. 20. The run pedestal subtracted ADC values are displayed for each time bucket for two events. For each time bucket, the ADC values were averaged across all of the anodes. In essence, the data consist of projections of the two-dimensional data in the top panel of Fig. 19 onto the time bucket axis, divided by the number of good anodes. The phase variations of the oscillations from event to event can be seen by comparing the event in the top panel with the event in the bottom panel. The amplitude of the oscillation also varies, but it can be as large as several channels. We concluded that this common mode noise dominated the RMS noise presented earlier. The sudden drop and rise in the ADC value at time bucket numbers 35 and 102 in the top and bottom panels respectively indicate the first and last capacitors of the SCA. This became clear in subsequent tests in which the position of the edges of the SCA was

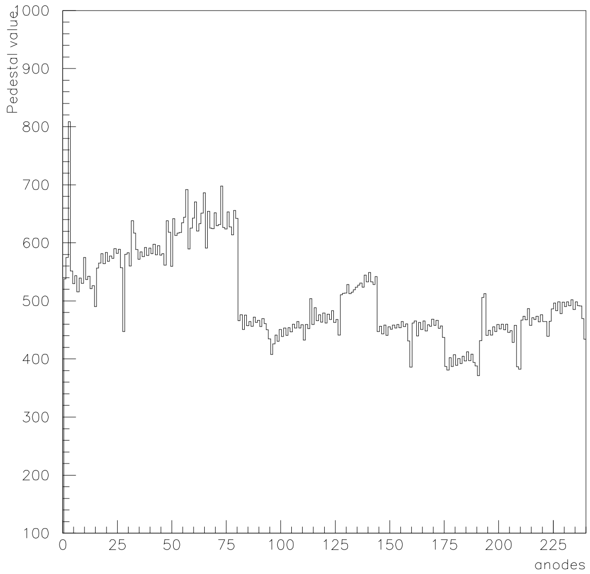

Fig. 17: Pedestal values for one time bucket plotted as a function of anode number.

added to the data stream.

Fortunately, the data shown in Fig. 20 may be employed to remove the common mode noise from the data. First, the pedestal is subtracted from the data. Next, the anode averaged time sequence is calculated as described above. We will refer to this as the "event pedestal." This is then subtracted from each individual anode, removing the common mode noise. Limits are placed on the ADC values allowed to contribute to the event pedestal in order to minimize the impact of actual particle hits on the common mode noise spectrum. Any drift in the pedestal average during a run will also appear in the event pedestal and be removed from the data by this technique. This procedure leads to a dramatic reduction in the common mode noise. The data displayed in the bottom panel of Fig. 19 resulted from the application of this technique to the data in the top panel. Note that the clusters appearing at longer drift times (later time buckets) are much more clearly defined after common mode noise suppression.

We can measure the noise in a single event by calculating the RMS deviation of the pixels in an event from the event ADC average. The average should be near zero following pedestal subtraction. This is a more accurate measure of the intrinsic noise of the detector electronics than the one shown earlier since the effect of event to event fluctuations in the overall pedestal level is removed. The results of this measure for each event in a run is shown before and after the common mode noise suppression in the top and bottom panels of Fig. 21. After subtraction of the run and event pedestals, the RMS noise has dropped to approximately 3 channels (6 mV), roughly

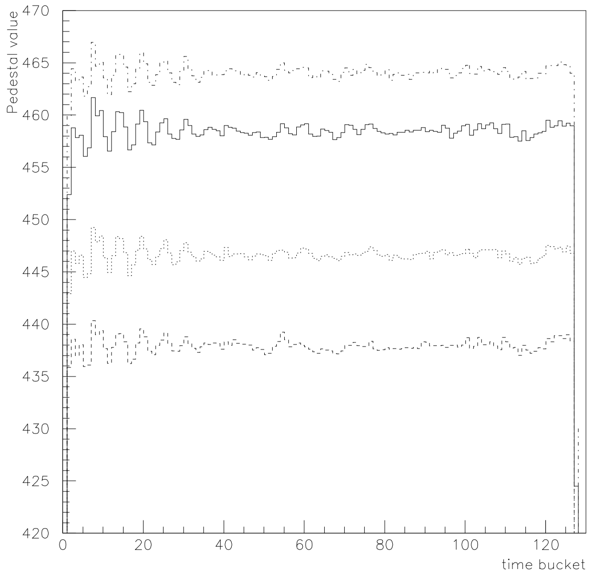

Fig. 18: Pedestals for several anodes as plotted as a function of time bucket.

equivalent to the values seen on the bench.

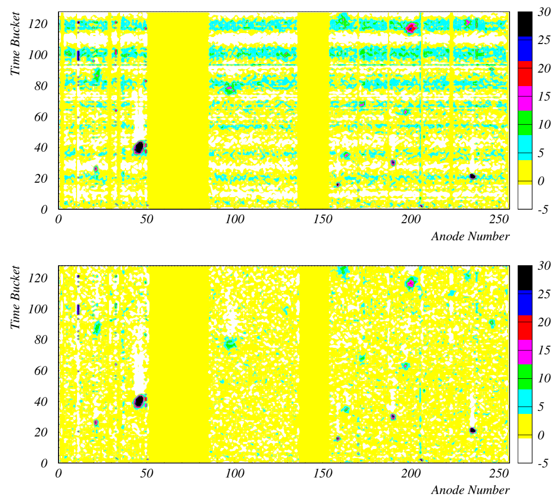

Fig. 19: ADC values (grey scale) are plotted as a function of time bucket and anode number for a single event after pedestal subtraction. The top panel shows the data before common mode noise suppression, the bottom panel shows the data after noise suppression.

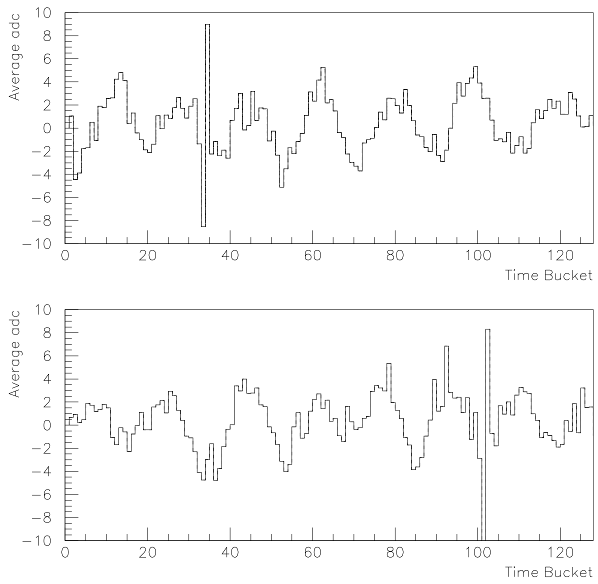

Fig. 20: ADC values after run pedestal subtraction are shown as a function of time buckets. The top and bottom panels contain data from two separate events in the same run. For each time bucket, the values were averaged across all of the anodes.

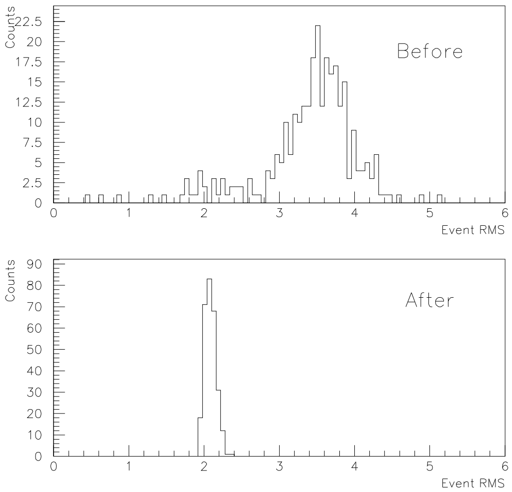

Fig. 21: The distribution of RMS deviations of pixel ADC values from their event average values. The top panel and bottom panels contain data from a single run before and after common mode noise suppression.

3.3. Cluster Analysis

3.3.1. Zero Supression and Cluster Locater

A. SDDA Zero Supression (SZS). Once the pedestal and common mode noise have been subtracted, the data are fed to the SZS. The zero suppression routine can be thought of as a rudimentary cluster finder. The algorithm requires five user-supplied parameters:

High Threshold: The SZS scans each pixel in an event, looking for pixels over the *High Threshold*. Once it has found a candidate, it proceeds further.

Low Threshold: Once the SZS has found a pixel that surpasses the *High Threshold*, it looks in the eight neighbors around this candidate pixel to see how many are above the *Low Threshold*.

Minimum Neighbors: It then compares the number of neighbors above the *Low Threshold* to the parameter *Minimum Neighbors*. If the number of neighbors equals or exceeds this parameter, the candidate pixel and additional neighbors will be saved for analysis using the Cluster Analyzer.

Anode and Bucket Buffers: To insure that the entire waveform is collected, additional pixels around the candidate pixel are saved. Any pixel within *Anode Buffer* anodes and *Bucket Buffer* buckets will also be saved.

In the data analyzed below, the *High Threshold* was set at 6 channels, the *Low Threshold* at 5 channels, the *Minimum Neighbors* to 4, the *Anode Buffer* to 2 and the *Bucket Buffer* parameter to 4. The *High Threshold* of 6 was chosen because it represents an ADC value just more than twice that of the common mode noise RMS. The buffer parameters were chosen to be asymmetric because the clusters are expected to be asymmetric. The lower threshold and the minimum number of neighbors were chosen to optimize the efficiency for finding good clusters and rejecting false ones. Fig. 22 shows a typical minimum ionizing hit at a short drift distance. Fig. 23 shows the same hit, with each number representing the ADC value of a pixel. These figures illustrate how the SZS works.

In the final step, non-zero pixels are assembled in a list fed to the SDDA Cluster Locater.

B. SDDA Cluster Locater (SCL). The time bucket sequences from the SZS are fed into the SVT Cluster Locater package as developed by David Read [5]. The SCL tries to assemble these sequences into clusters of two dimensions. In the simplest case, one cluster represents a single hit on the detector. However, there are cases where two hits may strike the detector close enough such that the charge clouds will overlap as they drift to the anodes. These clusters are said to contain merged hits, and occur with a low frequency (1 every 10 events) due to the low occupancy of the SDDA. In this analysis, the occupancy is small and cluster merging is neglected, i.e. all clusters are treated as single clusters.

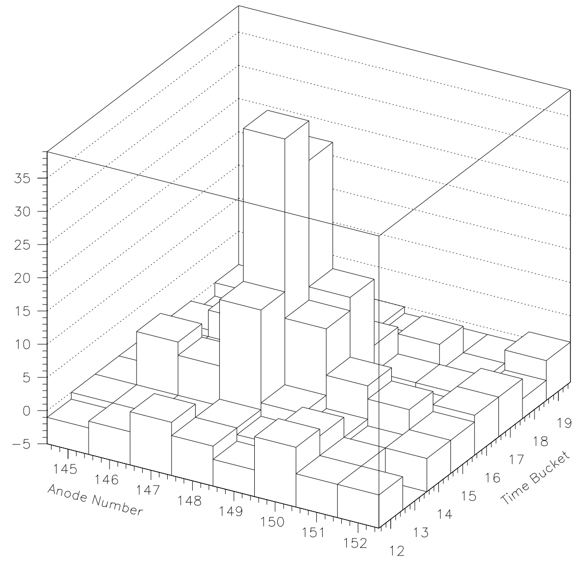

Fig. 22: Lego representation of a typical minimum-ionizing cluster

To construct a cluster from these sequences, the SCL begins with one sequence on an anode and calls it a new cluster. Adjacent anodes are examined for sequences. If one is found "nearby" [5], it becomes part of this new cluster. The search continues to adjacent anodes until no more sequences are found[1]. The data are now organized by cluster, and sent to the SDDA Cluster Fitter package for further analysis.

3.3.2. Cluster Analysis and Fitting

The SDDA Cluster Analysis and Fitter (SCAF) package is designed to give information about the position, size, and ADC sum of each cluster. There are two approaches to this problem. The first proceeds through a simple moment analysis while the second is based on a detailed fit of the signal shape. While the first method is certainly more time efficient, the latter provides a better description of the waveform.

A. Moment Analysis. Moment analysis has been used to evaluate characteristics of a cluster in a silicon drift detector before [5,6]. Let A_i represent the anode position of a pixel in a cluster, and B_j its time bucket position. An estimate of the particle's ADC channel sum Q is obtained using the zeroth moment of the distribution:

[1] In a sense, this has already been accomplished by the SZS for single hits. This analysis path was used to maintain a greater similarity to the SVT analysis path.

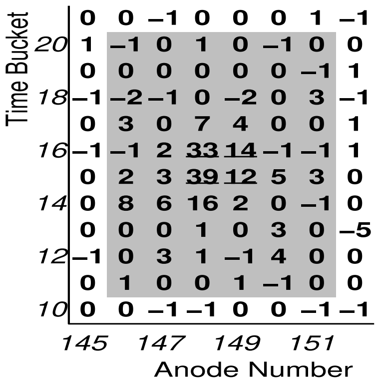

Fig. 23: Numerical representation of the same minimum-ionizing cluster as in Fig. 22. The shaded ADC values are kept by the SZS. The underlined pixels passed the SZS cuts.

$$Q = \sum_{i,j} u_{ij}, \; u_{ij} \geq u_{min}, \tag{2}$$

where u_{ij} is the ADC value in the pixel located by anode A_i and time bucket B_j and u_{min} is a threshold used to eliminate unwanted entries into the sum, such as the undershoot created by the PASA response function. The particle's hit coordinates in both the transverse A and drift B directions can be estimated using the first moments of the distribution:

$$A \simeq m_1 = \frac{\sum_{i,j} A_i \cdot u_{ij}}{\sum_{i,j} u_{ij}}, \; u_{ij} \geq u_{min} \tag{3}$$

$$B \simeq m_2 = \frac{\sum_{i,j} B_j \cdot u_{ij}}{\sum_{i,j} u_{ij}}, \; u_{ij} \geq u_{min}, \tag{4}$$

These first moments are also known as the centroids of the distribution. The second moments,

$$m_{11} = \frac{\sum_{i,j} A_i^2 \cdot u_{ij}}{\sum_{i,j} u_{ij}} - m_1^2, \; u_{ij} \geq u_{min} \tag{5}$$

$$m_{12} = \frac{\sum_{i,j} A_i \cdot B_i \cdot u_{ij}}{\sum_{i,j} u_{ij}} - m_1 \cdot m_2, \; u_{ij} \geq u_{min} \tag{6}$$

$$m_{22} = \frac{\sum_{i,j} B_i^2 \cdot u_{ij}}{\sum_{i,j} u_{ij}} - m_2^2, \; u_{ij} \geq u_{min} \tag{7}$$

provide some information regarding the electron cloud size upon its arrival to the anodes. If the u_{min} is set to zero or below (as it is in this analysis), anomalous pixels at the edge of clusters can significantly alter the calculation of these moments. Readers will recognize the second moments m_{11} and m_{22} as the RMS of the distribution in the transverse and drift directions, respectively.

In addition, the pixel with the largest ADC value within each cluster is located and recorded. This knowledge, along with the moment analysis, provides some information for comparison to simulation [6].

B. Cluster Fitting. Some of the clusters found by the cluster finder were subsequently fit in both the transverse and drift directions via separate one dimensional fits. No integration over anodes was performed while fitting in the drift direction nor integration over time buckets while fitting in the transverse direction. The two perpendicular axes of the cloud used for fitting were selected to be the longest in both the transverse and the drift directions. The fitting procedure in the drift direction has already been discussed in Section 3.1.2.

The signal in the anode direction was fitted both with a gaussian and with a gaussian integrated over the anode bin size which is given by the anode pitch. Only small differences were observed in the fitted signal width.

Results are summarized in Fig. 24. Six hits with an integrated charge greater than 1500 and smaller than 9000 ADC counts were studied. These correspond to integrated signals ranging from 3V to 18V. Note that a 3 V signal should be equivalent to eight times the signal of a minimum-ionizing signal. That means the minimum cluster in this analysis is about 8 times minimum-ionizing. The main reason to use these very rare large signals is that their charge spreads over many channels, therefore providing a signal on a number of channels large enough to perform a very reliable fit. The dimensions of the electron cloud generated by these signals is as large as 15 time buckets in the drift direction, and 9 anodes in the transverse direction.

As already shown in Fig. 13 the convolution of the PASA response function with a gaussian agrees well with the experimental signals. This agreement however has been reached by treating the two time constants, τ_{long} and τ_{short}, as free parameters. It should be noted that a variation in the time constants can produce a shift in the position of the signal peak. This is the reason why the fitted position along the drift is slightly different from the results of the moment analysis (Fig. 24, left panel). In particular the drift position obtained from fits is systematically shifted to lower values compared to moment analysis results. In fact the PASA introduces some delay in the amplified signal that is taken in account in the PASA response function, but not in the moment analysis.

3.4. Results

We focus on the cluster and waveform analysis to determine the functionality of the frontend as well as the wafer themselves. Results of the drift velocity and the signal distribution in drift and anode direction are compared to parametrized detector

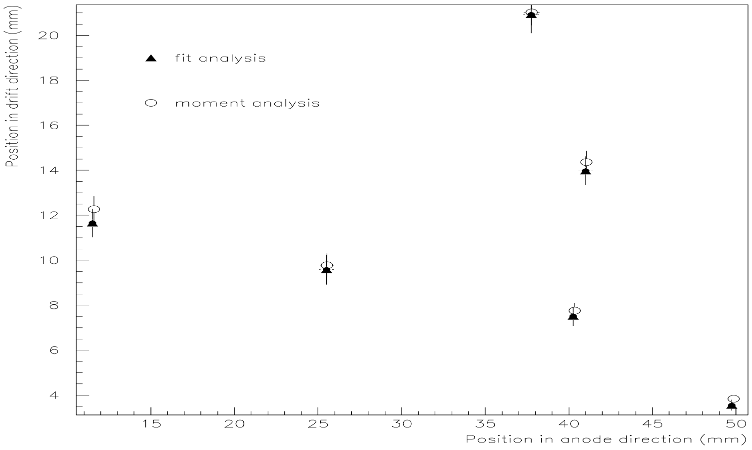

Fig. 24: Comparison of the results for hit distribution for 6 selected hits according to the fit analysis and the moment analysis. There is good agreement between the two methods.

simulations as presented in STAR Note #171. [7]

The analysis uses 4,270 central-triggered events from 5% and 10% interaction targets with the 4.2μs delay time and 889 pulser-triggered events with the same delay time. In Fig. 25, a comparison of central- and pulser-triggered events are presented. The pulser-triggered event clusters are not due to RMS noise of 4mV presented in Section 3.2.2, but are most likely generated from a combination of undetected warm anodes (see Section 3.1.1) and correlated noise (see Section 3.2.2). The SZS and SCAF algorithm parameters have been optimized to eliminate most of these false clusters. Fortunately, these false clusters peak at a summed ADC of less than 75mV, and can be cut later in the analysis. Moreover, since these false clusters taper quickly as one approaches a sum of 120 mV, a cut on this sum above 120 mV should insure good purity without sacrificing too much efficiency. Fig. 25(d) shows a comparison between relative purity and efficiency as a function of the threshold.

The RMS values of the distribution, namely the second moments m_{11} and m_{22}, are also calculated. The results of these calculations are shown in Fig. 26(a+b). In both panels the expansion of the electron cloud is obvious in both the transverse and drift directions as the drift time increases. Again, the lack of data at very short drift times (first four buckets = 160 ns = 770 μm) is intentional and due to the 4.2μs delay time. At long drift distances, the sharp reduction of the RMS in the drift direction is due to hits near the edge of the detector which are not completely stored in the SCA.

We can also get an estimate of the detector performance by examining how the highest pixel (amplitude) in the cluster varies with increasing drift distance. As the electron cloud drifts, Coulomb repulsion and diffusion expand the cluster, so that the cluster amplitude decreases with distance. Hits closest to the anode should show a larger pixel amplitude than hits at longer drift distances. Fig. 27(a) shows how the data behave as a function of drift distance. All charges are plotted here. Data shown in Fig. 27(a) are used to estimate the signal peak amplitude reduction versus drift distance for minimum-ionizing particles shown in Fig. 27(b). Data presented in (a) are "sliced" along the drift axis in 1.5mm bins. The most probable peak amplitude found in each slice is shown in Fig. 27(b) as solid squares. Error bars show an estimate of the RMS width (sigma) of the peak amplitude distribution.

By using the cluster finder and the moment analysis we determine the summed charge of each cluster and plot it as a function of drift distance, see Fig. 28(a). The data shown include a correction based on an effect due to the PASA response function. For increasing drift distance, the cluster size increases in the drift directions, as shown in Fig. 26. The leading edge of the signal produces an undershoot which reduces the amount of charge recorded at the tail end of the cluster. To correct for this effect, the summed ADC value of a cluster is multiplied by the factor $(1 + 0.2\sqrt{m_{22}^2 - 0.7^2})$. The most probable values of this corrected summed charge are shown in Fig. 28(b). The error bars represent an estimate of the rms width (sigma) of the summed charge distribution in each time slice. The result shows that the summed charge is independent of drift distance for a large part of the detector. The slight increase seen at short distances may be caused by a bias in the cluster finding. It is expected that the amount of charge is independent of the drift distance.

The vertical axis of the summed charge distribution shown in Fig. 28 is related to the number of electrons liberated by the passing of a charged particle. This relation involves a deconvolution of the PASA response function and is quite complex. However, an estimate of the number of electrons can be calculated using the following method. Since the number of electrons is independent of position (as shown in Fig. 28(b)), we can assume that a hit at the anode would produce the same number of electrons. These electrons do not require any time to drift, and so would be collected in a single pixel. Since the RMS noise is calculated to be about 6mV (see Section 3.2), our signal-to-noise ratio on a single pixel would be $350mV/6mV = 60/1$. This distribution is in agreement with that expected for minimum-ionizing particles, which was confirmed through bench tests utilizing a radioactive source.

This corrected summed ADC spectrum for all drift distances is shown in Fig. 29. As expected, this spectrum is consistent with a Landau distribution.

The centroids, or first moments of the cluster distributions, along the anode direction (Fig. 30(a)) and the drift direction (Fig. 30(b)) are shown. The number of clusters has been normalized to multiplicity per event. We conclude that the detector is hit by about 5 particles per event. The trigger setting corresponds to roughly the upper 20% of the geometrical cross section. A comparison with an ARC distribution (solid curve) with the proper centrality cut (upper 20% cross section corresponds to a impact parameter cut of b $\leq$ 6 fm) shows that the model predicts a much larger multiplicity per event. The average ARC multiplicty is 13 particles per event. We conclude that the disagreement with the model might be due to either uncertainties in the centrality selection of the trigger or the accuracy of the model prediction itself.

We expect the efficiency near the anodes to be 100%. Fig. 30(c) shows, irrespective of the normalization, that in the E896 run the efficiency decreased as a function of drift time. This efficiency drop is primarily due to electron cloud expansion which decreases the ADC value of any one pixel. These clusters are much harder to identify, as their amplitude could be as small as twice the noise RMS. Fig. 30(d) shows how the efficiency depends on the amplitude of the peak of the cluster. The data, represented by open circles, show that the efficiency drops dramatically around 30mV. This statement is confirmed by bench tests using a STAR 2.9 detector, as shown in the darkened circles. This second curve is an actual efficiency, not a comparison with a model. Software ineffiencies for high amplitude clusters are expected to be on the order of 1%.

In future runs, we expect 100% efficiency across the active area of the detector by:

1. Using the STAR 2.9 detector instead of the STAR 2.7. The STAR 2.9 detectors show higher peak amplitudes at long drifts due to a redesign of the Anode-E1 junction.
2. Drifting at higher voltages. This reduces the drift time and therefore increases the peak amplitude by decreasing cloud expansion.

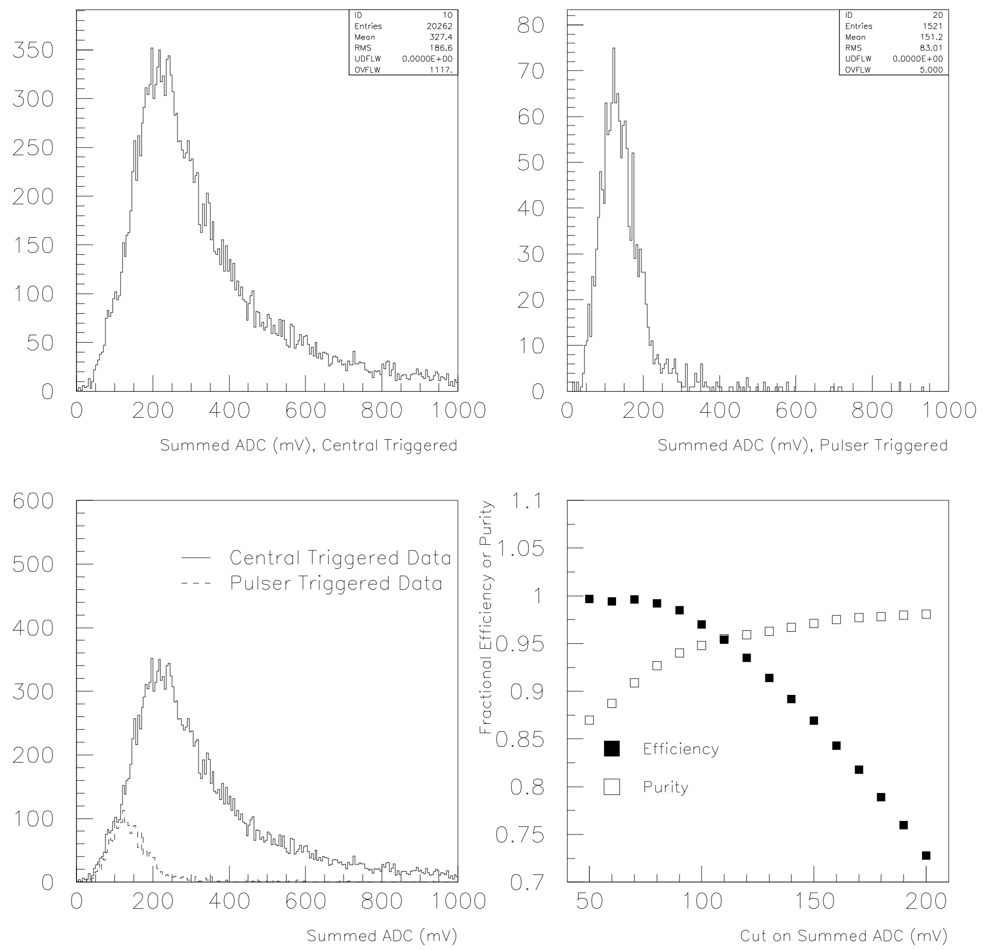

Fig. 25: Summed ADC Value for (a)Central and (b)Pulser Triggered Data. (c) For comparison the pulser triggered data has been scaled to the same number of events as the central triggered data. (d) Relative purity and efficiency as a function of noise threshold cut based on the pulser data.

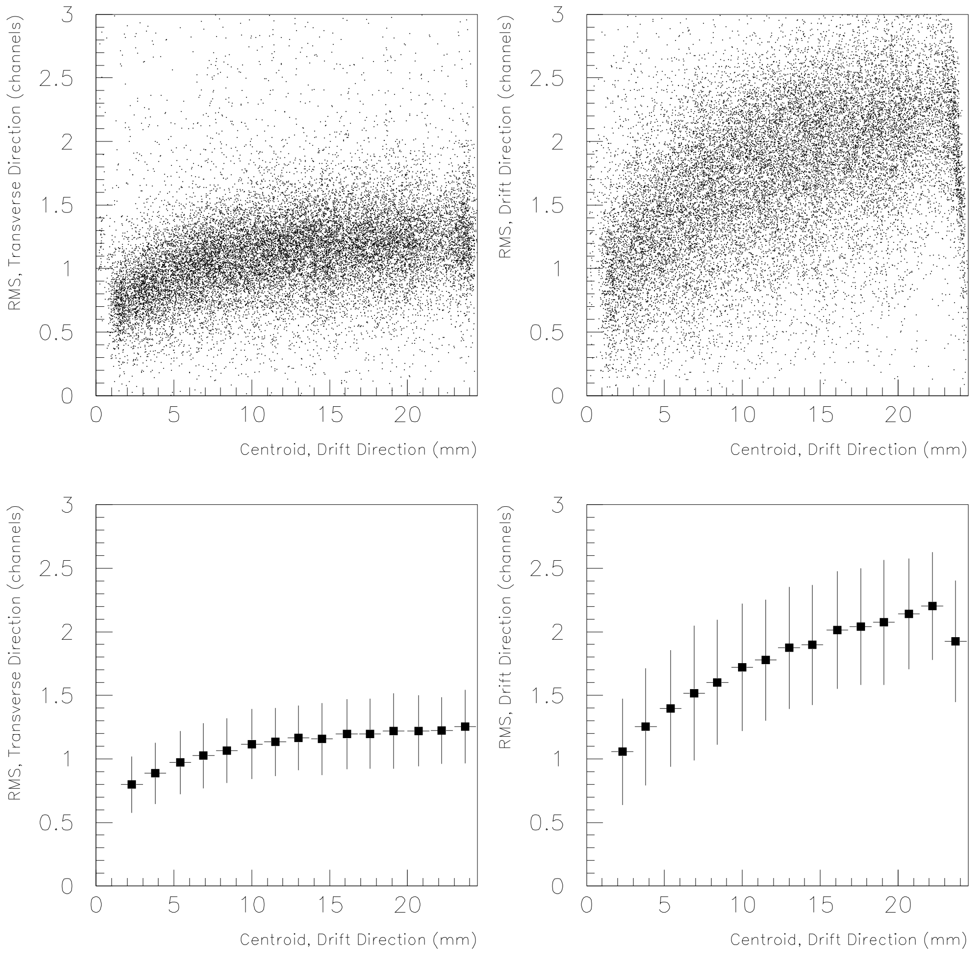

Fig. 26: (a+b) RMS Values of the Clusters in Tranverse and Drift Directions as a function of Drift Distance. A cut has been made on the summed ADC value of the cluster (see text). (c+d) The average RMS value of the data for sixteen slices in drift direction are presented. The error bars represent an estimate of the sigma of the distribution.

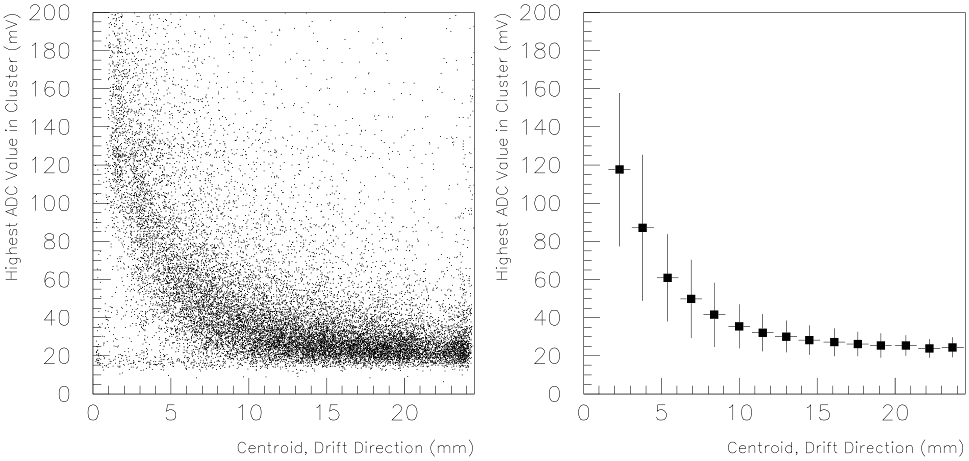

Fig. 27: (a) Highest ADC Value (amplitude) in a cluster as a function of drift distance. The SZS provides a cut on the highest ADC value at 12 mV. (b) The highest ADC value (amplitude) in a cluster now gated on minimum-ionizing particles. The error bars represent an estimate of the sigma of the distribution in each time slice.

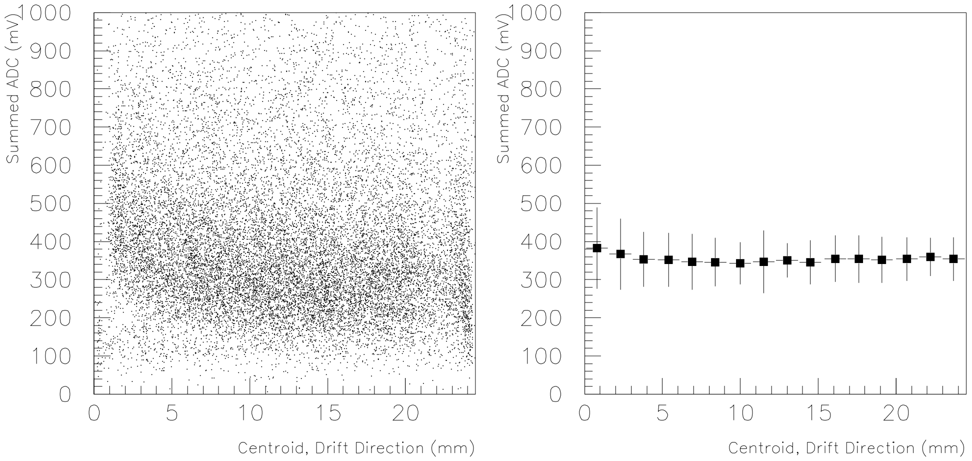

Fig. 28: (a) Summed ADC values (total charge) as a function of drift distance for all reconstructed hits. (b) Truncated means of the distribution in (a) in sixteen time slices, corrected for the PASA response function (see text).

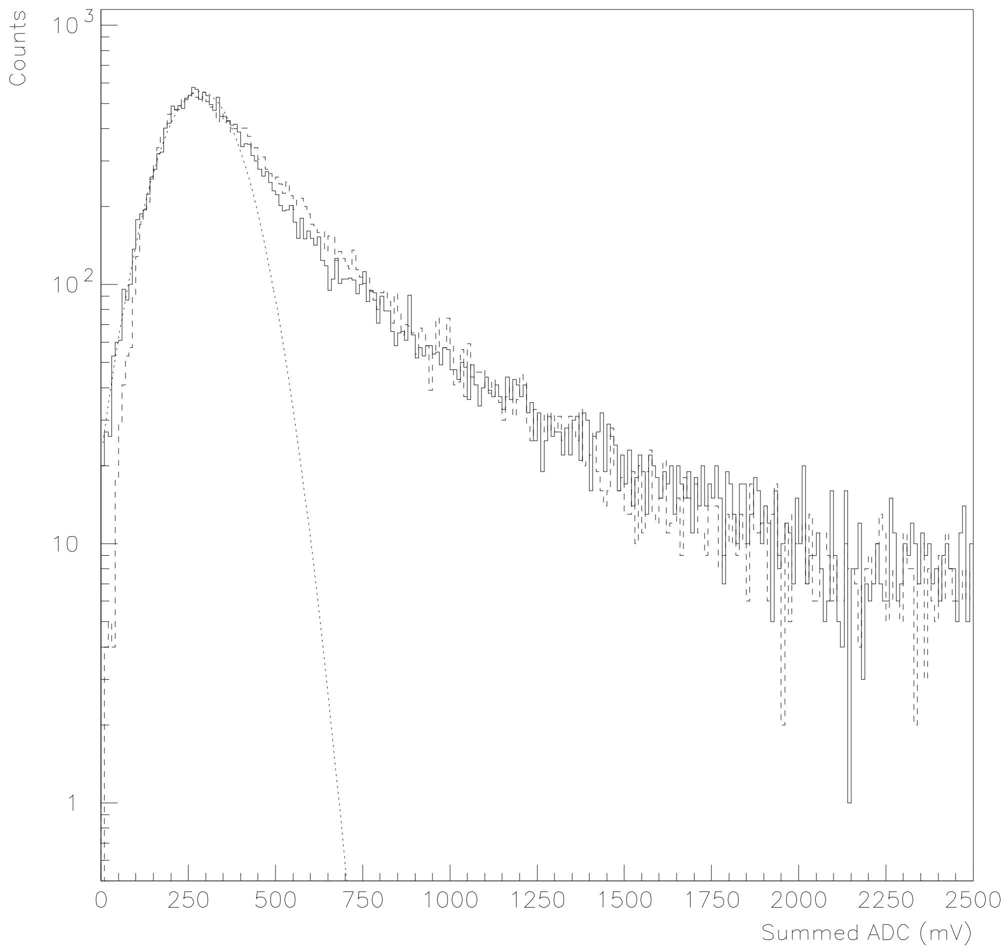

Fig. 29: Summed ADC spectrum for all drift distances corrected for the effects of the PASA response function. The data is represented by the solid line, a Landau distribution fitted to it is represented by the dashed line. A Gaussian distribution that fits the peak of the data, represented by the dotted curve, is shown for compairson.

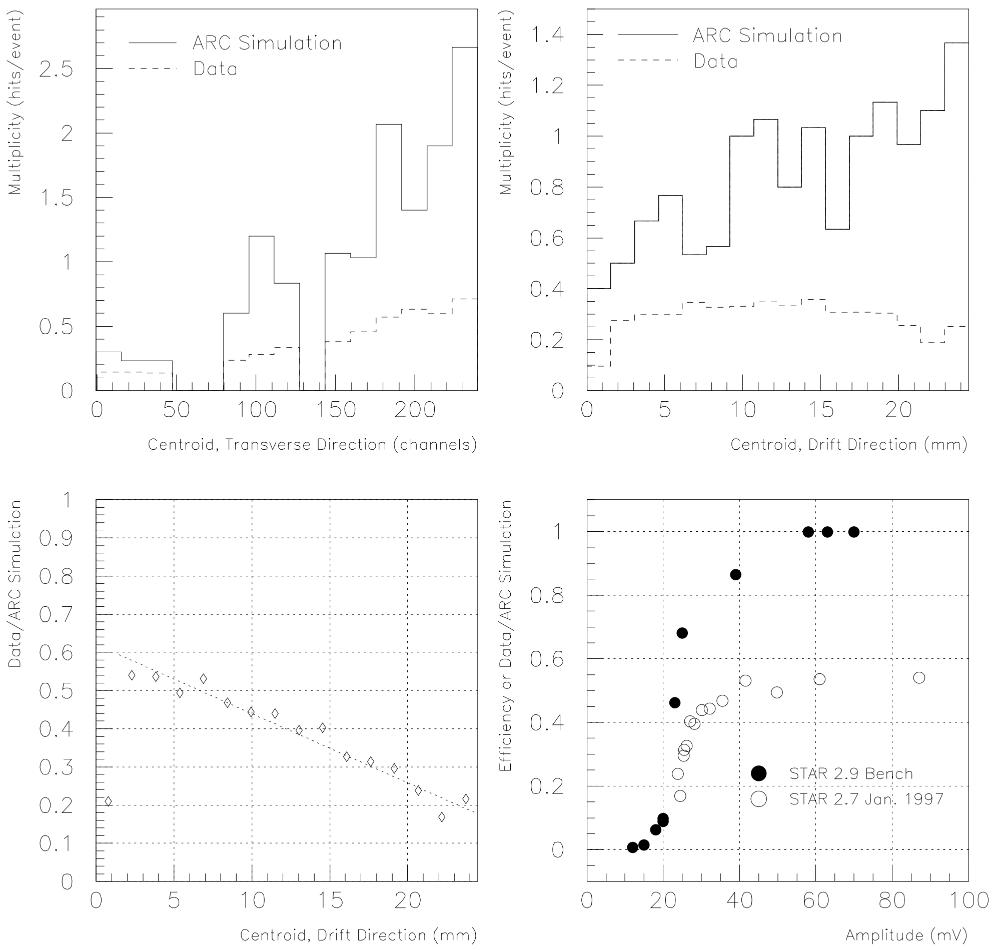

Fig. 30: (a+b) Centroids of the Clusters in Tranverse and Drift Directions. A cut has been made on the summed ADC value of the cluster (see text). In both panels, the dashed line represents the data while the solid line is the occupancy reported by the same number of ARC events run through GEANT. (c) estimates of detector efficiency based on comparison between ARC and data. (d) STAR 2.7 and STAR 2.9 detector efficiency as a function of amplitude. The STAR 2.9 bench test data was taken after the run. Both detectors show a marked decrease in efficiency as the amplitude of the signal drops below 30mV.

4. Problems and Future

4.1. Data Processing Problems

There were some data processing problems during the January 1997 beam time. The first problem involved the interface between the E896 trigger system and the VME processor (see Fig. 9). When an event is accepted, a busy signal issued by the VME control board is used to latch the SCA while they are read out. Triggers received during the conversion should be ignored. In practice, the busy signal was often missing and the SCA were re-triggered while being read out, resulting in an usable mix of events.

That means the call from the ROSIE board to the processor did not occur because ROSIE did not receive an "event-end." Instead a second trigger was issued during the collection of the first event. This problem results in events in which, after a certain time bucket, all of the ADC values stored are identical. This effect occurs in all the anodes simultaneously. Fortunately, because the values are identical, it is simple to write an algorithm which is able to determine if the data are being repeated. Because of this problem, two-thirds of the data stored are unusable.

A certain part of the events did not show this feature because a busy reset was hardwired into the logic to reset the busy every three seconds.

In addition, there were several other minor problems. In some of the data, the last two bits had been swapped. The first event of any run was unreliable and not used. The capacitor number of the first time bucket was stored, but was not reliable, and thus no pedestal subtractions by individual capacitor in the SCA could be performed.

4.2. Mechanical Structure

The SDDA support structure also had its share of difficulties. The water channels used for cooling were created by milling out grooves from the bottom side of the base plate, in between the detector grooves located on the top (and vice versa for the top plate). Additional aluminum pieces were added to the sides and bottom to close the channel. Because of the concern for the front end electronics, a strong sealant (Loctite Gasket Eliminator #518) was used liberally when attaching these pieces. Unfortunately, as the pieces were pushed into place, the excess sealant went directly into the water channels, especially in places where channels narrowed considerably at the crossover points between detector grooves. This material molds quickly, and large globs of the substance broke off inside, blocking the channel. This caused an intermittant problem for the cooling, which was partially solved "on the fly" by reversing the flow periodly to loosen the material within.

In addition, the final design of the PC boards used to hold the detectors was made wider after the construction of the structure was completed. To deal with this situation, the detector grooves were widened, but not completely through. The grooves were made tight to the dimensions of the PC board, sanded and cleaned before installation[2]. However, this resulted in difficulty placing the PC boards in the grooves before the experiment. In the future, these grooves will be widened even further to insure a smooth fit.

4.3. Future Improvements

4.3.1. Online Software

Before the next beam time in spring 1998, there will be some development in the online software. While the software used to monitor detector health was sufficient for this beam time, there are minor improvements to be made. It must be able to handle a greater number of detectors, check for problems such as repeating data, and subtract common mode noise on an event-by-event basis.

In the next run, more detectors (10-15) are expected to be used in the experiment. There will not only be a need to measure each detector's response, but also investigate things like alignment, tracking, and most of all, physics! It is expected that an online visualizing tracking package will be available for the next run. In addition, it is hoped that this tracking package can be combined with a vertex finding routine so that Λ, Ξ, and H_o particle candidates can be clearly seen during the experiment.

4.3.2. Data Processing

Most of the data processing improvements involve fixing the problems that occured during the Winter 1997 beam time.

The solution to the repeat problem is two fold. First, the ROSIE board has a timeout feature which can be synchronized to the proper event-end and can then be used as an error interrupt if false triggers should occur. Second, the reason for the large amount of false triggers was that the trigger board generated a level trigger, which is very susceptible to noise, rather than an edge trigger, which is planned for the next beam time.

In addition, it may be possible to move the test trigger card to the VME rack to facilitate the hold off. The SVT is supposed to have long PECL cables from trigger and those drivers are being used. Header information, including the capacitor number of the first time bucket and the DDC event number, has already been added.

Also, the use of a zero suppression algorithm in the Intel 80960 processors will be investigated to see if this leads to an enhanced event rate. The run control will be updated to include a more user-friendly interface. Finally, a dual tape drive system may be implemented if necessary to increase the amount of data stored during the run.

4.3.3. New E896 boards

The new E896 detector board has been modified to fit only one detector and two hybrids. The 80 pin card edge connector has been replaced by a 50 pin header type

[2]To insure good alignment, and thus good resolution, the grooves need to be made tight so that the PC boards will not move during the run.

connector which is more convenient from a geometric point of view.

The low voltage bias lines V1, V2 and V3 have decoupling capacitance added of $94\mu F, 94\mu F$ and $282\mu F$ respectively.

Pads are now provided for the eight injectors. Bias resistors ($1M\Omega$) and coupling capacitors ($1000pF$) connect to each pad. The coupling capacitors are connected directly together (on the interface side) and go to pin 23 (injector) of the header connector. The bias resistors connect to cathodes at a more negative potential.

Various changes were made in the board milling process to accomodate wire bonding access to the p side of the wafer and to assure no high voltage breakdown.

4.3.4. Slow Controls and Monitoring

The slow control that operated this year is only one piece of the more complete slow control that will operate in the next E896 run. Moreover, a Le Croy 1458 power supply mainframe will be used instead than the 1454 that was used for the January 1997 run.

The problem that occurred during this run, the crash of the 1454 high voltage power supply, has been interpreted by Le Croy to be generated by an unstable state of the mainframe. This problem, should it also be associated with the 1458, can be solved by checking, possibly with Labview, the state of the power supply before the beginning of the run.

The tasks of the full version of the slow control will be to control the high voltages and the focusing voltages, and to measure them and their currents, and moreover to measure the low voltages that supply power to the FEE and the temperatures and guard anode currents of each detector. The temperatures will be transformed into voltages by thermistors standing on the detector boards, and the voltages will be multiplexed and measured by an ADC. The guard anode currents will be transformed into voltages by a $100k\Omega$ resistor, then multiplexed and measured by an ADC. In figure 31 is shown a schematic of the slow control.

The program will measure all the data every minute, and will write them on the Macintosh disk and on the database every 10 minutes, or whenever a warning occurs. A warning would be issued whenever a measured quantity passes either the lower threshold or the upper threshold. Lower thresholds and upper thresholds are written on a file generated before starting the operation. The Macintosh is connected, through an ethernet line, to the database, which resides on the disk of a Silicon Graphics computer . The ADC and the multiplexer are controlled by the computer through GPIB lines.

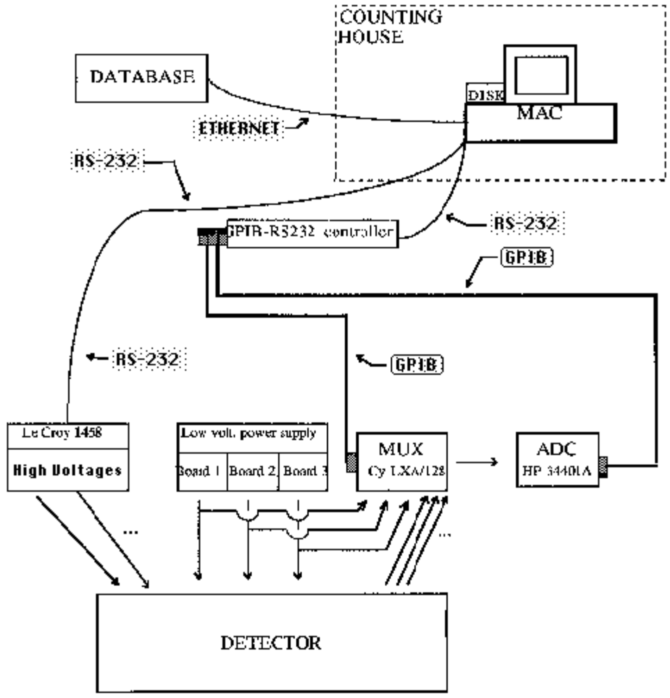

Fig. 31: A schematic of the slow control of the SDDA detector for the experiment E896.

5. References

[1] H. Crawford and the E896 Collaboration, *A Proposal to the BNL AGS: To Search for a Short-Lived H_o Di-Baryon, Short-Lived Strange Matter, and to Investigate Hyperon Production in 11.6 A GeV/c Au-Au Collisions*

[2] J. Schambach et al., *Data Acquisition Systems for E896 at the AGS*

[3] S. U. Pandey *et al.*, *Transport Properties of Electrons in Silicon Drift Detectors Measured in Large Magnetic Fields*, NIM in PRA, 383 (1996) 537-546.

[4] R. Bellwied *et al.*, *Anode Region Design and Focusing Properties of STAR Silicon Drift Detectors* , NIM paper to be published

[5] D. Read, *The Silicon Drift Detector and Its Use in a Vertex Tracker*, Ph.D. thesis, University of Texas at Austin, 1996

[6] V. L. Rykov, *Reconstructing the Parameters of Charged Particle Hits in the SDD* (STAR Note #223)

[7] J. R. Hall, A. French, R. Bellwied, C. Pruneau, V. L. Rykov, W. K. Wilson *STAR Silicon Drift Detector Tests* (STAR Note #171)