Production and Testing for the D0 SMT F-Disks

0.1 Introduction

To assemble the 12 planes of F-disk detectors we needed 144 trapezoidal shaped sensors which were attached to the hybrids (HDIs) containing the readout electronics to make F-wedges. The F-wedges are AC-coupled, double sided detectors with an angle of $\pm 15°$ with respect to the vertical direction. The pitch is 50 μm and 62.5 μm on the p- and n-side respectively and constant over the wedge, whereas the strip length varies across the face of the detectors. The n-side has 768 strips that are connected to 6 SVXIIe readout chips. The p-side has a total of 1024 strips to be read out via 8 SVXIIe chips. This makes for a total of over 258,000 channels contained in the F-disk detectors or about one third of the entire SMT project. The sensor production and testing is described above. Here we describe the construction of the hybrids, assembly and testing of the wedges, and assembly and testing of the disks.

F-wedges are made from a silicon piece glued to the electronics hybrid called the HDI. The silicon sensor is qualified separately before it is assembled. This production and testing is described above. The HDI assembly is made from two flexible HDI pieces as well as a jumper which are laminated onto a beryllium substrate. The HDI lamination, population, and testing processes are described in the next section. The silicon sensor and HDI are each tested thoroughly before being glued together. The gluing process is controlled for alignment using a coordinate measuring machine. After gluing, the channels on each of the sensors are connected by wirebonding to the readout electronics. Once each of the wedges is qualified, it is mounted on a beryllium ring using a coordinate measuring machine. For each disk, the wedges are alternated for either n-side up or p-side up as shown in Figure 0.1. Each of these rings were produced by laminating two flat beryllium pieces together so that they can enclose the cooling channel. Once all of the wedges have been aligned onto a ring an F-disk is produced and tested and prepared for installation into the final setup. These procedures are described below.

0.2 HDI Assembly

The HDI assembly is composed of over 20 pieces which are laminated together. The most important pieces of the assembly include: a) the 6-chip side flex circuit, b) the 8-chip side flex circuit, c) the kapton jumper (pitch adapter) for the 6-chip side, and d) the beryllium substrate. Each of these four pieces was qualified separately before the HDI was assembled. The F-wedge HDI is then fabricated as

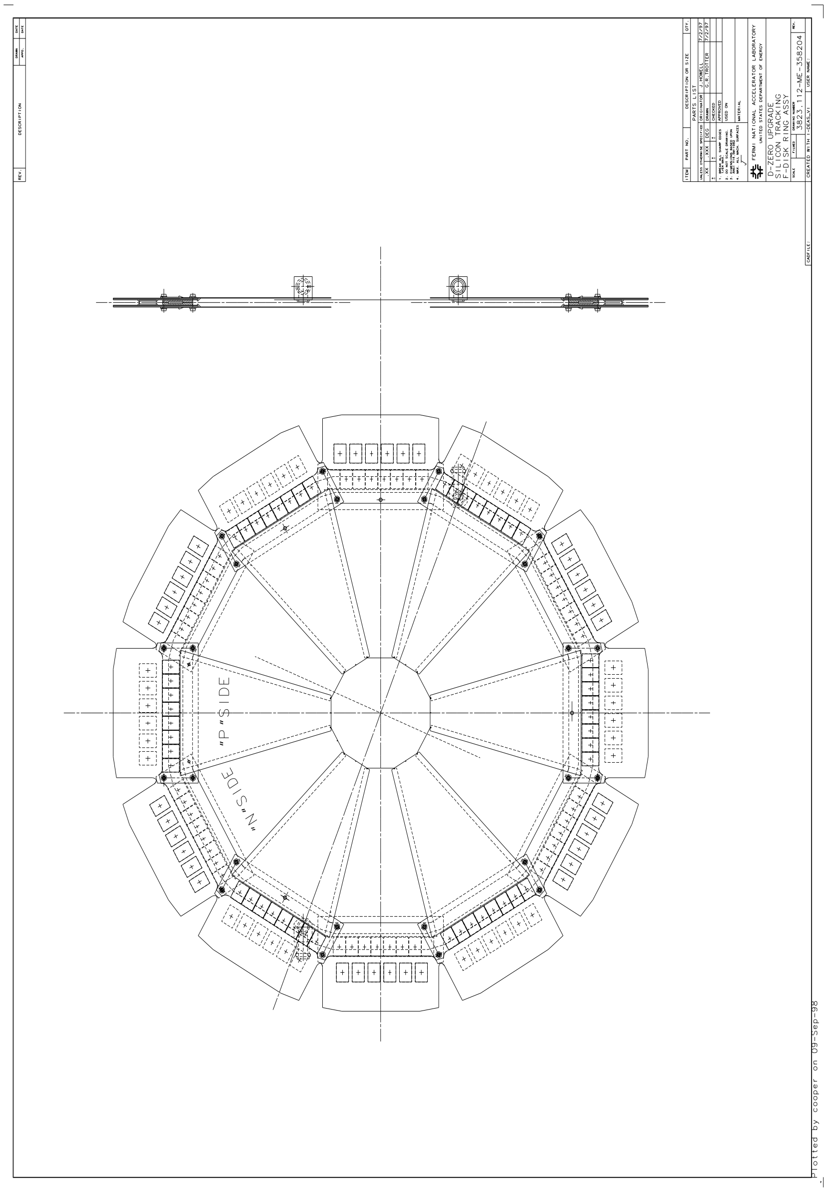

Figure 0.1: View of an F-disk detector with alternating wedges having p-side or n-side to the front.

a double sided device starting with a lamination procedure. Figure 0.2 shows the specifications for the HDI on the n-side.

The 6-chip and 8-chip side flex circuits were fabricated by Compunetics[1]. After receipt of the circuits, each was probe tested as described in the earlier sections. Due to problems from the surface mount procedure further in the process, each of the circuits was then visually inspected before lamination. This visual inspection searched for raised sections of the cover layer on the flex circuit itself. While it was not clear that these imperfections in the flex circuit were causing our surface mount problems, we erred on the side of conservatism for accepting the bare circuits. Early on in the production process, almost 20% of these circuits failed the visual inspection. After iteration with the vendor, 99% of the circuits passed this visual inspection.

The beryllium substrates were inspected for dimensional correctness and flatness (less than 20 microns over the whole surface). A slot test was used to test for flatness as well as a Coordinate Measuring Machine (CMM) measurement of 6 points over the surface of the beryllium. Later in the production, beryllium substrates were recycled for use in the detector. After recycling, the beryllium substrates sometimes needed to be flattened or recleaned before they passed the flatness requirement.

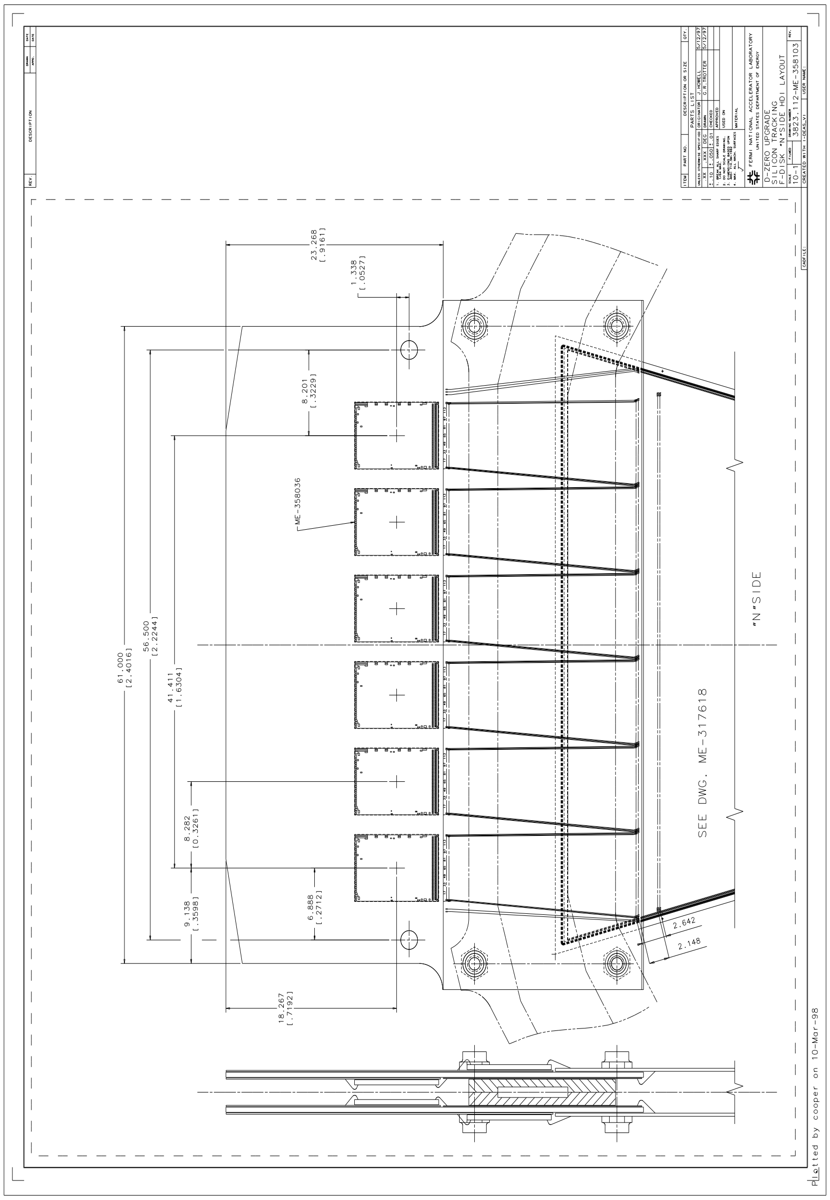

Figure 0.2: Drawing of the n-side of the HDI

0.2.1 Jumper

To adapt the pitch from the sensors to the SVXIIe readout chips for the n-side, we used a flexible circuit manufactured by Max Levy corporation.[2]. It is a rectangular circuit 21.25 mm by 65.80 mm. On this circuit there are 768 traces arranged in 6 groups of copper traces for each of the channel readouts. In addition there are 5 more traces, for the bias voltages. The traces are roughly 21 mm long with a width of 24 micrometers and are made of copper with nickel-gold plating so that the resistance between bond pads is 10 ohms or less. The substrate is .001 inch thick Polyimide film. A drawing of this circuit mounted on the HDI can be found in Figure 0.2.

In addition to functioning as a pitch adapter, the wedge is mounted to the cooling ring on this jumper surface. Therefore the jumper needs to provide a surface that will allow heat transmission and that is flat. The fragile traces on this jumper must therefore be protected against mechanical abrasions by some sort of cover layer.

Early prototypes for these circuits were probed for trace resistance and continuity as well as bondability. These early circuits worked quite well. A cover layer was added by the company to help protect the traces. After this cover layer was added, we were no longer able to bond to these circuits. An extensive testing program was instituted to try and recover the bondability by cleaning the circuits. Mechanical and chemical cleanings were tried, but the success was very unpredictable. We iterated with the company to produce these circuits without their cover layer. Without the cover layer, we once again could bond to the circuits. This caused us a delay of about 6 months in the production of the F-wedges.

Once we received coverlayerless jumpers, we visually inspected them but did not reprobe them for shorts and opens. We developed a method to apply our own cover layer of 100 MT kapton with 8043 adhesive on the jumper by essentially taping it using specially machined fixtures to maintain the alignment.

0.2.2 Lamination

The parts of the HDI were laminated together by technicians working in the Sidet facility at Fermilab. With the number of parts for this device due to its dual sides, these procedures became quite complicated. Figure 0.3 shows the lamination assembly required for the F-wedge HDIs. Some of the tedlar and Ablefilm pieces required careful cutting and trimming by hand. Later in the production process, we sent these pieces out for laser trimming to save the manpower involved in this step. After assembly on a precisely machined fixture the lamination assembly was inserted into a vacuum jar where it was evacuated and later heated to complete the lamination. Figure 0.4 shows the fixture used for this process.

The procedures for lamination required extensive development and training. We trained two technicians to perform this process. There were two vacuum

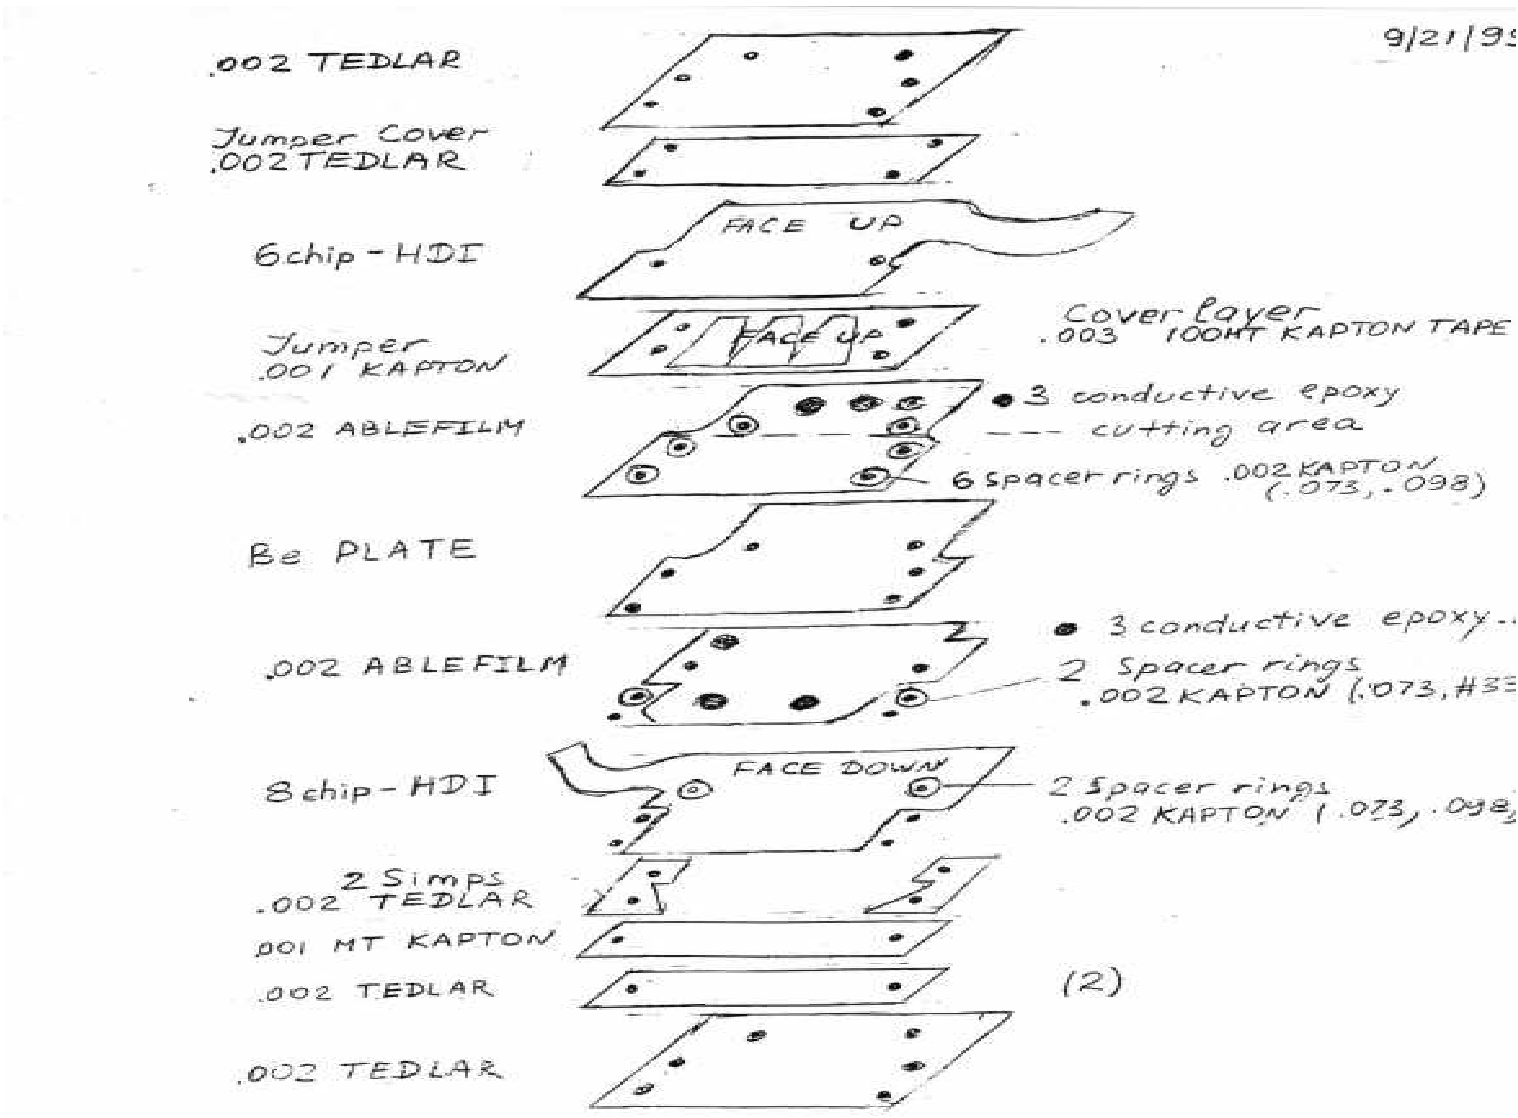

Figure 0.3: Parts that are assembled together for the lamination of the F-wedge HDIs.

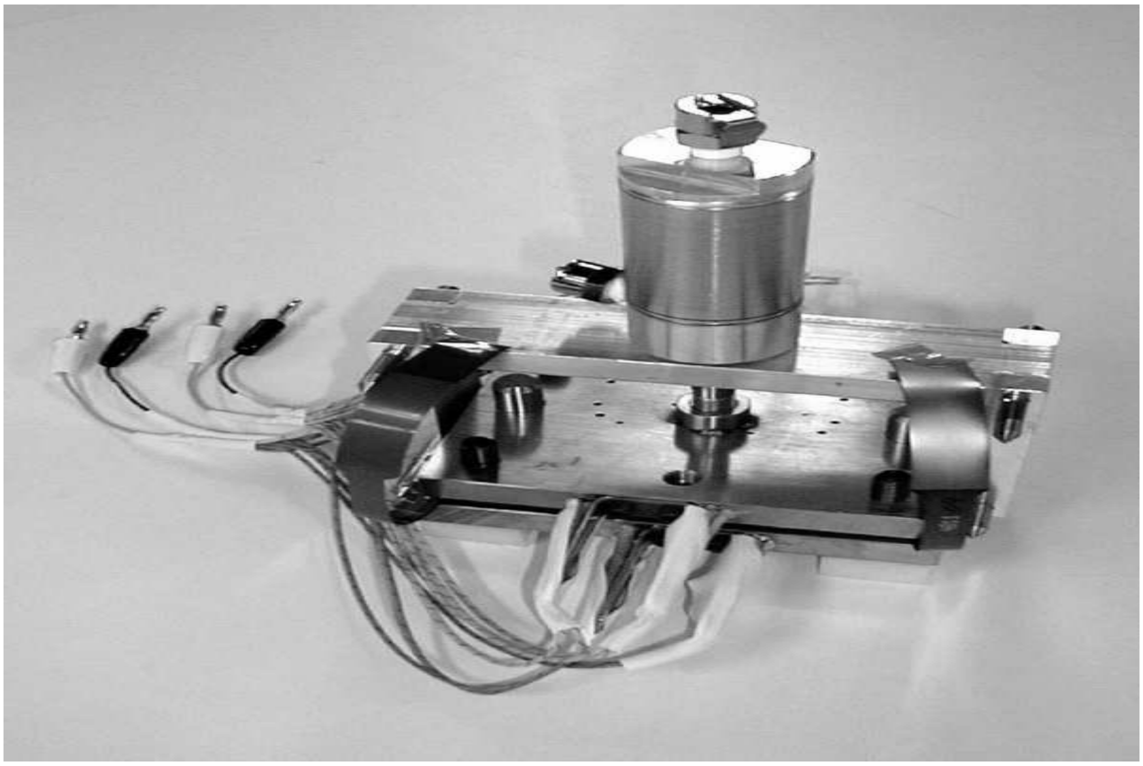

Figure 0.4: Fixture with heating elements that was inserted into vacuum jars for lamination.

Figure 0.5: HDI readout assembly mounted on carrier bar.

fixtures allowing two laminations to be done simultaneously. The setup and lamination process took approximately 4 hours so that we could produce about 4 laminations per day. After development the lamination yield was approximately 90%.

0.2.3 HDI Population

The HDIs were shipped to Promex[3] to: a) surface mount components on both sides, b) attach the SVXIIe die to both sides, c) wire bond the appropriate die to the HDI substrate bonds, and d) to attach the power and ground traces from one side to the other.

Because these units were double-sided, extensive development of fixtures and shipping boxes was needed. The HDIs were held for shipping and handing via a carrier plate mounted to the 6-chip side. This carrier plate mimics the cooling channel in the real detector. Figure 0.5 shows the HDI mounted on the carrier bar. The vendor based their wire-bonding fixtures on our shipping box design which held the HDIs by mounting pins similar to how the wedges are attached onto the rings.

Problems occured from the surface mounting procedure due to lack of precision in the control of the temperature for the surface mount oven. Several units bubbled on the second-side surface-mount step that were not recoverable. To try and limit this bubbling, several steps were taken: a) better oversight of control on the surface mount oven, b) a bake out before surface mount, and c) a visual inspection on the flex units as described above. Constant vigilance was needed at the vendor to insure that these quality control procedures were followed. Despite this vigilance, in addition to the unusable units due to bubbling, the company lost 14 units that were never found.

After completion of the procedures at Promex, units were shipped to Universities to be initially tested. There were several problems found with the quality assurance oversight of the vendor in terms of packing and shipping that also limited the yield. At the beginning of production, the yield for working units was only 50%. This yield went up to 85% later in the production process. Most of these units were later repaired and made to work. Due to the low yields and loss of parts, a reorder of parts was needed to complete the detector. This delayed production by about 4 months.

0.2.4 HDI testing

The F-wedge HDI testing was conducted in various steps using the electrical test setup described above. A preliminary verification and diagnosis of the readout via pedestals and injection of charge was conducted and subsequently the HDI was either sent to burn-in or to repairs for physical repairs or for further analysis. After burn-in the further verification of a detailed raw data check was performed. This further test was required as the SVXIIe chips were encapsulated at this point making any further repairs and diagnosis of the HDI difficult. The encapsulation was done by a skilled technician manually using a small opening syringe. The encapsulant was cured using an ultraviolet lamp. Subsequent testing was performed after the encapsulation to ensure that none of the wire bonds had popped during the encapsulation. At this point the HDI was glued to the silicon sensor and wire bonded to the sensor.

Early on it was discovered that there was a fold-over on the HDI which connected the AVDD2 power lines on the p and n side. It was concluded that this extra connection was redundant and could cause problems as the AVDD2 voltage could be set differently on either side which would result in a potentially damaging current flowing via this connection. Consequently it was decided to sever this connection on all the HDIs.

Another problem encountered was that of shorts appearing in the HDI tail of the traces on the tail to either the ground mesh or to adjacent traces. In the case of the low voltage supplies the HDIs would draw very high currents when powering up, although this was most often linked to an improper insertion of the HDI tails but could also happen due to chip failures. It was unclear whether the two effects were linked, i.e. whether the improper insertion of the tail or actual shorts between the traces could cause a failure in chips. In order to minimize this effect, especially in the later stages of production, the resistance between the low voltage supplies and ground was measured before powering the SVXIIe chips. The physical shorts on the tail were often due to fraying at the end of the tails or to contacts forming at the tip of the tail. The problem could be easily solved simply by cutting the end of the tail or by finely sanding the end. A further example would be the voltage supply tripping due to high currents even when applying as low as 1 V in negative bias. This would be linked with a short on the p side tail between the secondary bias line and ground.

Defects and failures in the chips with the subsequent replacement of these chips could cause further complications. Bonding on the same bond pads could be problematic especially after several iterations. A consequence of this was that bonds wouldn't stick or would pop after a short period of time. Also in replacing a chip the procedure would be to bond the chip to HDI bonds first and test the functioning of the HDI. Subsequently the chip would be then bonded to the sensor (via the jumper for the n-side). In this case there was the chance that the wedge would go through debugging with a chip not connected to the sensor. The general symptoms would be low noise (about 1 ADC count) in a extended number of consecutive channels together with a very flat pedestal.

The n-side of the HDI would have a further problem linked with the jumper. In various HDIs the jumper traces would be damaged or disconnected over a number of nearby channels. This showed up either as low noise or else as regions of variable noise together with a different pedestal than for the rest of the chip. The jumper traces could be easily verified by visual inspection. These channels could be connected by bump bonding on the traces themselves, however this wasn't possible if the trace was broken close to the bond pad. For the jumper traces which couldn't be fixed in this way the channels were disconnected.

0.3 Gluing process

Qualified Sensors and HDIs were selected to be glued together. The HDI is mounted on the same carrier bar that has been used for testing and shipping. This bar is also used to help handle the HDI for the gluing procedure. There are two important considerations for the gluing process: a) relative alignment of the sensor to the HDI, and b) the thickness of the glue joint such that good contact is insured. The process is performed on a specially machined fixture that is mounted on a CMM. This fixture is shown in Figure 0.6.

To control the thickness of the glue, the thicknesses of the HDI substrate as well as the silicon sensor must be known. We use the silicon sensor thickness as measured by the manufacturer and measure the HDI thickness directly. Once the thicknesses are known, the correct thickness guage blocks are chosen to mount on the fixture. The HDI is mounted on top of these guage blocks in the fixture on the CMM. A coordinate system is defined from the HDI which now is fixed on the gluing fixture.

The sensor is mounted on a flat pedestal and secured via a vacuum connection. This pedestal can be translated or rotated relative to the HDI using two micrometer adjustable screws. The plane height is fixed relative to the HDI (the z direction). The alignment of the silicon relative to the HDI coordinate system is performed using two fiducial marks on the sensor. The first fiducial is in the center of the sensor right next to the HDI, the second is at the centermost part of the wedge as it would be placed on a disk. An iterative procedure is performed using the micrometer screws to allow both fiducials to line up at the correct place. This alignment is done to better than 5 microns in this x-y plane.

The same epoxy was used to glue the wedges as the ladders. Care was taken to make sure that the mixing was done within the proper time window and the glue was metered out on the HDI kapton covered substrate. The HDI was held down onto the silicon using spring loaded screws and left to cure overnight. The next morning, the alignment was double checked before the silicon vacuum was released so that the new F-wedge could be dismounted from the gluing fixture

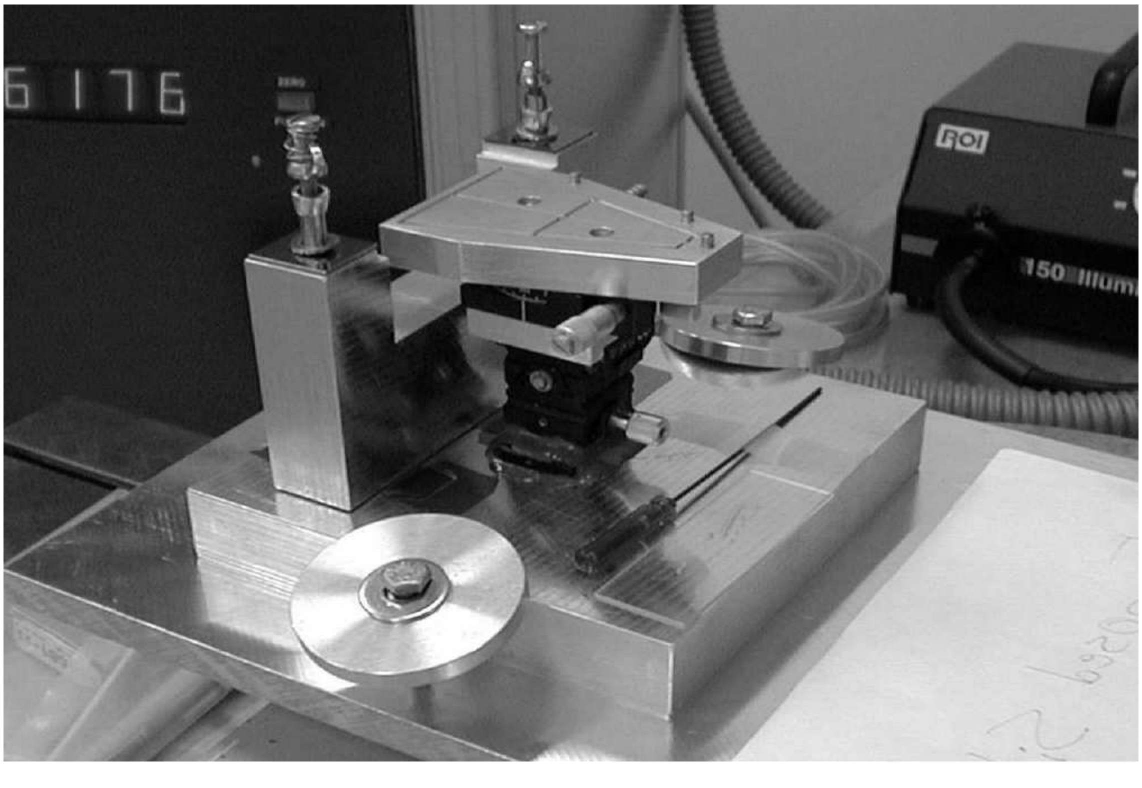

Figure 0.6: The fixture that is mounted on a CMM to allow alignment of the sensor to the HDI during the gluing procedure.

and stored in its handling box. Because there are bias bonding pads on the very corner parts of the wedge, an extra step was taken to inject glue at these points. A trained technician using a small bore syringe injected the glue here.

A total of 4 gluing fixtures were made so that 4 F-wedges could be glued per day. The alignment and gluing procedure for these wedges could be completed in one hour during the afternoon. The subsequent corner gluing and visual inspection took another hour during the morning. Once the procedures were developed, there were very few problems during gluing. The main problem was assuring that the sensor did not come into contact with the alignment pin on the fixture during dismounting. The yield for this procedure over the entire production ended up being over 90% .

0.4 Channel bonding

The F-wedges have a double-sided silicon sensor with a double-sided HDI with 8 chips on one side and 6 chips and a jumper on the other side. We had four production wire bonding operations for each wedge: a) bonding the 8-chip side bonds from the chips to the sensors (147K bonds), b) bonding the 6-chip side bonds from the chips to the jumper (110K bonds) , c) bonding the 6-chip side bonds from the jumper to the sensor (110K bonds) , and d) bonding the bias bonds to the detectors on both sides (1152 bonds). The total number of wire bonds for these wedges was about 370K.

The 8-chip (p-side) channel bonds were bonded using the KS8090 automatic wirebonding machine. A wedge on the fixture used for bonding is shown in Figure 0.7. This procedure took approximately 20 minutes per wedge. Minor problems were encountered with this operation only during development when the automatic bonder operating parameters were being determined. Care needed to be taken so that if a problem was found the bonder would not leave a row of tool marks more than 2-3 long on the silicon sensor.

We found that the wedge-to-wedge height differences and non-similar jumper surfaces made the wirebonding on this side more problematic. Because of the awkwardness in bonding to the jumper, both of the bonding steps on the 6-chip side of the wedge (n-side) were performed on the semi-automatic KS1478 machines. The wedge was placed into a vacuum fixture for these steps. Bonding parameters occasionally had to be adjusted during this process, but the bonding was accomplished with very few unbondable channels. The two wirebonding steps on the 6-chip side took on average 45 minutes to an hour to complete for each wedge.

The bias bonds were located on the corners of the sensor where the amount of glue underneath the sensor could contribute to broken sensors. Therefore the bias bonds were bonded using a manual bonding machine.

Aside from the wirebonding technician, we had another technician help to

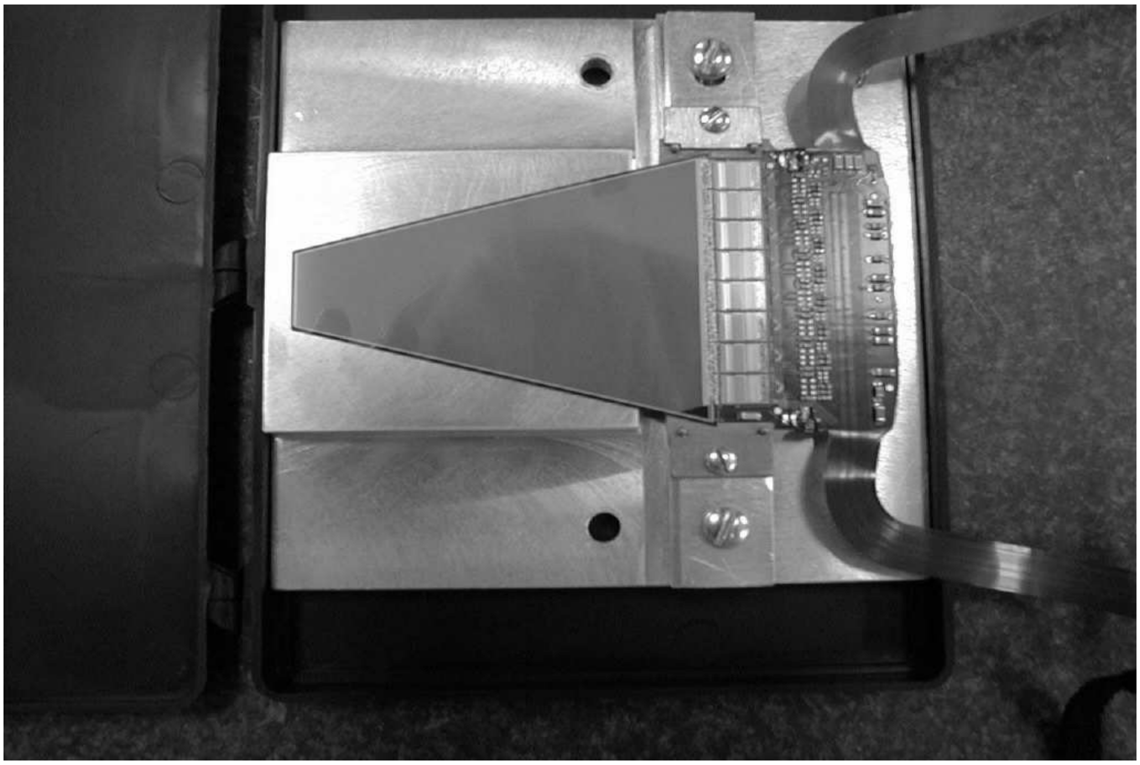

Figure 0.7: An F-wedge mounted on the 8-chip side wirebonding fixture.

mount and dismount wedges from the fixtures and perform the quality assurance procedures. The majority of problems encountered were due to mishandling of the wedges in the fixtures. Only one misalignment on the wirebonding machine broke a wedge. We encountered only very minor problems with quality assurance of the bonding itself that included tool marks and some wire bonds that popped after bonding.

0.5 Wedge Debugging

Each wedge was tested after bonding with bias voltages applied. A description of the setup used is described above together with the standard procedure for the diagnosis of the common problems encountered in the debugging of the ladders and wedges. The procedure for the debugging of the F-wedges was similar to that for the ladders but had various extensions added due to the specific characteristics of the wedges. Due to the nature of the design of the HDI tails it was more likely that a bad connection to the Hiroshi connector was made than for the ladders. Therefore the currents had to be monitored carefully when powering up the SVXIIe chips using the low voltage supplies.

When applying a positive bias from the high voltage supply the wedges were biased to 100 V or the voltage at which the micro-discharge region was reached, whichever was lower. This voltage was generally higher than that applied to the ladders. The justification for this was that it was unclear at the time of testing as to the period of time required for the wedges to become type inverted. This combined with the uncertainty of the design of the runIIb silicon tracker meant that the wedges might have had to be biased to high voltages in order to be fully depleted after a couple of years of running. Also the wedges, especially the ones produced by Micron, had a tendency to fail burn-in as pinholes would occur during the first few runs tripping the high voltage supply. In order to limit the failure rate at burn-in the wedges were left biased for 30 minutes at high voltage (equivalent to the voltages used in burn-in) and the currents were monitored. If they passed this test they were then sent to burn-in. Although his did not prevent pinholes from occurring during burn-in, it did limit the frequency of pinholes occurring and generally reduced the number of burn-in iterations needed for each F-wedge.

A further step which was employed in the case of download problems was to do a visual inspection of the chip bonds. For other problems a visual of the bonds linking the chips to the sensor (in the case of the n-side the chip to jumper and the jumper to sensor bonds), together with a visual of the silicon sensor were performed depending on the nature of the problem. It was found that often download problems of the HDIs were associated with missing and popped bonds on the chips and these could be diagnosed and re-bonded without having to go through the more time-consuming repair process.

0.5.1 Common Problems and Diagnosis

The F-wedges had many problems which were distinct from the other SMT detector elements. Some of these problems were general to all the wedges and others were specific to the wedges from one of the manufacturers. The vast majority of the general problems were linked with faults in the HDI manufacturing and problems resulting from chip replacements. These problems are described above.

The micron sensors were found to have two main problems: regions in which pinholes would seemingly propagate to adjacent strips and problems related to the stability of the currents when biased. The first problem had a very distinctive signature, the wedge in positive bias would exhibit a high current and the diagnosis would clearly locate a pinhole. This channel would be disconnected and when tested again under bias initially the currents would be lower and then one of the adjacent channels would seemingly develop a pinhole. This channel would be disconnected and the pinhole would seemingly propagate to the next strip. This area could extend to 10-15 consecutive strips or even more for the worst cases. When examined under a high power microscope p-stop implantation faults could be seen in the areas affected, however this correlation wasn't complete as some wedges affected didn't exhibit visible faults. The general impression though was that this problem was linked to faults in the masking procedure of the silicon. A solution which generally seemed to limit the area affected was to pull an extra bond than needed when a region started to develop. This was also done to try to limit the number of iterations required for each wedge.

The second problem encountered with Micron sensors had to do with the reproducibility and stability over time of the bias currents. The stability was an issue in negative bias as the currents seemed to increase over time as described in the burn-in section. As a result, the definition of a safe operating voltage in negative bias would be modified to a lower voltage which seemed to reduce the drift. The problem in the definition of a safe voltage in secondary bias was linked to the Micron sensor characteristics. The I-V curve in negative bias would not exhibit a clear junction breakdown, although above a certain voltage the measured current would tend to increase more rapidly with time with an increase in voltage. Once this point had been reached the currents measured previously at a lower voltage would be higher as well. As a result the decision as to the voltage at which the junction breakdown occurred was somewhat arbitrary. It was consequently defined based on the voltage at which the current wouldn't drift significantly (more than 5 - 10 μA) over the period biased together with a decrease in the rate at which the current would increase with time. This effect was seen also in some Eurysis wedges but the effect wouldn't be as pronounced and the I-V curve always had a clear junction breakdown voltage.

The problem which was observed with the reproducibility was that the currents (especially in positive bias) could suddenly jump by a factor of 4 to 5 without showing any indication of pinholes or other signs for such a high discontinuity.

Table 0.1: Average Number of Bonds Pulled for the Various Sensor Types in Debugging

Sensor Type	n-side	p-side
Micron High (MH)	5.5	3.3
Micron Medium (MM)	5.5	2.8
Eurysis (E)	4.3	12.0

The currents would be stable but much higher. This variation would not always be permanent and sometimes after leaving the wedge in a dry box for a couple of weeks the wedge would have an I-V curve similar to that before the effect was seen. However in a limited number of cases (about 4) the currents remained high and stable. In total the number of wedges which exhibited this were less than 15. The measured current under positive bias was much more stable and had little tendency to drift even at high voltages (90-100V).

The Eurysis sensors generally exhibited less problems than the Micron sensors. While the Micron sensors generally had problems on the n-side, the Eurysis sensors exhibited a larger number of pinholes on the p-side. This was correlated with the fact that the wrong set of bond pads had been used to probe the sensors, thus the chips were bonded to the pads which had probe marks. The I-V characteristics were also more defined as a clear junction breakdown was visible when applying negative bias. The only real problem encountered was that the wedges would generally exhibit a micro-discharge voltage at around 70 V in positive bias, which signified that the wedges could not be biased to as high a voltage as would have been ideal. As the Eurysis sensors were low depletion devices this could only become an issue when the F-wedges become type inverted and the depletion voltage increases past the micro-discharge voltage.

0.6 Wedge Burn-In

One of the fundamental tests which all the SMT components had to pass was an extended period of continuous read-out under a DØlike operating configuration, in regards to readout and electronics, temperature and humidity and bias voltage.

Initially the burn-in had many problems in trying to download the SVXIIe chips of the F-Wedges. The reason for this difficulty was probably as a consequence of various factors and was not well understood, especially as the other detector types didn't encounter the same problems. Initially when trying to download the F-Wedges at least 20% would trip the low voltage supplies or not download. This wasn't a problem in itself but the low voltage fuses on the interface boards were set to a lower current than the low voltage trip on the power supplies. This meant that each time a burn-in was attempted between 2-4 interface boards would have to have fuses replaced. The explanation for these problems was linked to

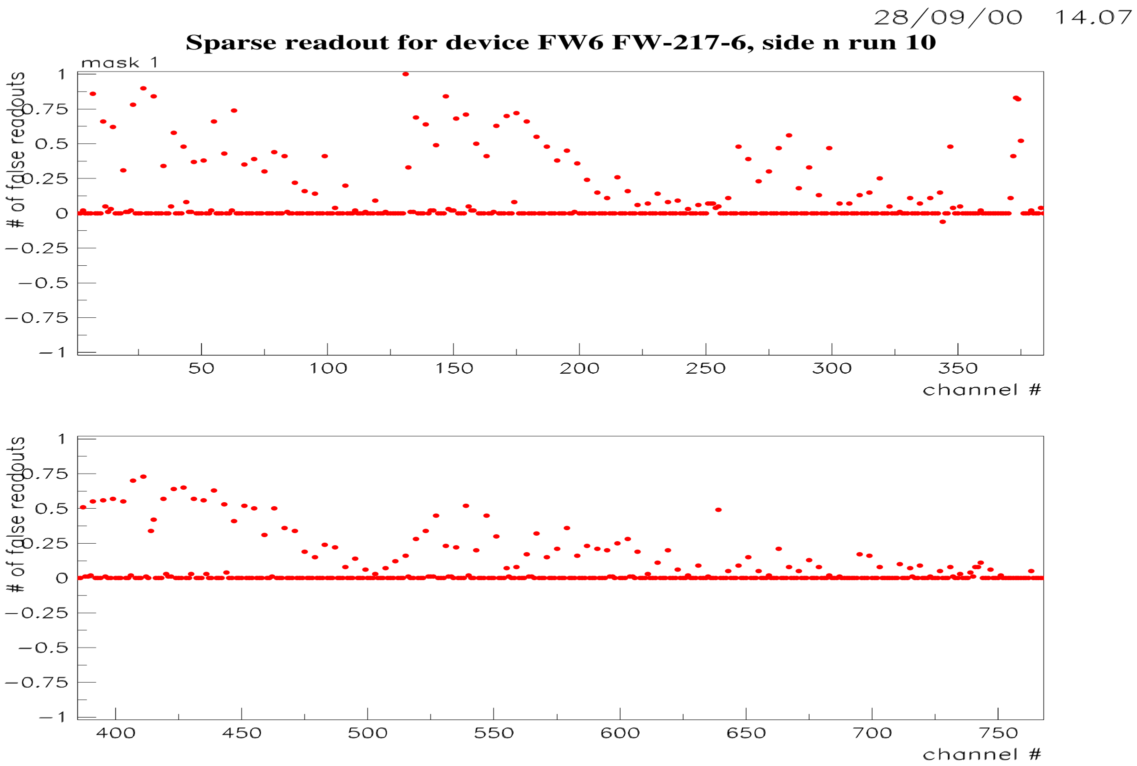

Figure 0.8: 6 Chip Sparse Readout During Burn-in

multiple factors. One factor was probably some degree of tester inexperience as the insertion of the HDI tails was more complicated than that for the other detectors [see debuggin section]. A further component was that the F-Wedges, having two HDI tails, provided a connection between the two channels on the interface boards. Theoretically there should be no coupling between the two sides of an F-Wedge (except for a common ground) and also the two interface board chains should be decoupled, however the extent of the decoupling wasn't clear. A further issue was that the spreadsheet used to download the F-Wedges in burn-in was different to that used in electrical debugging and probably some of the bad downloads seen in burn-in and not in debugging were linked to this difference. With time modifications to the interface boards (such as a larger fuse for the low voltage supplies) and to the spreadsheet meant that the download problems stopped being a major issue although they still sporadically occurred.

There were problems which did not improve with time and remained an issue throughout the whole period of testing. They can be classified into two broad categories, readout problems and high voltage problems.

The main issues with the readout were, for the most part, problems with the interface boards. Four different errors were observed on a regular basis. The primary was that the sparse readout would nearly always be very irregular for the 6 chip side, an example of this behaviour is shown in Figure 0.8 with periodic variations on a chip-to- chip basis. This feature would not be seen for the HDI burn-in. The reason for this was imagined to be linked with the fact that in burn-in the low voltages were optimized for the first chip on the 8 chip side and not for the 6 chip side as there was only one low voltage power source for both chains. Whereas this would not have an impact for an HDI the presence of the silicon sensor would amplify this difference and cause the sparse to be bad. A check was performed using the same spreadsheet as the one in burn-in on another test stand using various wedges. The results showed the sparse readout to be acceptable when the 6-chip side was powered up and downloaded on this test stand.

The other features often seen were all linked to problems with the interface boards used in burn-in. For example the gain for chips was monitored to be 0 although the chips would respond correctly to an injection of charge. Also the noise pedestal would be displayed as 20 ADC counts for certain chips but which chips displayed this and the total number of such noisy chips would change from run to run. A similar effect was seen with the pedestals under charge injection, although in this case the pedestals would fluctuate from their normal value to 0 on a run to run basis. The last 2 problems would often be seen on the same wedge during burn-in.

There were various problems seen in burn-in related to the bias currents. There would be a tendency for the wedges to develop pinholes and trip the high voltage supplies even though the debugging procedure would test the wedges under higher voltage. This would usually occur during the first run of the burn-in. However this could occur also while applying the voltage to the sensor (even at 5 V). This lead one to believe that some times the power supplies would have a surge in voltage when turning them on which could blow the capacitors. This would not be worrying in itself as one of the functions of burn-in is to make sure that the wedges are in a stable state. However only one capacitor would fail at a time, making some wedges go through the burn-in procedure multiple times and also a couple of wedges developed high leakage currents after burn-in.

The leakage currents also had a tendency to increase over the course of the burn-in cycle as shown in Figure 0.9. This was more pronounced for the Micron wedges even though the Eurysis sensors would also show this to a lesser degree. The currents would usually increase over the course of the first few runs and then plateau to a stable value for the rest of the test. There were various wedges however which increased at a relatively constant rate throughout the whole burn-in even though there were no signs of bad capacitors. The increase was regarded to be temperature dependent as the temperature registered in burn-in was higher than the theoretical value (20 degrees compared to around 3). This was confirmed by various tests which showed that the increase was less significant when at a lower temperature. There were some Micron wedges which even at a lower temperature would still have increasing bias currents and seemingly not plateau.

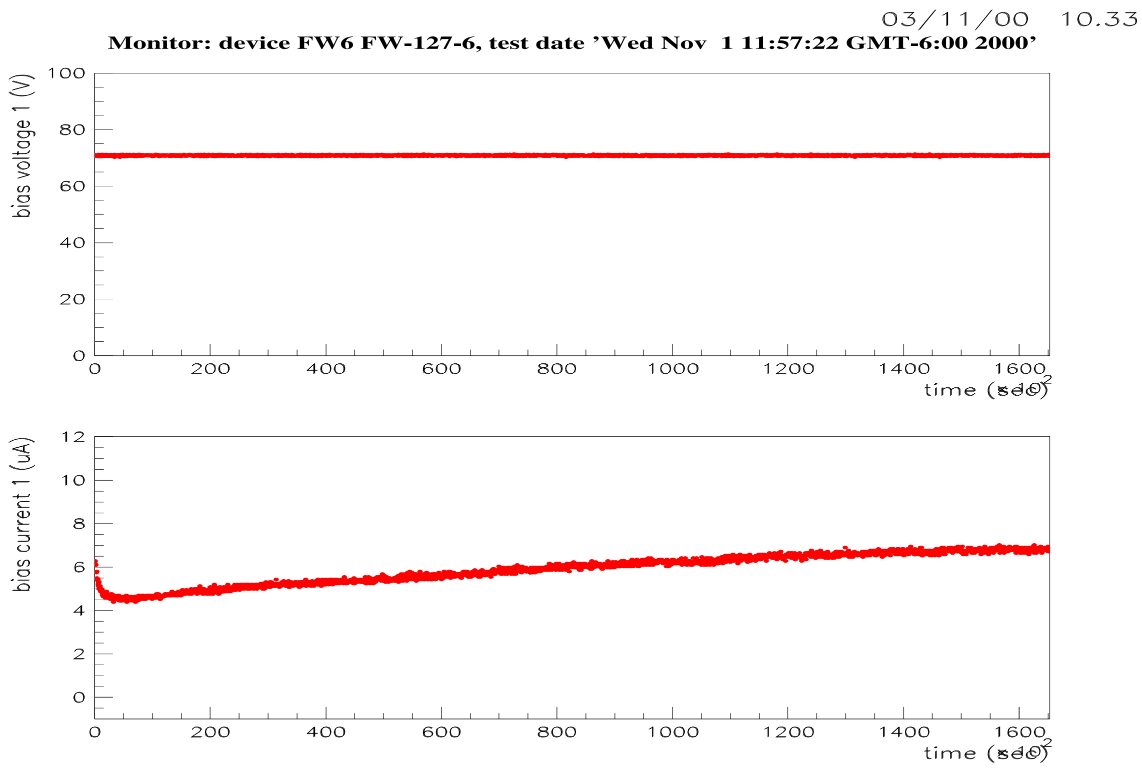

Figure 0.9: Bias Currents for a Micron Sensor During Burn-in

Table 0.2: Typical Values of some Parameters Monitored Duing Burn-in. The label W refers to wedges and A to HDI. The numbers 6 and 8 refer to the number of SVXIIe chips on the HDI tail being readout. A range is given for the gain as the gain profile had a slope starting at the lowest channel and rising linearly with channel number to the final value.

	FW-6	FW-8	FA-6	FA-8
Pedestal (ADC counts)	100	100	100	100
σ_{ped}	2	2	1	0.75
Charge Inject (ADC counts)	110	100	100	100
σ_{cal}	2	2	1	1
Gain (ADC / fc)	8-12	9-12	12-14	13
σ_{gain}	0.5	0.25	0.25	0.25
AVDD (mA)	330	400	340	350
AVDD2 (mA)	140	180	140	180
DVDD (mA)	40	60	40	60

Table 0.6 gives some typical values found for the parameters monitored during burn-in.

0.7 Ring Assembly

The F-disk cooling rings are delicate assemblies, which serve as the support structure and heat sink for silicon detector electronics assemblies. The cooling rings, which are made of beryllium, were machined in (2) halves and then adhesively bonded together to form a ring with a 1 $mm \times 7.5$ mm cooling channel. The bonded ring is 2.5 mm thick with an average diameter of 240 mm. Final machining was performed on the bonded ring to produce tooling holes and tapped holes later fitted with titanium studs used to hold the wedge assemblies. Carbon fiber posts and fiducial balls supports were bonded to the ring face to provide mounting and measurement features. Two general styles of rings were assembled. One style attached to a barrel bulkhead and another style attached to other rings to form the 3-ring end-ring assembly. Figure 0.10 shows a photograph of the beryllium ring assembly.

Hardric laboratories of Hudson, MA machined the beryllium rings. The rings were made from S-200F hot pressed beryllium block. The rings were machined in two stages. In the first stage, ring halves with the final inner and outer perimeters were machined. The ring halves were bonded together at Fermilab and then returned to Hardric laboratories for the second stage. In the second stage two tooling holes and 36 tapped holes were machined. The first stage machining presented the most difficulty. The local and overall flatness of the rings were the

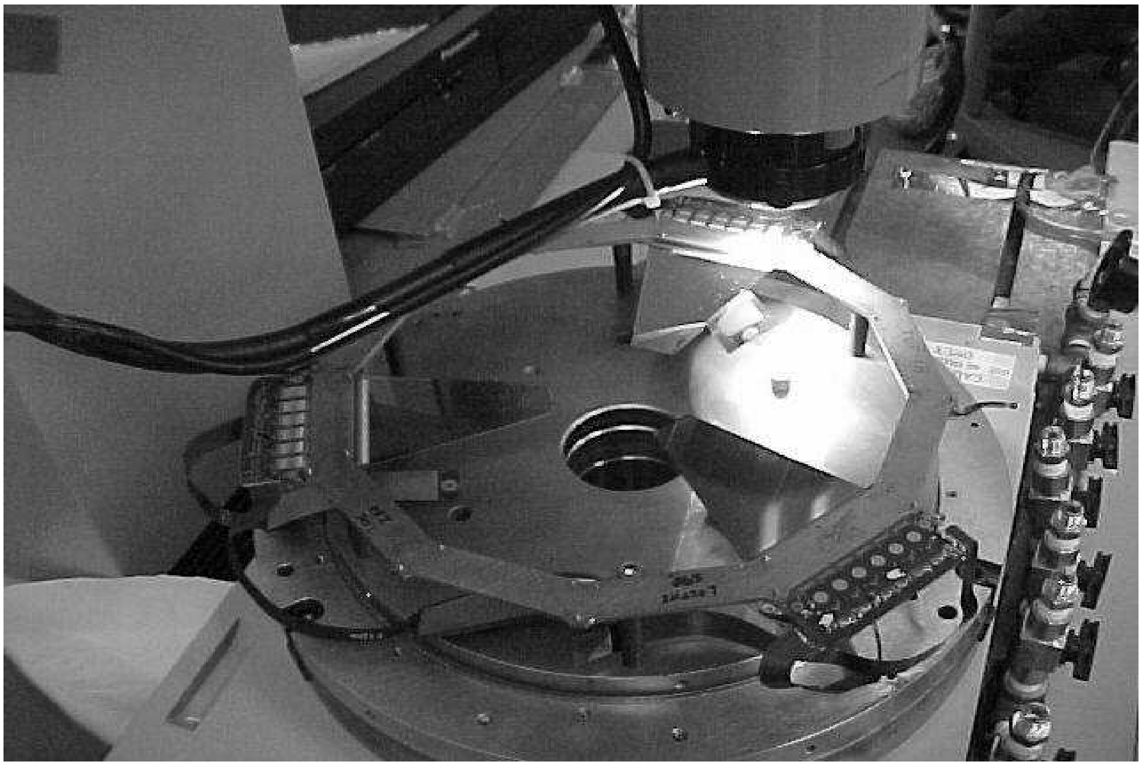

Figure 0.10: Photograph of the beryllium F-ring with 3 wedges mounted on it.

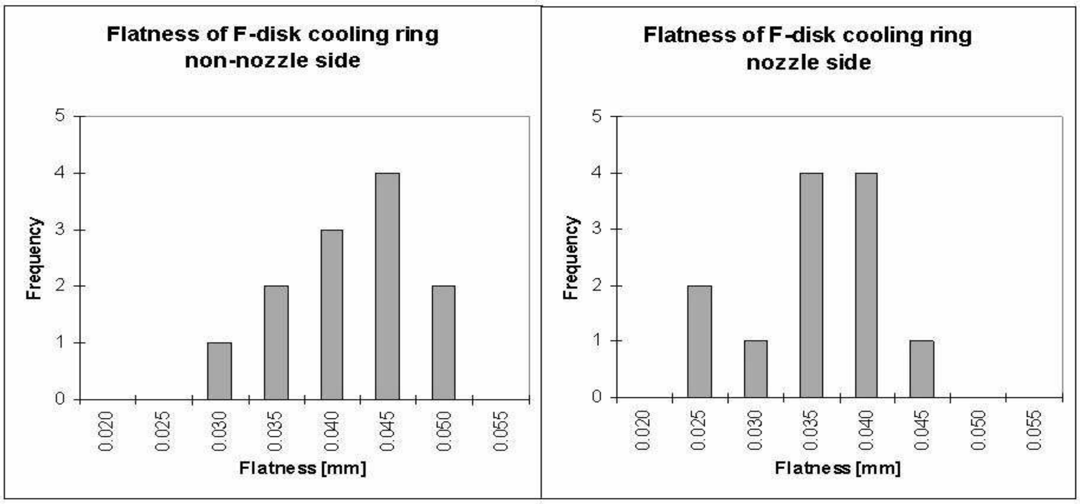

Figure 0.11: Distributions showing flatnesses of beryllium halves used in the ring assembly.

most important features and the most difficult for Hardric Laboratories to achieve. The overall flatness was specified as .07 mm and the local flatness was specified as .02 mm for an area the size of a single wedge. After a few trials Hardric laboratories developed a fabrication sequence that met specifications. The sequence involved initial heat treatment, milling and lapping, and a final flattening heat treatment between beryllium platens. Any machining, including deburring, after the final heat treatment would cause distortion of the thin 1.25 mm rings. The rings halves had two 1.5 mm diameter holes that would later be used to align two halves for bonding. Half of the rings had two holes where the inlet/outlet nozzles would be attached later. A dye penetrant inspection to check for cracks was performed as the last step of the first stage. No significant problems were encountered in the second stage machining. A beryllium surface passivation treatment was performed as the last step of the second stage. Figure 0.11 shows the distribution of flatnesses for each beryllium half of the assembly. They range in flatness from $20 - 40\mu m$.

The two primary requirements of the bonding operation were to obtain a quality leak free joint and to produce a full ring with an overall flatness of less than .07 mm. The bending stiffness of a complete ring is 8 times greater than the bending stiffness of ring half. If the ring halves could be bonded in a flattened state then spring back from that state would be minimal. The rings halves were bonded with .075 mm thick Ablefilm 563K B-stage epoxy film adhesive. This film adhesive requires an oven cure at 120 C for two hours. Squeeze out of the film adhesive was controlled by distributing glass beads on the film adhesive around the bond areas. The glass beads were sieved to a .038-.043 mm range. Initially .05 mm thick film adhesive was tested but the local waviness of the test rings produced an unsatisfactory bond joint with this film thickness. The cure characteristic of the film adhesive had been studied with a microscope and an instrumented heat platen. This study revealed that the film adhesive became soft at 80 C and cured to a hardened state within 10 minutes of reaching 120 C. Trapped gas or evolved gas expanded during the soft state and locked in size shortly after reaching 120 C. The cure cycle was chosen to eliminate or minimize the size of any gas bubbles.

The equipment developed for bonding the ring halves involved a laminating press with two volumes separated by a plastic sheet barrier. The two volumes could be independently evacuated or pressurized with any differential pressure displacing the plastic barrier. The two ring halves with film adhesive and glass beads were placed on the flat platen of the laminating press. Two 1.5 mm titanium pins aligned the ring halves. The barrier plastic and laminating press cover were positioned and sealed. The laminating press was used in a convection oven with a port that permitted thermocouples and two gas lines to feed to the outside. At the beginning of the cycle both volumes of the lamination press were evacuated. The press was heated to 80 C in 45 minutes while the interior was evacuated. The top volume was vented thus applying pressure through the plastic barrier to the cooling ring in the second volume. The pressure of both volumes was increased in stages as the temperature rose. Through the balance of the cure cycle the top volume pressure was maintained at .19 barr with the bottom maintained at 1.3 barr. The .6 barr differential pressure was used to flatten the rings. The overall flatness of the bonded rings were .04 mm or better on the face against the platen and .05 mm or better on the face away from the platen.

After a final inspection of the stage 2 machined rings the ring assembly process began. First, short titanium threaded studs were installed and bonded into the ring. A wicking Loctite 292 anaerobic adhesive was used to bond the studs into the ring. Next carbon fibers post and ruby fiducials holders were bonded to the face of the rings. Six carbon fiber posts and six 3 mm ruby fiducial balls holders were bonded to one face of each cooling ring. Two styles of posts were used. One style would attach to a barrel bulkhead and another style would attach to another ring to form the 3-ring end ring assembly. For cooling rings that attached to a barrel bulkhead the carbon fiber posts were 35 mm long with a 4 mm diameter aluminum tip insert that would engage holes on the cooled bulkhead. For cooling rings that would attach to another cooling ring, the posts were 23 mm long with a cone shaped aluminum tip that would engage a 4.75 mm ruby ball. The ruby ball holders were 1 mm thick cantilever arms that positioned a 3 mm ruby ball in between wedge detectors near the outer radius of the ring assembly.

For the attachment of either style of post the cooling ring was clamped to a fixture base and aligned to the base with pins through the two tooling holes. For one style of ring the posts were positioned with a top fixture plate that simulated the engagement holes that the posts would eventually mount to on the barrel bulkhead. The engagement hole-to-hole alignment of the top fixture plate was the most important feature, as the wedges would later be aligned to a best-fit center of the posts. For the second style of posts small inserts with 4.75 mm balls were mounted in the top fixture plate. For bonding to the ring, the posts were clamped to the top fixture plate. The top fixture plate was positioned and aligned relative to the ring mounted to the base plate. Ground spacer blocks controlled the gap between the top plate and base plate. The posts hung like bells above the ring with a .075-.125 mm gap between the end of the post and the face of the ring. The gaps were filled with a structural epoxy 3M 2216 Grey and allowed to cure at room temperature overnight. The positioning of the posts as a group relative to the ring was not as critical as the post-to-post alignment. Ruby ball fiducial holders were bonded with the same structural epoxy. The ruby balls were positioned in small V-shaped cups located on the fixture base. The alignment of the ruby fiducials was not critical.

0.8 Disk Assembly

Each F-disk has a fluid cooled Be ring, on which 12 detector elements, wedges, are mounted, 6 to a side, with a little overlap. Wedges are built up on a Be cooling plate and incorporate the silicon sensors and the readout electronics. The cooling plate extends beyond the sensor where it meets the cooling ring. These "Ears" have two holes on each side. The holes match tapped holes on the cooling ring, which hold 2mm dia. Titanium threaded rods for attaching the wedges. The holes are 0.24 mm oversize to allow adjustment of the wedge positions.

The wedges are assigned positions on each disk determined by their operating bias voltages as described below. Before wedges are mounted on rings, the bias bonds on the corners of the sensor are encapsulated for protection. This encapsulation is done by hand using a syringe to apply the encapsulant. The encapsulant is heat cured for an hour. We started off encapsulating all of the channel bonds on the sensor using the UV cure encapsulant used to encapsulate bonds from the chips to the HDI (discussed earlier). Each qualified wedge is retested immediately before placement on the ring to insure that it is still functioning correctly. Leakage current tests indicated that there was some problem with the interaction of the encapsulant with the double-sided detector so this step was eliminated for all wedges but those mounted on the first disk.

F-disk positions are referenced by a set of 6 sapphire balls mounted on stalks at the periphery, in the small gap between readout electronics plates. The balls are touch probed by a CMM through holes in the support cylinder for. The ball positions are known accurately in the system defined by fiducials on the 12 wedges. The locations of the 6 balls were originally measured in a system centered on the mounting rod locations, to assure adequate room for alignment of wedges. This was done on a touch probe CMM. When installing wedges, the same 6 balls are measured optically on a Zeiss UPM 850 CMM. This is done by measuring about 6 to 8 points around the equator of each ball. It is not easy (and, in principle, not possible) to focus on that periphery, as each off-axis ray is tangent to the sphere at a different focus depth. We achieved excellent results by using mostly vertical on-axis rays. The through-the lens illumination was reflected by a piece of Scotchbright reflecting material, which sends light back to where it came from. The ball positions differed by only 2 to 3 microns, typically from those measured earlier by touch probe on a different CMM. Fore every new setup we re-measured the 6 balls, and fit a planar coordinate system (X, Y, rotation) to minimize the offsets of the 6 balls from their nominal locations. The remaining offsets were about 2 to 3 microns rms.

Typically we mounted two opposing wedges at a time, using heat coupling grease and aluminum nuts that fit the 2mm Ti rods. The Be cooling plate was so crowded with components that there were no places left to grab it with a positioning device. Hence, all alignment was done by a skilled technician, who developed methods to push and rotate the Be plate with prods and small screw drivers while watching the CMM cross hair, while the nuts were snug up just enough to provide a little friction. Each wedge is aligned using the same two sensor fiducial marks as used for gluing the sensor to the HDI. When the alignment was satisfactory (we required less than 6 micron offset in any direction on any fiducial), the screws were slowly tightened using a pre-set torque nut driver. We installed and aligned/tested about 2 wedges per day. Installed wedges were electronically tested whenever a few had been installed. While the yield for wedge placement was high, some wedges were destroyed by slipups in this process that needed to be replaced on the disks. Due to time constraints, another method could not be developed that would have ensured more safety in the procedure. However, during the disk assembly, most of the wedges were covered with a plexiglass plate.

Completed F-disks were re-measured, using machine vision, on an automatic optical CMM (made by OGP) to get an independent verification of the wedge placement accuracy. We found that wedges were placed typically with an average offset of 4 to 8 microns, and a scatter of about 6 microns rms. In addition to the placement errors by the installer, these numbers include the errors of measuring the fiducial balls on both machines, measuring the fiducials on the silicon, and random and systematic machine errors. The OGP also measured out-of plane deviations, which were mostly due to non-flatness of the silicon sensors. The inner tips of the wedges were bent away from the active surface by 450 ± 75 microns.

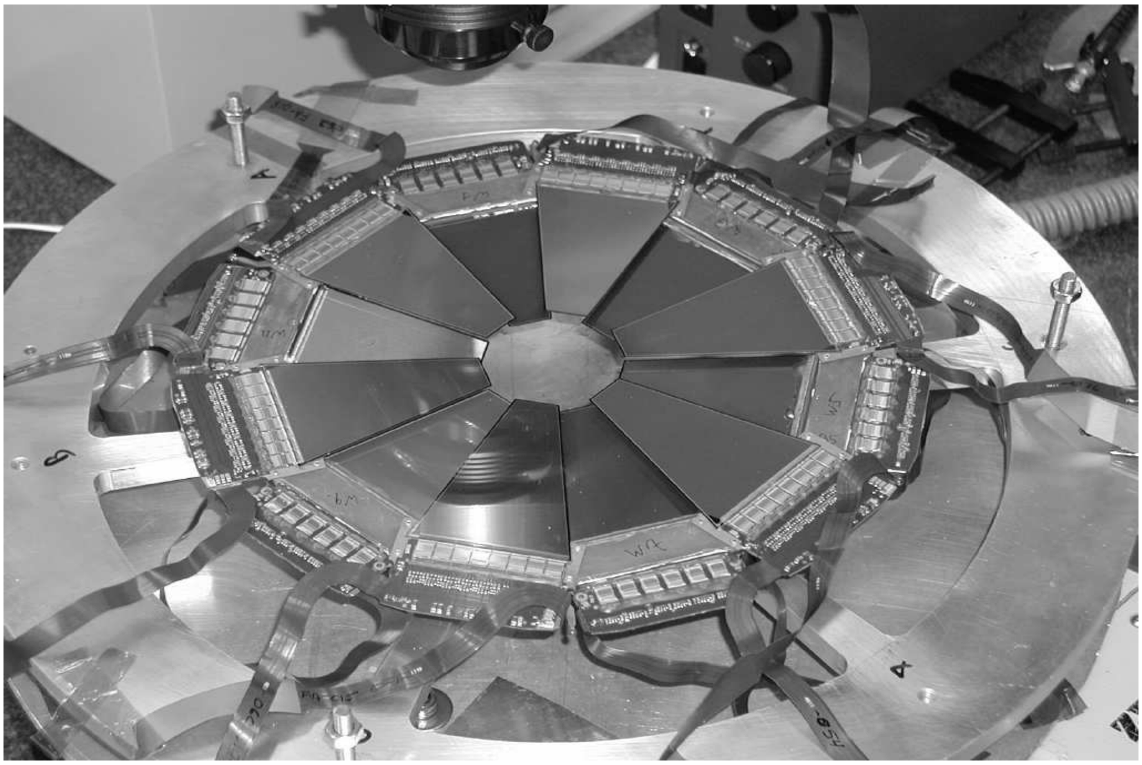

Figure 0.12: First prototype F-disk Assembly

Working disks were dismounted from the CMM disk assembly fixture and placed in individual boxes. These boxes were made from static free plexiglass and included grounding straps for extra static protection. The disks were stored in a dry box until they were ready for mating with a barrel or with another end-disk.

Before this mating procedure, the disk would be tested again. For the disks mating with barrels, a dimpled 1 mil kapton diaphram was carefully mounted over the posts to try to minimize heat transfer from the disks to the barrel detectors mounted close by. The HDI tails were carefully dressed into a pinwheel design to hold them for their mating.

0.9 Overall Performance

All 12 disks were assembled and delivered in time for assembly into the detector. Figure 0.12 shows the first prototype disk that was completed.

After disk assembly there were a total of 1521 dead channels on the p-side (1.03 percent) and 1486 dead channels on the n-side (1.15 percent). Figure 0.9 shows the distribution of dead channels on the p and n sides for all of the wedges. The highest number of dead channels was 4.9% for one wedge on the p-side and also for another one on the n-side. Most of the dead channels were due to leakage

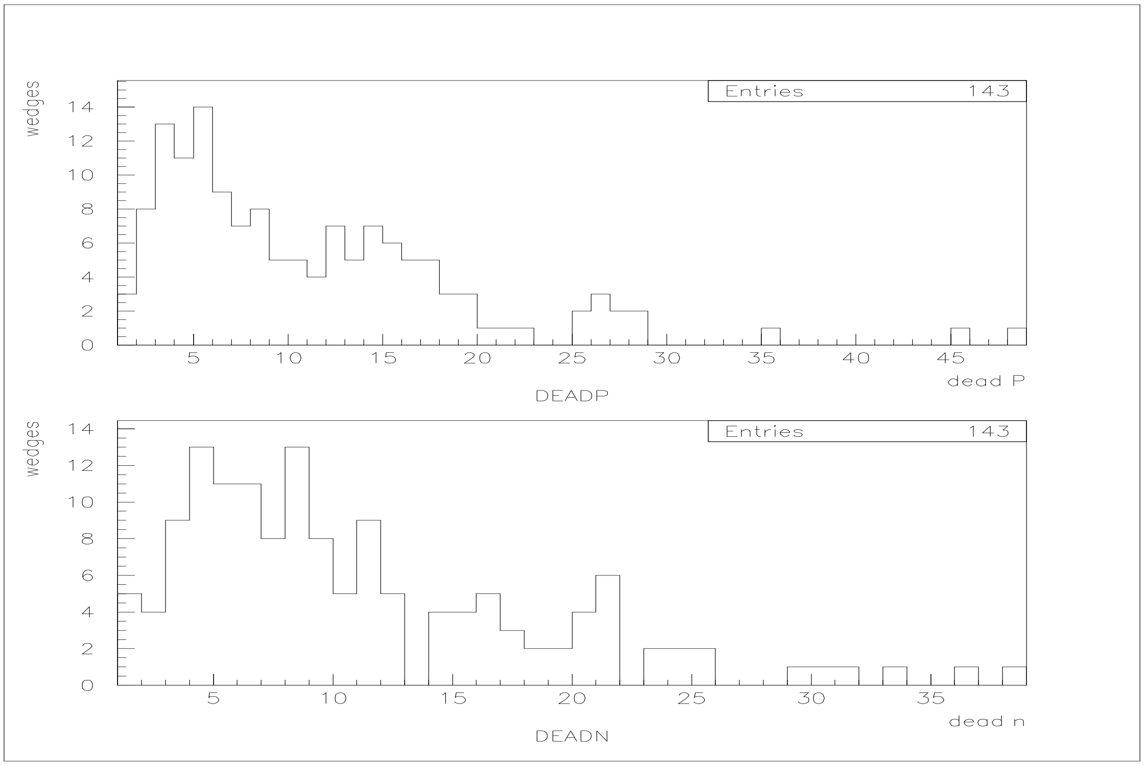

Figure 0.13: Number of dead channels per wedge for a) the p-side (total of 1024 channels per wedge) and b) the n-side (total of 768 channels per wedge).

current problems and not wirebonding errors.

A noisy channel is defined to be one where the sigma on the pedestal is greater than three ADC counts and is mainly due to problems with the HDI. This is described elsewhere in this paper. The number of noisy channels was determined to be 419 on the p-side (0.28 percent) and 416 on the n-side (0.32 percent). The noisiest p-side has 22 channels, while the noisiest n-side has 14 channels. Figure 0.9 shows the distributions for noisy channels.

The disks were assembled so that each disk had wedges with similar bias voltage requirements. We ended up with three different classifications of disks: those built with Micron high depletion voltage wedges (MH), those built with Micron medium depletion voltage wedges (MM), and those built with Eurisys wedges. These disks were distributed so that they are symmetric about Z equals 0. The two end-disks farthest away from Z=0 are constructed using MH as is the third disk from z=0. Micron medium depletion voltage wedges went into the disk closest to z=0. Eurisys wedges went into the remaining disks. Each disk has three high voltage groups so that four wedges are supplied the same bias voltages per disk. The voltage is applied as a split bias voltage so that there is a maximum negative voltage and a positive operating voltage determined using the depletion

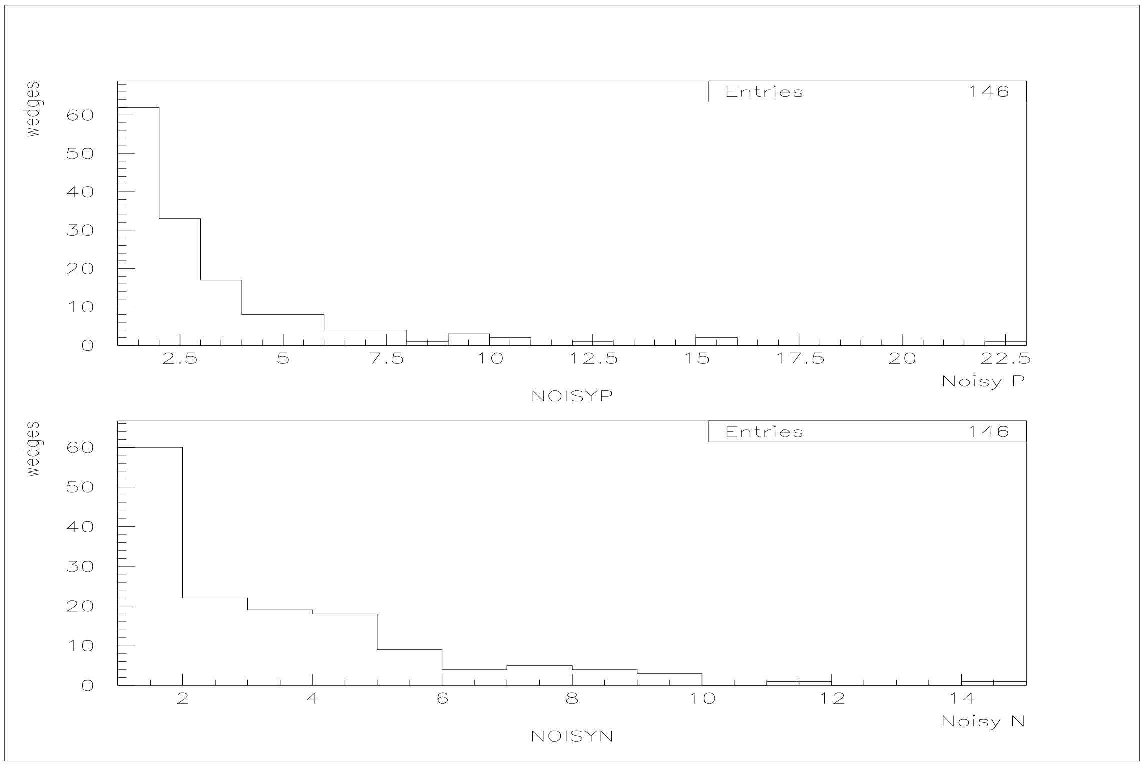

Figure 0.14: Number of noisy channels per wedge for a) the p-side (total of 1024 channels per wedge) and b) the n-side (total of 768 channels per wedge).

Table 0.3: F-disk operating operating voltages in Volts. The wedge types are: MH - Micron High Depletion, MM - Micron Medium Depletion, E - Eurisys.

Disk No.	**Wedge type**	**Group 1**		**Group 2**		**Group 3**	
		V-	**V+**	**V-**	**V+**	**V-**	**V+**
1	MH	20	50	20	60	10	80
2	MH	10	70	5	80	25	65
3	E	20	25	15	30	20	35
4	MH	30	50	20	55	5	75
5	E	20	25	20	35	15	40
6	MM	15	35	20	30	10	45
7	MM	10	55	15	35	10	40
8	E	15	45	25	25	20	30
9	MH	10	75	15	65	30	50
10	E	20	25	15	35	20	30
11	MH	20	60	20	70	15	60
12	MH	5	75	30	50	20	60

voltage information. The values of these voltages are determined from by the wedge tests described above. Table 0.3 gives these negative and positive voltages for each of the 3 high voltage groupings in each of the disks. These operating voltages are also displayed in Figure 0.9.

With these operating voltages, the leakage currents for each wedge have been determined. Figure 0.9 shows the distribution of leakage currents for the wedges mounted in disks. Most of the wedges operate with less than 20μA of current although we have a couple of wedges with over 100 μ A of leakage current.

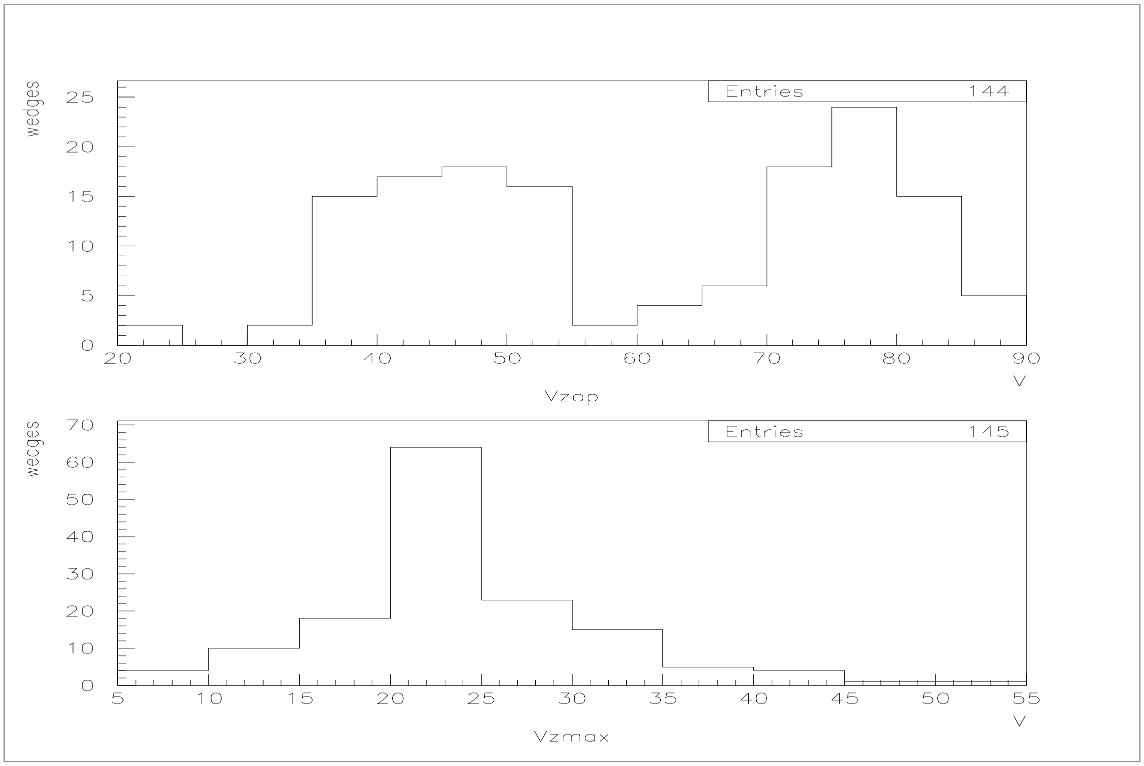

Figure 0.15: Number of wedges with a) the positive bias voltage, and b) the maximum negative bias voltage.

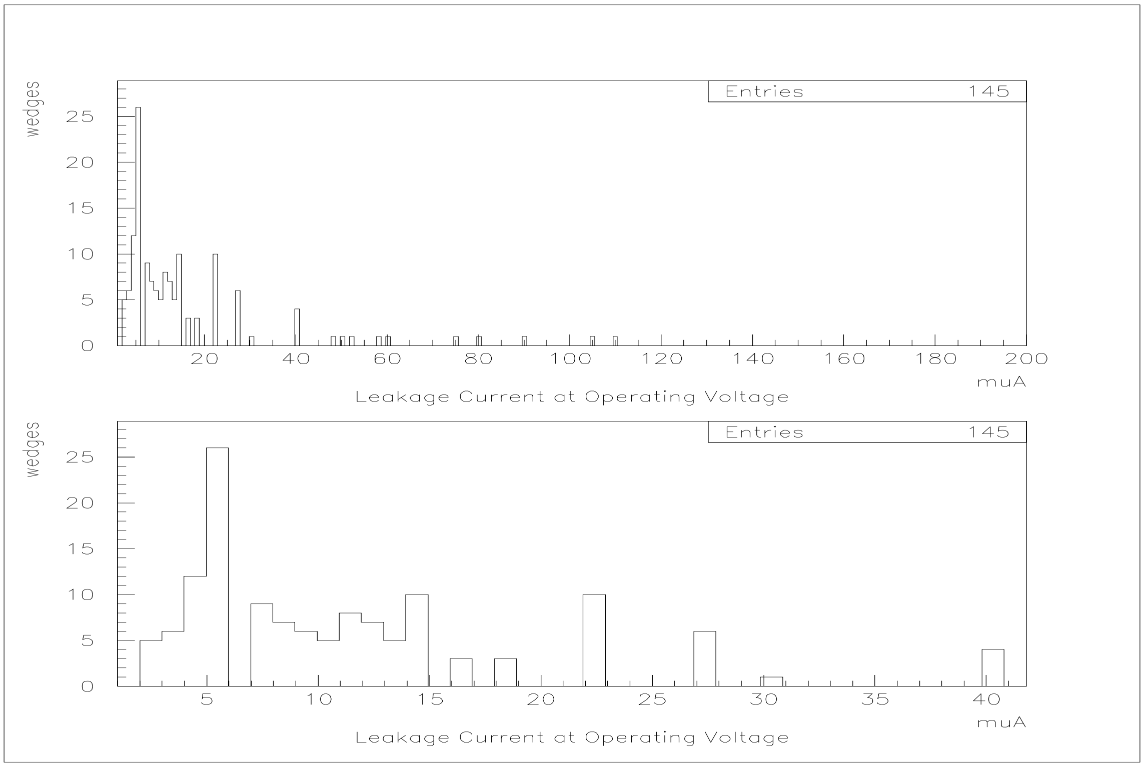

Figure 0.16: Leakage current found for each wedge in μ A. The bottom plot shows the lower leakage current wedges on an expanded scale.

Bibliography

[1] Compunetics Inc., Monroeville, PA.

[2] Max Levy Corporation in Philadelphia, PA.

[3] Promex Industries in Santa Clara, CA.