5.3.5 Cooling

The silicon strip and pad detectors is in a temperature controlled environment maintained by a single closed circuit cooling system. Preliminary calculations and lab tests indicate that air-cooling is feasible. The advantages of an air-cooling system include low cost and minimal material in the path of the particles. Furthermore, the possibility of coolant spills in the region of the silicon is eliminated.

A plenum of 50-μm thick aluminized mylar will surround the hybrid circuits, and provide a path for the cooling air to flow around them. It will also help electrically isolate the hybrid circuits. A pressure difference within the plenum will force the air flow from the barrel hybrids to the endcap hybrids.

The electronics mounted on the hybrid circuits will dissipate 5 mW per channel, and there are 256 channels per hybrid. The cooling system, including a fan and heat exchanger located outside of the support cylinder and connected to it at both ends, will maintain the operating temperature of the electronics at less than 50 °C, where components fabricated in CMOS become temperature sensitive. It should be noted that the operating temperature of the magnet coils is between 30 and 40 °C. However, the MVD is not in direct thermal contact with the mmagnet. The silicon wafers operate at room temperature and are physically separated from the electronics by the cable. Therefore the electronics are the main heat source in the system, and determine the cooling requirements. It is also crucial to control the humidity inside the enclosure to avoid condensation problems on the detector wafers or any other elements.

Preliminary calculations and measurements indicate that with a total input heat of 100 W and a flow rate of 40 CFM it is possible to maintain the temperature in a test plenum to 45 °C with a temperature difference between air inlet and outlet of several degrees. The specific heat of the air is 0.61 W per cubic foot per degree Celsius. It is important to note that the cooling of the electronics is accomplished through the temperature difference between the electronics and air in the plenum; not by a large temperature change in the flowing air.

We expect an operating temperature for the electronics below 50 °C, and 200 W from the 40,000 channels at 5 mW. This implies flow rates in the vicinity of 100 CFM and a few degree change in the air temperature.

Elements of the first level trigger, and low level data collection interface electronics are located behind the silicon pad detectors. It is still not fully determined how much area these elements fill, nor how much power they dissipate. Cooling of those elements is provided by the same MVD cooling system. Should higher cooling power than can be provided by extending the air system be required, it is possible to water cool using the same closed circuit cooling system which cools the heat exchanger in the air cooling system.

5.4 Electronics

5.4.1 Front-end

The silicon detector systems will utilize front-end electronics custom designed for PHENIX, and fabricated using a CMOS process. This work is currently underway at Oak Ridge

National Laboratory. Since the silicon pads range in capacitance from 10 to 25 pF and the silicon strips are of order 10 pf, the same front-end electronics can be used. This allows development of only one preamplifier circuit and use of identical hybrid circuits incorporating these preamplifiers throughout the silicon system.

The requirements for the silicon detector electronics are to provide the pulse height from each strip or pad along with the information of which element was hit by a charged particle. The charge must be collected rapidly to ensure that the electronics are ready for a possible interaction in the next bunch crossing. The silicon detectors also provide information on the charged particle multiplicity density, which is derived from the analog sum of pulse heights on multiple elements. Consequently, a good gain match is required for all the elements which contribute to one sum. It has been shown that channels on a single silicon die in an ASIC are well matched, [3] but will require gain matching over multiple dies. The pulse height resolution required to separate multiple charged particle hits on a single element implies signal to noise of at least 10:1.

The front-end electronics chain consists of an amplifier/shaper chip and an analog memory plus digitization, or AMU/ADC, chip. The AMU is a switched capacitor array with simultaneous read/write capability; it pipelines the data pending a low level trigger decision, and makes the data available for use in higher level trigger decisions. The AMU/ADC circuit is under development for all of the PHENIX subsystems, as described in Chapter 11; the silicon detectors will use a variant of this ASIC which is modified for lower resolution and smaller power dissipation. We require eight bits of resolution in the ADC to enable fitting the pulse height spectrum of the silicon pads to determine the number of charged particle hitting a single pad, from zero to four. The AMU output amplifier is optimized for low power consumption, consistent with the 8-bit ADC, which is less than in the other PHENIX detector subsystems.

The beam bunch crossing time implies short shaping times on the preamplifier. We expect a shaping time of 25–30 ns, and baseline restoration by 120 ns. With this rather short shaping time, it is feasible to achieve signal to noise of approximately 10:1. Each premamplifier chip contains 32 channels at a 100-μm pitch and will be mounted as bare (unpackaged) silicon on the hybrid circuit described below.

5.4.2 Trigger and Data Acquisition

The silicon provides information to the LVL-1 trigger through the analog sum formed on each 32 channel preamplifier chip. These sums are formed either by using a fraction of the charge from each channel, or via an analog sum of digital signals, namely summing discriminator outputs. The analog sum signals from all eight preamplifier chips resident on one hybrid circuit are then further summed and driven as an analog signal to one of the trigger electronics boards located behind each endcap motherboard. In this location the analog sum from each hybrid circuit is digitized and further summing is done digitally. This arrangement allows a straightforward formation of left and right side sums. The digital summing of signals from different hybrids allows gain factors to be used. Consequently the gain matching requirement among preamplifier chips is only very stringent among the eight chips present on one hybrid circuiit.

Following the AMU/ADC, zero suppression is performed on the hybrid circuit on a Data Acquisition Interface Level 1 (DAI1) element. This element will either be an ASIC or Field Programmable Gate Array servicing all 256 channels. This element also includes a multiplexer/FIFO to drive the signals on a moderate speed (10 MHz) bus connecting six hybrids to the next level data acquisition electronics. These electronics are located either outside the central magnet, or behind the endcap motherboard near the LVL-1 trigger electronics.

5.4.3 Hybrid Circuit

As shown in Figs. 5.5 and 5.6, each detector wafer has 256 channels and is connected via kapton cable to one hybrid circuit. A layout of the hybrid circuit is shown in Fig. 5.8. The signals enter the preamplifier/shaper/summer and then the AMU/ADC which both have 32 channels per die. A 32 channel buffer chip follows the ADC, and may possibly be incorporated into the AMU/ADC. The large size of the AMU/ADC dies mandates the staggered arrangement shown. The first level data acquisition functions are shown in the figure as a separate DAI1 chip with 256 channels, following the AMU/ADC. The read and write functions in the AMU are controlled by a single heap manager per hybrid circuit. This chip is common to all of the PHENIX subsystems. The DAI1 presents data to the data bus. Each hybrid also contains an analog sum driver which collects the analog sums from the eight preamp/shaper/summer chips, further sums them, and drives the resulting signal to the LVL-1 trigger electronics residing behind the endcap motherboards.

The hybrid circuits are fabricated either using a standard alumina thick film process, or a process from Texas Instruments called High Density Interconnects (HDI). This process looks like a better choice for the trace densities required in our circuit. The process uses alumina substrates with milled wells for the silicon dies. A polyimide coating is laser-drilled at each die contact and then metallized to provide electrical contact and interconnect traces. This is a proven process that allows pitches of less than 100 μm. Hybrid circuits fabricated using the HDI process can be repaired by removing and replacing defective parts then recoating and remetallizing.

5.4.4 Testing, Alignment, and Calibration

As discussed above, extensive testing of the components of the MVD is required prior to and during the assembly process. The front-end electronics must be tested to ensure each chip is functional before assembly of the hybrid circuit. The gain matching requirement for chips on a circuit implies measuring of the gain during this testing step. Each completed hybrid circuit must then be tested again to ensure functionality of all channels before bonding to the detector. Our choice of hybrid manufacturing technology offers the possibility of repairing hybrids, which should reduce the cost by improving the yield of usable ones. Each silicon detector wafer must also be tested prior to assembly. We are currently designing testing protocols for the detectors, and plan to use a laser diode and computer controlled $x-y$ table to automate these tests. Specifications and testing protocols for the hybrid circuits will be

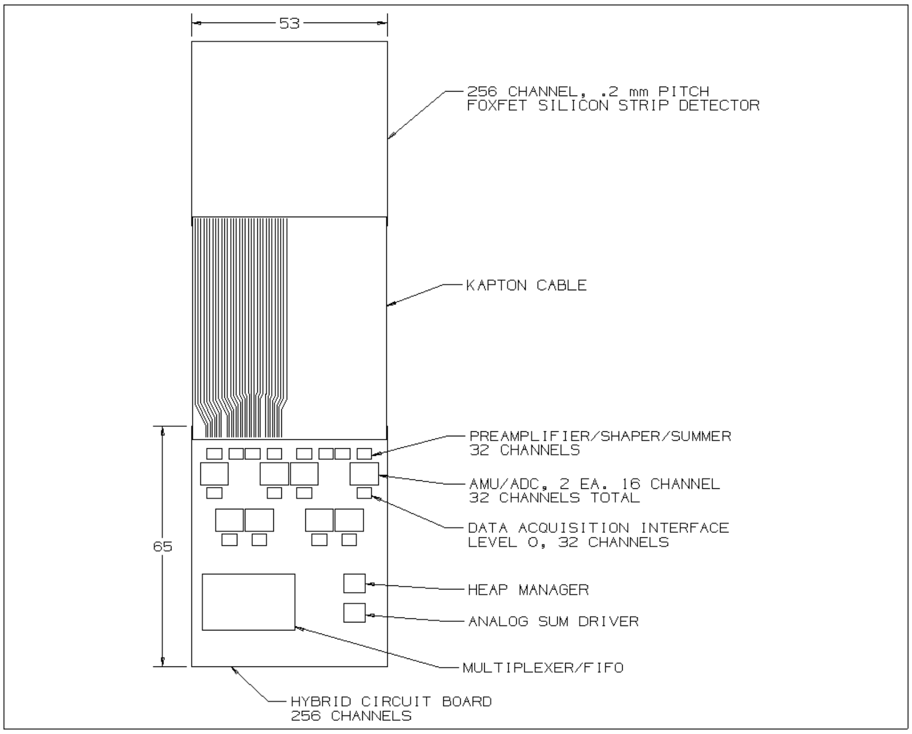

Figure 5.8: Layout of the hybrid circuit. The various components are labeled.

developed with the hybrid design engineers and the manufacturers. We expect some of this testing to be performed by the manufacturer before releasing the product.

The silicon barrel will be assembled into azimuthal sectors as described above, and the pads into endcap halves. These sectors can then be further assembled into the enclosing half-cylinders and tested with a laser diode. The laser diode chosen is a 1064-nm infrared laser, manufactured by Laser Diode, Inc. It is coupled to a 150-μm single mode fiber to deliver a 25-μm spot size, enabling the study of individual silicon pixels. The laser diode can produce charge equivalent to that deposited by a minimum ionizing particle in several silicon detector layers. At the assembly stage this system will be used to test the detector wafers and completed detector-plus-hybrid assemblies. Its use as a calibration system is feasible and is currently under study.

The vertex position requirement in the z direction is 2 mm; in the plane transverse to the beam, the definition is the beam size ($\sigma_T \approx 0.45$ mm for Au + Au [4]). Some chip placement aberrations result in a track ending up in a strip neighboring the expected one, reducing the peak value of the correlation function between the two layers, which is the quantity used to locate the vertex. Preventing this kind of error defines most of the constraints on placement of the chips in three directions, z (along the beam axis), R (radial), and s (circumferential, or Rϕ), plus the three rotations around these axes. The chips lie at radii R_1 and R_2, which are 5 and 7.5 cm, respectively, in the calculations below. The size of a chip in the z direction

Table 5.3: Summary of tolerances in placement of inner/outer chip pairs relative to each other.

Axis	Max. error in relative position	Max. error for relative rotation about this axis
z (beam)	0.7 mm	0.2°
R (transverse)	0.4 mm	0.3°
s (circumference)	0.5 mm	0.3°

Table 5.4: Summary of tolerances on the placement of chip pairs relative to the beam.

Axis	Max. error	Max. error for rotation about this axis
z	2.0 mm	1°
R	1.3 mm	16°
s	no limit	1°

(Z_{chip}) is 5.3 cm.

The vertex finding constrains only the relative placement of chip pairs, one chip on the inner barrel layer and the corresponding chip(s) on the outer layer. Here, one example of the determination of the alignment constraints is given. The other constraints are determined similarly and are summarized in Table 5.3

Consider the displacement in $R(\Delta R)$ of one of the layers relative to the other. When scaling the outer pattern by the nominal R_1/R_2, a radial displacement would result in a pattern that is improperly scaled. A calibration procedure could find the actual R_1/R_2, but the nominal ratio may be used in a trigger-level vertex finder. In order to limit the error such that in the worst case, a track is displaced by one strip (200 μm),

$$200\mu m = \frac{R_1}{R_2} \times \frac{Z_{chip}}{R_2} \times \Delta R$$

This is satisfied if radial displacement of the outer chip is less than approximately 0.4 mm.

The tolerances above are between inner/outer chip pairs only, but there are further constraints on the positions of chips pairs relative to the beam. Table 5.4 summarizes the constraints on the positions of pairs of chips relative to the beam.

Calibration of the system is achieved by the injection of a known amount of charge into each preamplifier channel. Such a calibration feature must be built into the ASIC, and is required for each front-end electronics channel in PHENIX. The charge injection allows testing of the entire electronics chain through the data acquisition system, different locations