Proc. 8th Int. Workshop on Vertex Detectors (VERTEX '99)
DeKoog, The Netherlands
June 20-25, 1999

BNL-66963

The STAR Silicon Vertex Tracker: A Large Area Silicon Drift Detector

D. Lynn [a,1], R. Bellwied [f], R. Beuttenmulller [a], H. Caines [c], W. Chen [a], D. DiMassimo [a], H. Dyke [c], D. Elliot [a], V. Eremin [a], M. Grau [a], G. W. Hoffmann [e], T. Humanic [c], I. Ilyashenko [a], I. Kotov [c], H. W. Kraner [a], P. Kuczewski [a], B. Leonhardt [a], Z. Li [a], C. J. Liaw [a], G. LoCurto [c], P. Middelkamp [a], R. Minor [b], M. Munhoz [d], G. Ott [a], S. U. Pandey [f], C. Pruneau [f], V. Rykov [f], J. Schambach [c], J. Sedlmeir [a], B. Soja [a], E. Sugarbaker [c], J. Takahashi [d], K. Wilson [f] R. Wilson [c]

[a] *Brookhaven National Laboratory*

[b] *Lawrence Berkeley National Laboratory*

[c] *Ohio State University*

[d] *University of Sao Paulo*

[e] *University of Texas, Austin*

[f] *Wayne State University*

Abstract

The STAR-SVT (Solenoidal Tracker At RHIC-Silicon Vertex Tracker) is a three barrel microvertex detector based upon silicon drift detector technology. As designed for the STAR-SVT, silicon drift detectors (SDDs) are capable of providing unambiguous two dimensional hit position measurements with resolutions on the order of 20 μm in each coordinate. Achievement of such resolutions, particularly in the drift direction coordinate, depends upon certain characteristics of silicon and drift detector geometry that are uniquely critical for silicon drift detectors hit measurements. Here we describe features of the design of the STAR-SVT SDDs and the front-end electronics that are motivated by such characteristics.

[1] Corresponding author. Address: Brookhaven National Laboratory, Bldg. 510A, P.O. Box 5000, Upton, NY 11973. Email:lynn2@bnl.gov

1 Introduction

STAR is a heavy ion experiment at the Relativistic Heavy Ion Collider (RHIC) at the Brookhaven National Laboratory on Long Island, New York [1]. The main detector component of STAR consists of a large Time Projection Chamber (TPC) surrounded by a solenoidal magnet. In order to improve primary vertexing and to add secondary vertexing capabilities (primarily for strangeness measurements), as well as to expand low momentum tracking capabilities, a silicon vertex detector to be placed around the interaction region is being constructed. In order to handle the expected high charge multiplicities (1500-3000 charged tracks/event) and to minimize the number of readout channels, silicon drift technology [2] was chosen.

The detector consists of a total of 216 SDDs arranged in three barrels around the beampipe at radii of appoximately 6.9cm, 10.8cm, and 14.5cm. In this paper we give an overview of silicon drift detector technology and describe how the technology's unique requirements influence the wafer and front-end design and specifically that of the STAR SVT.

2 General Characteristics of Silicon Drift Detectors(SDDs)

The fundamental geometry of a linear silicon drift detector is shown in Fig. 1a. The detector is fabricated on a thin n-type silicon wafer that typically has a thickness on the order of 300 μm. P+ strips are implanted symmetrically on both the top and bottom of the wafer as shown in the figure. Segmented N+ anodes lie at the edge of the detector and run in a direction parallel to the P+ strips. The detector may be conveniently thought to consist of a drift region (most of the detector area) and a focusing region (the last few millimeters before the anodes). In the drift region the P+ strips are symmetrically biased (top and bottom) with a potential having a gradient in the drift ($\hat{X}$) direction as illustrated in Fig. 1b. A potential valley (as seen by an electron) is formed by the P+/N junction such that ionized electrons generated by the passage of an energetic charged particle drift to the middle (in $\hat{Z}$) of the detector and towards the segmented anodes. The valley thus serves to constrain the electrons away from the surface of the detector to insure full charge collection. In the focusing region an asymmetric potential is applied to the P+ strips to focus the electron cloud towards the readout anodes as illustrated in Fig. 1b. The current signal is then read through appropriate preamplifier electronics. The hit anodes determine the $\hat{Y}$ coordinate, and the drift time from the initial hit to readout of the signal determines the $\hat{X}$ coordinate.

The electron drift velocity v_e is dependent upon the applied electric field E and

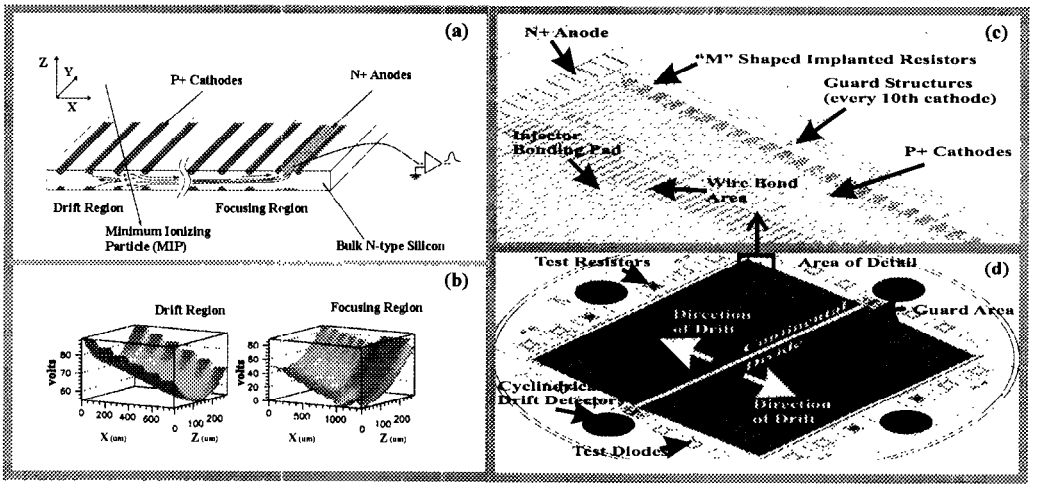

Fig. 1. A) Illustrative principles of SDDs. B) Typical STAR-SDD potential distribution as seen by an electron. C) Magnified view of the corner of the STAR-SVT detector shown in D). D) Enlarged view of a STAR-SVT detector.

is given by $v_e = \mu_e E$, where μ_e is the electron mobility in silicon. The electric field usually is limited to less than 1000 V/cm, below which the electron mobility is independent of E and is equal to 1350 cm^2/Vs at 300K [3]. For typical detector dimensions on the order of a centimeter, this gives characteristic drift times t_{drift} on the order of a microsecond.

The pixel-like readout of the silicon drift detector makes it a good choice for high multiplicity environments (a "pixel" in a drift detector is defined by the anode segmentation in one coordinate and the drift velocity divided by the sampling frequency in the drift direction coordinate). In this case, however, to avoid hits in the detector from interactions separated in time, the interaction rate should be much less than $1/t_{drift} \sim$ 100kHz. For lower multiplicity interactions (proton-proton, for example), this constraint may be eased as the lower multiplicities may permit proper association of the hit to its initial beam interaction.

Hit resolution in the anode (or $\hat{Y}$) direction is primarily a function of the signal level, electronic noise, and level of charge sharing on adjacent anodes. The desirability of charge sharing to improve hit resolution indicates an appropriate anode pitch on the order of the width of the charge diffusion. At the μs scale of the drift time, expected charge widths are on the order of $\sigma \sim 100\mu$m according to the diffusion equation $\sigma(t) = \sqrt{2Dt}$, where $D \approx 3.5\mu m^2/ns$ is the silicon diffusion constant [4]. For charge depositions greater than a few MIPs instant coulomb repulsion results in widths greater than that calculated from the diffusion equation.

The primary challenge of silicon drift detectors is to achieve the best possible hit resolution in the drift direction. This depends upon maintaining a highly linear drift velocity across the detector. The drift velocity is primarily a func-

tion of three quantities- the silicon temperature, the applied electric field, and the applied magnetic field. We consider the implications for the detector design of the temperature and electric field, while a detailed discussion of the effects of the magnetic field may be found in [5] .

The electron drift mobility in silicon is proportional to $T^{-\beta}$, where $\beta \approx 2.4$ at room temperature [6]. We may then write $v_e(T) = \mu_{e0}(T/T_0)^{-\beta}E$, where $T_0 = 300\text{K}$, μ_{e0} is the electron mobility at 300K, and v_e is the drift velocity. Thus the error Δx between the actual hit position x and the reconstructed hit position (at assumed drift velocity $v(T_0)$) is $\Delta x = (x\beta/T_0)\Delta T$ for small changes in temperature about T_0. For a drift distance of 1 cm this gives $\Delta x/\Delta t = 80\mu\text{m/K}$ at room temperature. From this calculation we can conclude that control of the silicon temperature to less than a degree Celsius is necessary in the absence of any online drift velocity calibration. We can also see that it is important to minimize temperature gradients across the detector to much less than a degree Celsius. Finally, to monitor any changes in drift time it is desirable to have online calibration of the drift velocity through direct injection of charge into the silicon bulk at well defined distances. Thus an injection mechanism such as one described in [7] is valuable.

The linearity of the electric field first depends on the linearity of the voltage distribution applied to the P+ strips. Note, however, that as seen in Fig. 1b, the electric field in the center (in $\hat{Z}$) of the wafer is much more uniform than at the surfaces. The surface voltage distribution may be supplied by a resistive voltage divider either implanted onto the detector or located off-detector, or by a combination of the two. In any case to maintain good linearity the dividers must be stiff enough to be unperturbed by leakage currents through the P+ implants. This requires small resistor values. However, in the case of resistors implanted into the detector itself, small resistors result in large localized thermal power dissipation which may degrade drift linearity due to the temperature effects previously described. An appropriate compromise is necessary.

A highly homogenous donor concentration in the n-type silicon is necessary to avoid local non-linearities in the electric field. This necessitates the use of Neutron Transmutation Doped (NTD) silicon wafers as opposed to Floating-Zone (FZ) silicon, where the inhomogeneity is on the order of 3% as opposed to 10-15%[8].

A feature of drift detectors is their small anode capacitance (typically < 500 fF). This is the result of the small geometric size of the anode and the depeletion of the silicon bulk. In order to take advantage of this small input capacitance to minimize the electronic series noise, it is necessary to have the input of the front end preamplifier as close to the anode as possible, and to

minimize the size of the interconnect between the two. This may put important constraints in the location and size of the front-end electronics.

As noted earlier, the charge from a hit particle is spread on many "pixels", with the maximum charge on any pixel from a minimum ionizing particle as low as a few thousand electrons for long drift times. Thus to achieve reasonable cluster finding and hit resolution, single-pixel noise levels below 1000 electrons rms are generally required.

3 STAR-SVT Silicon Drift Detector Design

The main features of the STAR-SVT SDD design are illustrated by the design files shown in Fig. 1c and Fig. 1d. Fig. 1d shows the layout of the SDD on a 4" Wacker NTD wafer with a thickness of 280μm and a resitivity of 3 kΩcm. The detector size is 63 mm $\times$63 mm (after laser cutting). It consists of two half-detectors separated by the "continental divide", the central cathode that receives the maximum voltage bias. The half-detectors drift in opposite directions from one another. Test resistors outside the square detector are used to test implantation quality. Four cylindrical drift detectors are also fabricated on the wafer for use in other applications.

Fig. 1c shows a magnified view of one corner of the SDD. The segmented anodes are at a 250 μm pitch and are 200 μm $\times$ 200 μm in size. The pitch is appropriate over the range of signal gaussian widths 70μm$< \sigma <$ 200μm expected for drift distances between 1mm and 3cm and ionization from 1-10 MIPs. There are 240 anodes for each drift direction. The P+ cathodes are implanted at a 135 μm pitch that is sufficient to maintain a suitably linear electric field in the bulk. Adjacent cathodes are connected via the "M" shaped implanted resistors at the edge of the cathodes as seen in the figure. The guard strips (P+ implants with aluminum overlay) that connect to every 10^{th} cathode serve two purposes. First, each 10^{th} cathode of one half-detector connects to that of the other half detector; thus only one half-detector requires external bias and the other half is automatically biased. Second, the guard strips provide a controlled voltage step-down gradient between the high voltage implants near the center of the detector (i.e. the continental divide) and the cut edge of the detector. This is necessary to prevent breakdown voltages in this region where the voltage gradients on the detector are the greatest. Much effort was devoted to minimizing the guard area, resulting in a design with a 94.5% active area. Four MOS(Metal-Oxide-Semiconductor) injection lines per half detector provide online drift velocity calibration and are described in [9].

There are a total on 220 cathodes on each half-detector, and one common center cathode (i.e. the continental divide). Externally biasing each of these

cathodes via wirebonds, while possible, would significantly increase the complexity of the wire bonding and the mechanical support. The chosen design in which only every 10^{th} cathode is wire bonded resulted in an approximate 2/3 reduction in the number of wire bonds. However, this necessitated having an internal on-detector resistor chain to bias the cathodes between every 10^{th} bond. The values of the internal resistors were chosen to be 500 kΩ. These values were sufficiently large to minimize power dissipation and thus temperature gradients below 0.3°C. Together with a sufficiently stiff external voltage divider (15 kΩ between every 10^{th} cathode) the resistor chain maintains the necessary uniformity in the electric field to permit 20μm resolution.

The nominal working voltage of the STAR SDD is to be -1500 volts. Since each half detector has a drift distance of approximately 3cm, this results in an electric field $E = 500$V/cm, well within the range where the electron mobility is independent of E. The drift velocity then is about 6.75μm/ns, with a total drift time $t_{drift} \approx 4.5\mu$s

Fig. 2 shows linearity measurements made on a SVT SDD using a laser mounted on an x-y table. An injected signal of approximately 10 MIPs was used to eliminate noise influences in the measurements. These measurements were part of the initial SVT system test that featured a half-ladder of four detectors mounted with prototypes of all front-end components. Fig. 2a shows measurements of drift distance vs. drift time and appears highly linear. Measurements are made in 100μm steps along the drift coordinate. Fig. 2b shows the difference between a fit line and the data in Fig. 2a. Two distinct slopes are seen that correspond to the focusing and the drift regions of the detector. Fig. 2c shows a histogram of the deviation of those data points in the drift region from a fit line. An rms of 17μm is obtained. Similar rms values are obtained across all anodes in a detector, though the slopes of the fit line vary from the center anode to the edge anodes due to the thermal gradients produced by the implanted resistors. A similar measurement was made in the anode direction and an rms of 8μm was obtained as seen in Fig. 2d.

In order to account for drift time variations (both temperature-gradient induced and intrinsic) to obtain maximum hit resolution, the drift profile of each detector mounted on a ladder is measured by scanning a laser across the detector at various drift distances prior to installation in the SVT. The mapping of the detectors is performed under similar environmental conditions of air and water cooling provided in the installed SVT. During running conditions, the mapping is used to correct for the detector non-linearities. Small changes in the drift velocity due to deviations in the absolute detector temperature are monitored with the detector injectors. These provide a single correction factor for each detector for changes in the mean drift velocity. In the system test of our half-ladder using the prototype cooling system this correction was found to be unnecessary, as the temperature changes resulted in drift coordinate

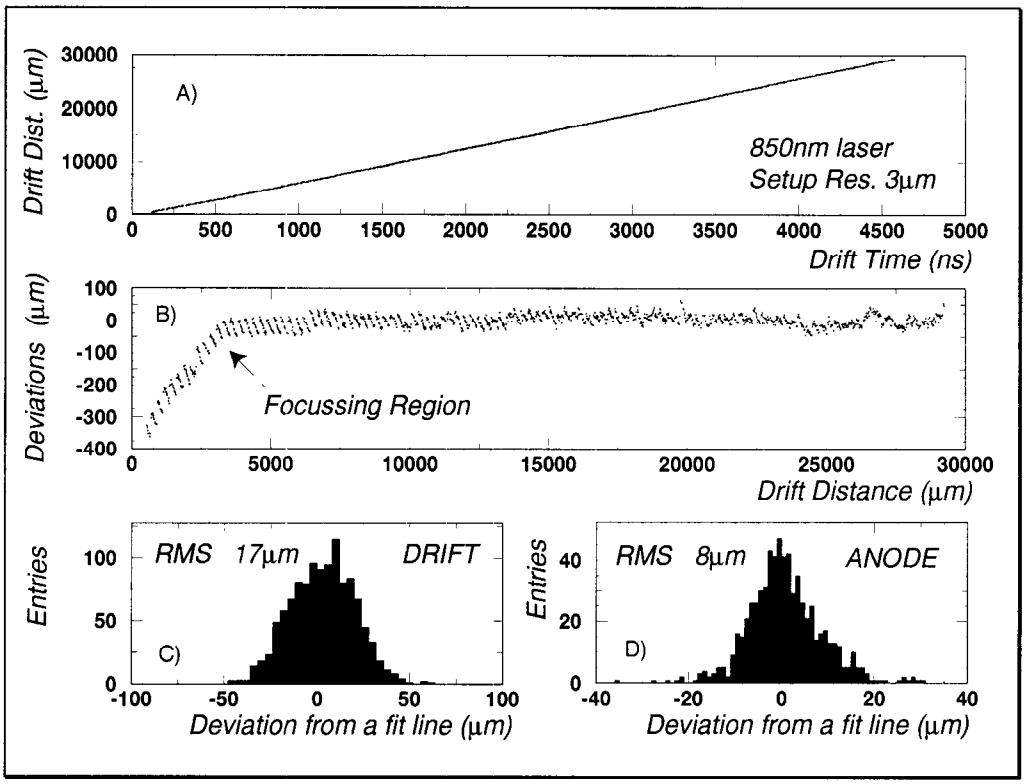

Fig. 2. A)Laser measurement of drift time vs. drift distance on a SVT SDD in a system test. B) Difference between data points in drift region and a straight line fit. C) Histogram of difference between data points and straight line. D) Histogram of difference between actual and reconstructed position of laser injection hit in anode direction.

deviations of only 5μm rms once the ladder had reached its mean operating temperature after power-up. These deviations were measured over a two hour period.

The above linearity measurements yield the intrinsic resolutions of the SDDs. To obtain measurements of hit position resolutions, these measurements were repeated using single MIP laser injection, taking one measurement at each 100μm step in the drift direction. Thus these measurements include the effects of electronic noise (ENC = 530 e- rms as described in the following section) and detector drift non-linearites. Typical values obtained across the detector were 20μm and 25μm in the anode and drift directions, respectively.

4 Front-end Multi-Chip Module Design

The primary purpose of the front-end electronics is to amplify and shape the detector current signal arriving via the detector anodes such that time-of-arrival and total charge information are preserved. The front-end must also

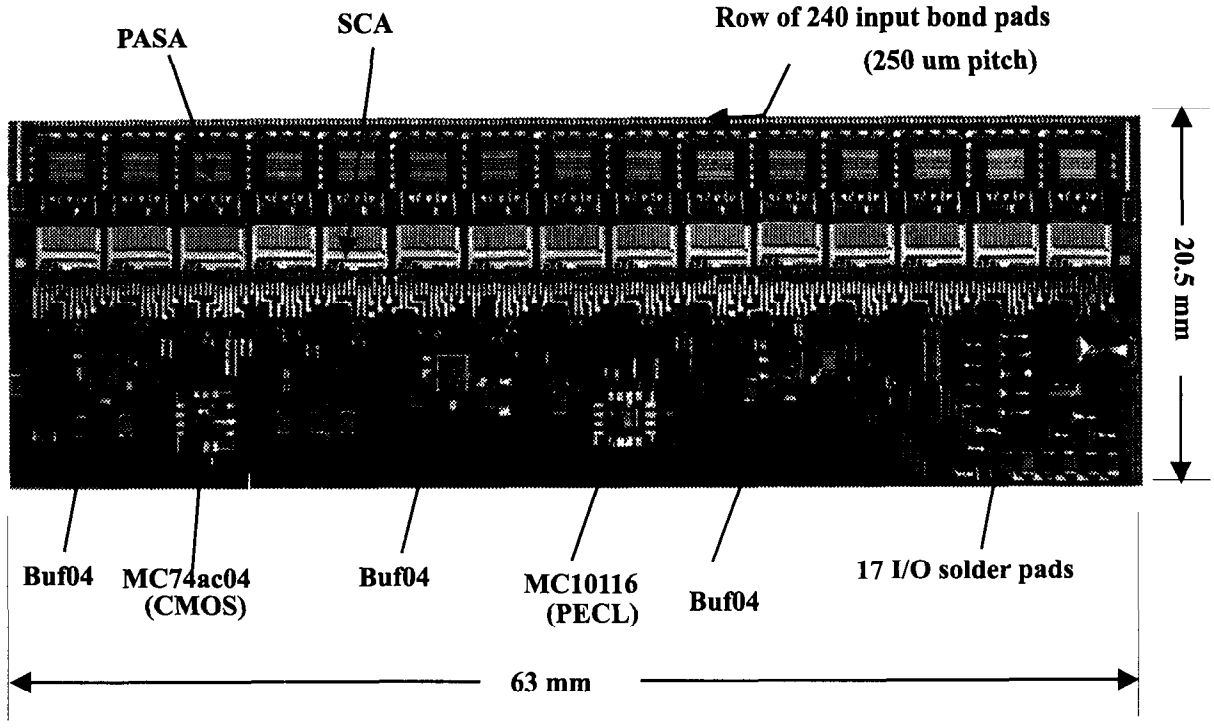

Fig. 3. SVT Front-end Multi-Chip Module

minimize the noise introduced into these measurements. A photograph of the STAR-SVT front-end multi-chip module(MCM) is shown in Fig. 3. We will only discuss in general terms the MCM design as a detailed discussion may be found in [10].

As previously stated, the detector anode capacitance of an SDD is much less than 1 pF. In order not to significantly increase this input capacitance, the interconnect between the anode and the front-end transistor should be no longer than approximately 1 cm (a general rule-of-thumb gives 1pF/cm for a variety of interconnects, e.g. bond wires, printed circuit signal traces, etc.). Fig. 3 shows a row of 240 input pads that are spaced at the detector anode pitch of 250 μm and that are intended to be wire bonded to the SDD anodes. Each of these pads has a second wire bond to an input pad of the PreAmplifer ShAper (PASA) die. Measurements indicate each of these bond wires contribute about 170 e- (electrons) rms noise, though only about an additional 30 e- rms when added in quadrature to the other noise sources. This is an indication (though, of course, preamplifier dependent) of the importance in minimizing this interconnect.

The low input capacitance permits a fast preamplifier-shaper response due to the lower series (capacitance dependent) noise. A fast impulse response is desirable for good separation of charge from two hits arriving on the same

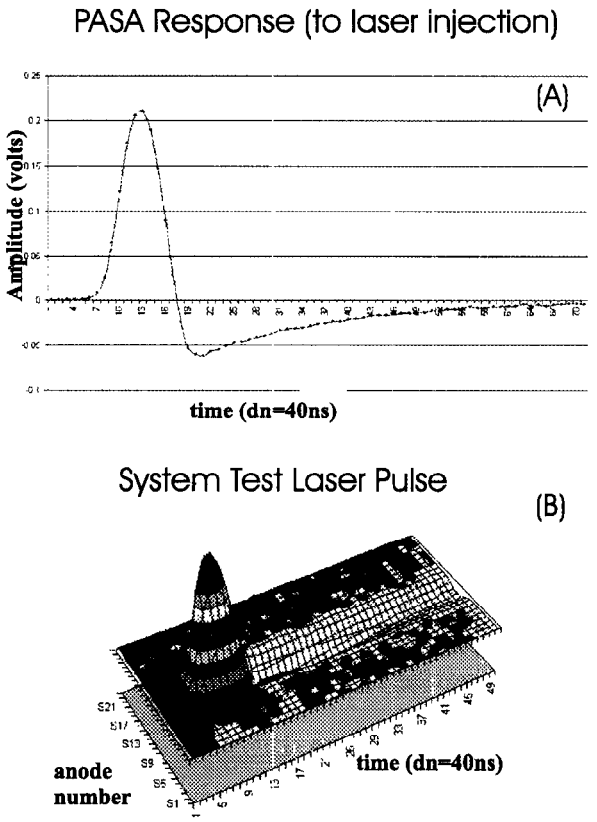

Fig. 4. A) PASA response to laser injection hit. B) Three dimensional reconstruction of laser pulse.

anode(two-track resolution). "Fast" is therefore defined relative to the expected separations and widths of adjacent hits and in our case refers to a risetime (10-100%) of 50 ns. The PASA was thus implemented in bipolar technology over CMOS for its superior noise-power performance at such speeds. The shape of the impulse response is shown in Fig. 4a. It was chosen to be bipolar in order to avoid baseline shifts due to the high hit multiplicity of a heavy ion environment. Details of the 16 channel PASA are listed in Fig. 5.

It is necessary to have waveform sampling to preserve both timing and dE/dx information. The number of required samples is determined by $N \times 1/f_{sampling} > t_{drift}$, where N is the number of samples, $f_{sampling}$ is the sampling frequency, and t_{drift} the total SDD drift time. In order to avoid aliasing of the output waveform the sampling frequency should be greater than twice the PASA bandwidth of about 10 MHz. We chose $f_{sampling} \approx 25$MHz, which together with our drift time of about 4.5μs suggests the number of samples to be greater than 112. The front-end is thus designed for 128 samples which, as a power of two, is naturally implemented.

In order to reduce the number of output channels per MCM (nominally 240)

FEE Parameters and Specifications

Multi-Chip Module	
Technology	Thick Film on beryllia substrate Dupont QM silver based pastes
Size	63 mm x 20.5 mm
No. Channels	240
No. Time Samples	128/channel
Total No. "pixels"	128 x 240=30,720
Output channels	3 (for 10,240 pixels)
Sampling Rate	25 Mhz
Readout Rate	2MHz 10,240/2MHz = 5 ms
Power Dissipation	12 mW/Channel = 2.9 W
Noise	ENC = 490 e- rms (no detector) ENC = 530 e- rms (w/ biased detector)
Radiation Length	1.4%

PASA	
Technology	Maxim SHPi NPN-BJT process
Die Size	3.3 mm x 2.4 mm
No. Channels	16
Power Dissipation	3.8 mW/Channel
Gain	7.2 uV/e-
Noise (C_{in}< 0.2 pF)	ENC = 380 e-
Peaking Time (10-100%)	50 ns
Dynamic Range	2 V (275,000 e-)

SCA	
Technology	Orbit 1.2 um CMOS
Die Size	3.8 mm x 2.8 mm
No. Channels	16
No. Capacitor/Channel	128
Power Dissipation	6 mW/Channel
Noise (C_{in}< 0.2 pF)	2.2 mV (ENC = 300 e- rms)

Fig. 5. FEE Parameters and Specifications

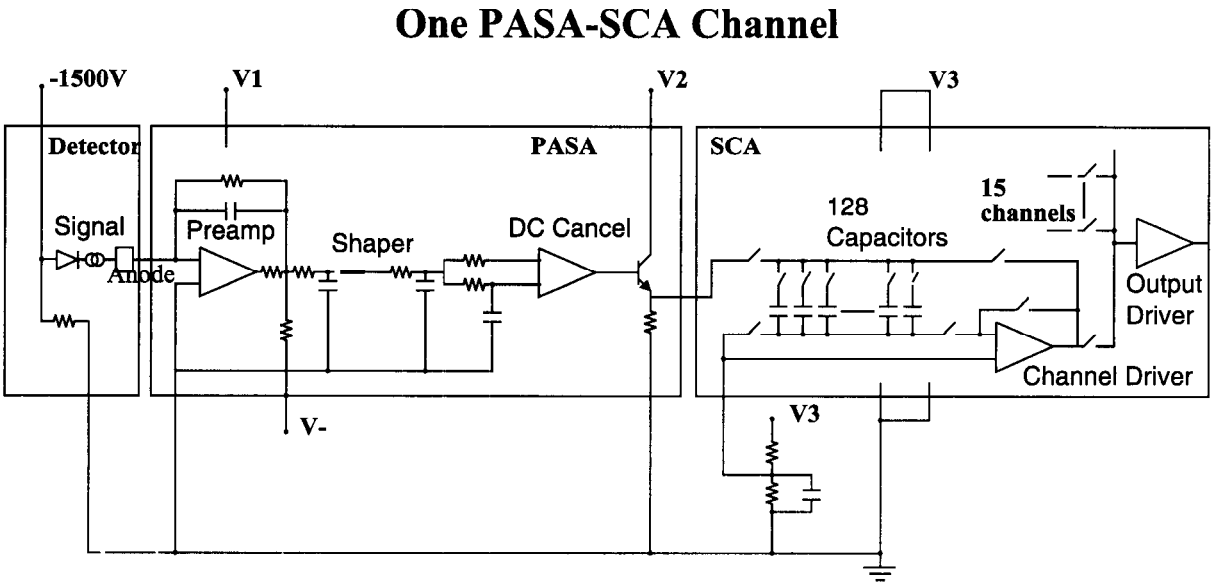

Fig. 6. Simplified schematic of one PASA-SCA channel

it is desirable to have a level of multiplexing at the MCM. However, due to the waveform sampling, there are $128 \times 240 = 30,720$ pixels per half-detector or MCM. A solution involving front-end digitization before multiplexing is prohibitive due to the large number of pixels. An analog multiplexing scheme is required and is implemented with Switched Capacitor Array integrated circuits [11]. Detailed specifications of the SCA are given in Fig. 5.

A simplified schematic of one PASA-SCA channel is shown in Fig. 6. A PASA channel contains preamplifying and shaping circuitry followed by a DC cancellation stage. This stage results in zero gain at dc (direct current), thus making the PASA output independent of the detector leakage current. A final emitter-follower stage provides higher current driving capability. The output PASA waveform is sequentially sampled on each of the 128 capacitors on SCA channel. After 128 capacitors have been written an internal counter resets and sampling continues, overwriting previous values. The SCA and thus the detector is continuously active until a trigger is received and readout of the SCAs begins. The signals on all $128 \times 16 = 2048$ capacitors per SCA are multiplexed onto one common analog output. An arbitrary number of SCAs may be multiplexed together. In the implementation chosen for the SVT MCM, groups of five SCAs (80 channels or 10,240 pixels) are multiplexed. Therefore there are three analog outputs per MCM. These are driven by the Analog Devices' Buf04 analog drivers labeled in Fig. 3. A commercial CMOS chip for the SCA control circuitry and an ECL chip for the differential clock to the SCA are also shown in Fig. 3. A reconstructed laser-injected signal from a SDD read through an SVT-MCM is shown in Fig. 4. The main specifications of the MCM are summarized in Fig. 5.

5 Ladder Components

Fig. 7 shows the main components of one complete ladder of detectors and electronics. Thirty six such ladders (with either seven, six, or four detectors) arranged in three barrels constitute the full SVT. The ladder in the figure contains seven SDDS. The cathodes of the three SDDS on the left connect via wire bonds on the top and bottom surfaces. The SDD furthest to the left is bonded to the High Voltage Interface board (HVI). The HVI supplies the approximately 30 voltages needed to bias the SDDs. These voltages are derived from a stiff voltage divider located about 15cm away on the Transition Board(not shown) that is connected to the HVI via a small gauge HV cable capable of withstanding over 12kV. The stiff divider is located away from the detector in order to avoid coupling its 7.5 watts power dissipation into the detectors and introducing drift non-linearities. The four detectors on the right are similarly connected and biased, but the left and right sides of the ladder are decoupled in order to avoid large ground loops in the design. Thus each half-ladder, including its associated electronics, are electrically isolated from each other.

The wafers are epoxied to a beryllium frame that is connected to beryllium end-rings via mounting posts at each end of the ladder (note the mounting holes visible in each HVI). Each SDD generates about 100mW due to its internal resistor chain. This heat is absorbed by a flow of air maintained at room temperature to within a degree Celsius and at 50% relative humidity.

The MCMs for each detector are epoxied to a carbon-fiber electronics carrier containing a small channel through which water flows to sink the heat generated by the front-end electronics. Thin printed circuit cables (PCCs) are epoxied to the underside of the electronics carrier. These cables contain digital differential receivers which receive and drive the clock and control signals to operate the SCAs. They connect to each MCM via a custom fabricated "wrap-around" cable made of small gauge magnet wire. The wrap-around cables connect the MCM input/output with the PCC. The analog outputs from each MCM are driven single-ended through the PCC and via the Low Voltage cable to the transition board, where, in addition to the external detector voltage divider, analog differential drivers are located. These drivers convert the single-ended analog signals to differential levels and drive the signals via shielded twisted-pair cable a distance of about 10 feet to the readout system where the signals are digitized. The digitized data are driven via fiber optic cable to the data acquisition system. The data acquisition system zero-suppresses the data and writes it to tape.

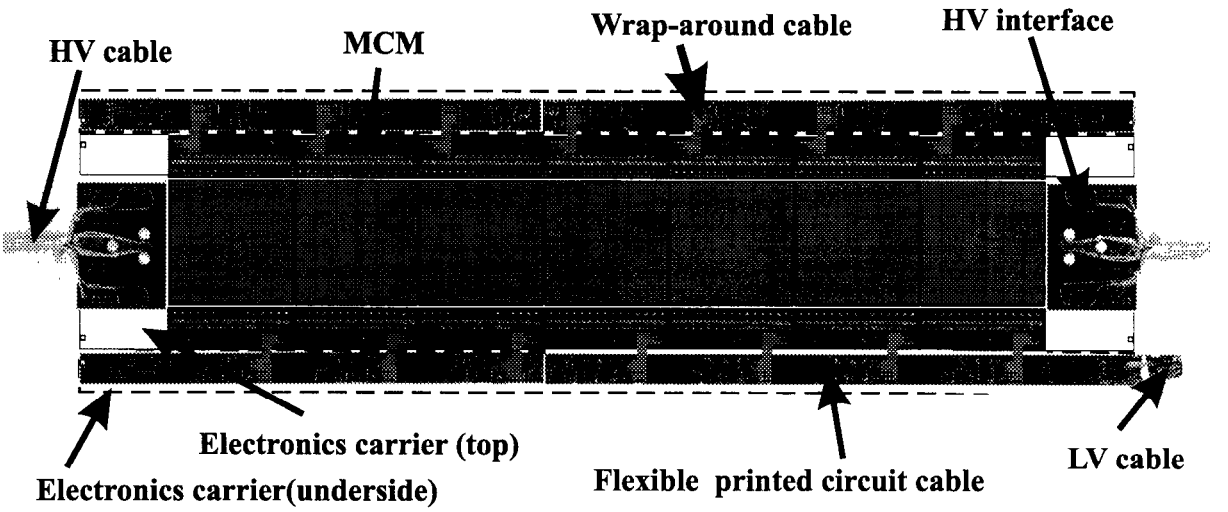

Fig. 7. Components comprising one of 36 STAR-SVT ladder assemblies

6 Conclusions

The STAR Silicon Vertex Tracker is a microvertex detector implemented in the novel silicon drift detector technology. It is intended for the high multiplicity/low rate environment of the new Relativistic Heavy Ion Collider at Brookhaven National Laboratory. Drift detectors permit high resolution unambiguous two-dimensional hit position measurements with readout of only one edge of the silicon wafer. The design goal is to achieve position resolutions of 20μm in both coordinates.

Experiment E896 at the Alternate Gradient Synchrotron at BNL recently built a 15 layer silicon drift detector that was run in a 6 Tesla field in a heavy ion environment [12]. This detector used SVT detectors and MCMs as well as other features of the SVT design. The experiment demonstrated that drift detectors are now a mature technology.

Construction of the STAR SVT is now underway, with completion scheduled for the second half of the year 2000. When completed, the detector will consist of 216 SDDs containing over 13 million pixels multiplexed onto just 1300 readout channels.

7 Acknowledgements

This work is supported in part by the US Department of Energy(under contract No. DE-AC02-98CH10886) and the Robert A. Welch Foundation.

References

[1] J.W. Harris and the STAR Collaboration, *Nucl. Phys.*, **A566** (1994) 277c

[2] E. Gatti and P. Rehak, *Nucl. Instr. and Meth.*, **225**, (1984) pp. 608.

[3] W.R. Leo, Techniques for Nuclear and Particle Experiments, 2^{nd} Ed., (Springer-Verlag, Berlin, 1994)

[4] E. Gatti and P.F. Manfredi.,*Revista Nuovo Cimento*, **9**, 1 (1986)

[5] S.U. Pandey *IEEE Trans. on Nucl. Science*, **44** (1997) 610

[6] S.M. Sze, Physics of Semiconductor Devices, 2^{nd} Ed. (Wiley, New York, 1979)

[7] E. Gatti *et. al., Nucl. Instr. and Meth.*, **A295** (1990) 489-491

[8] G. Vilkelis, Silicon Drift Detector R & D for the STAR Experiment at RHIC, Ph.D. Thesis, Univ. of Pittsburgh, 1995

[9] R. Bellwied *et. al.,Nucl. Instr. and Meth.*, **A416** (1998) 70-80

[10] D. Lynn *et. al.*,"A 240-Channel Thick Film Multi-chip Module for Readout of Silicon Drift Detectors", *Nucl. Instr. and Meth.*, submitted

[11] S.A. Kleinfelder *IEEE Trans. on Nucl. Science*, **37** (1990) 1230

[12] S.U. Pandey *et. al.*, "The 15 Layer Silicon Drift Detector Tracker in Experiment E896", *IEEE Trans. on Nucl. Science*, submitted