/CDF/DOC/SEC_VTX/PUBLIC/5049

Bench Tests of SVXII Quarter-Ladder Q005

Martin Rehn [1] and M. Garcia-Sciveres
Lawrence Berkeley National Laboratory

June 30, 1999

1 Introduction

This note is an excerpt from the master thesis of Martin Rehn [1] with minor editing, and summarizes work done at LBNL during the first quarter of 1999. It is being posted as a CDF note in order to make potentially useful information more readily available to members of CDF. This is not an exhaustive or definitive study of the operation of SVX II ladders, nor is it the first or last study of its kind. A complete understanding of the details of SVX II ladder operation will ultimately arise from combining the results from all such studies.

The device tested for this note is quarter-ladder Q005. It is a functional layer 0 prototype. It has one grade B production sensor instrumented using prototype hybrids and radiation-soft prototypes of the readout chip (SVX3C). A full readout unit, known as a half-ladder, would have two sensors end-to-end. The sensor was classified as grade B due to a large pinhole count, but was otherwise of production quality *b*efore the assembly of Q005. However, on Q005 the sensor exhibits higher than acceptable leakage current, which is presently believed to be caused by a reaction between the glue used to mount the hybrids and scratches in the sensor passivation (but this was not known at the time of this study).

2 Biasing

2.1 One Sided vs. Two Sided Sensors

In a one-sided silicon sensor it is desirable to have as large a depletion region as possible, but one-sided sensors can be operated even if not fully depleted. The depletion region is the active part of the sensor; the larger the depletion region, the larger the signal generated by charged particles passing through. The efficiency increases linearly with the width of the depletion region. In the case of a double sided sensor there is in addition another mechanism to the effect that the ohmic side of the sensor only works if the sensor is *fully* depleted. That is, the depletion region extends all the way through the silicon. The reason for this is simple: Undepleted doped silicon is a conductor. If the bulk silicon surrounding the strips on the ohmic side is not depleted of charge carriers, those channels will be shorted together.

2.2 Split Bias

The simplest way to connect bias to the sensor would be to apply bias voltage from one side and ground the other side. Grounding here means connecting to the same ground that the readout chips use. This would put all of the bias voltage across the coupling capacitors of the first side. The highest available bias voltage will then be set by the breakdown voltage of those capacitors alone. That is an unnecessarily low limit. Instead, the SVXII detector has been designed to allow *split bias*, where the bias voltage is applied partly from one side of the sensor, partly from the other, with the ground level somewhere in between. At best, this doubles the maximum bias voltage available. Bias voltage is connected to each implant strip individually through a 2 MΩ resistor.

[1] visitor

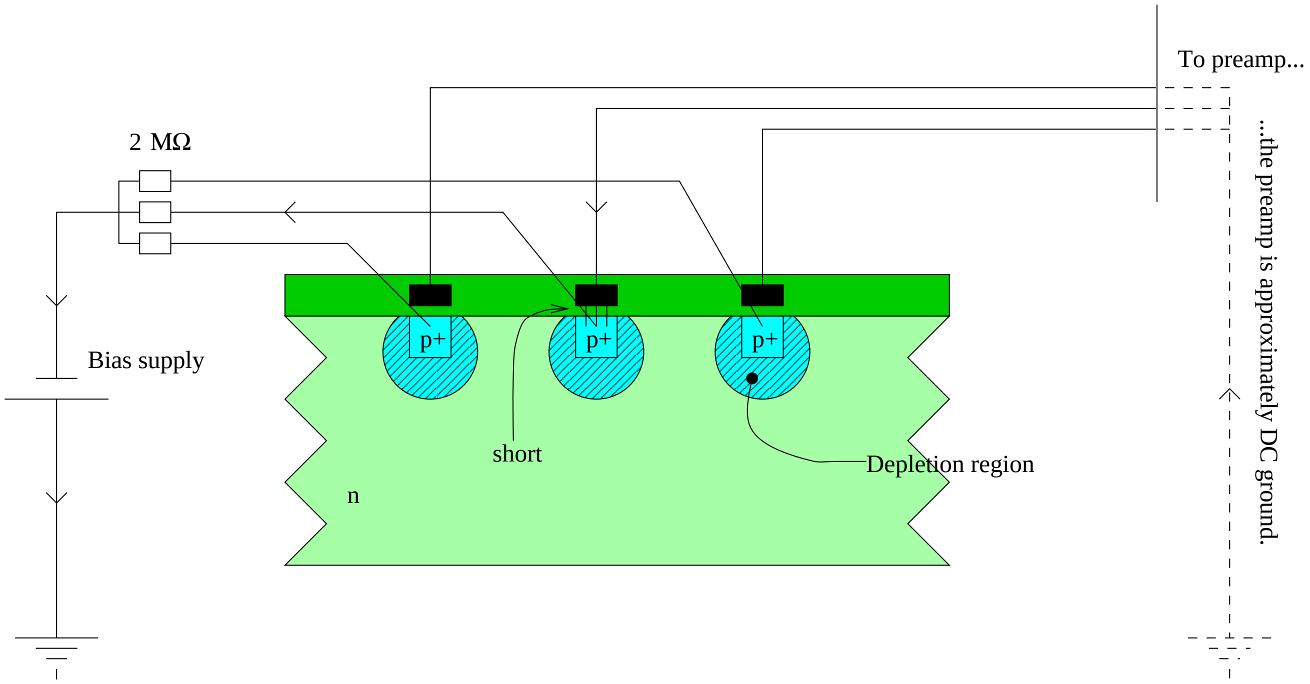

Figure 1: A pinhole on the junction side (currently p-side) of a sensor. The current flowing through the short is limited by the 2 MΩ resistor.

Another approach altogether would be to use *floating electronics*, meaning there is no common ground level. Each side's readout electronics would regard its own bias line as the ground level. That would nullify the limits imposed by voltage drop across the coupling capacitors. Floating electronics is problematic because of the need to separate ground levels and was not chosen in the design of the SVX II detector.

2.3 Pinholes

A pinhole is a short through a coupling capacitor. It can be a defect of the manufacturing process, caused by electrostatic breakdown, or by mechanical damage, for instance damage that occurs when a sensor is wire-bonded to its readout chips. A pinhole causes the strip implant to be directly connected to the readout electronics, in a DC rather than an AC sense. To protect the SVX3 chip from the high voltage, a protection device inside the chip will be activated. The protection is implemented as as diode that connects the input channel to ground or power, depending on polarity. For the bias voltage, a pinhole therefore looks like a connection to ground, as illustrated in figure 1. To remedy the situation, the wire-bonds connecting strips suffering from pinholes to the SVX3 chip are removed, *pulled*.

2.4 Effects of Radiation

Layer 0 of the SVX II detector will be subjected to a radiation dose on the order of $2\,MRad$ during Run II. The radiation will have several effects as far as biasing is concerned. Firstly, it will reduce the breakdown voltage of the coupling capacitors (100 volts is the limit for an irradiated sensor. For a non-irradiated sensor the rating is 120 volts). Secondly, it introduces defects into the silicon, yielding a larger leakage current. Finally and most remarkably, it introduces p-type doping into the silicon itself.

The latter phenomenon has little effect on the heavily doped strip implants, but it significantly alters the properties of the lightly n-doped bulk silicon. At first the effect is to reduce the strength of n-type doping. When the original doping has been completely eliminated, the silicon undergoes *type inversion*, turning the n-type silicon into p-type silicon. This does not mean that the sensor will cease to work. There will still be a junction side and

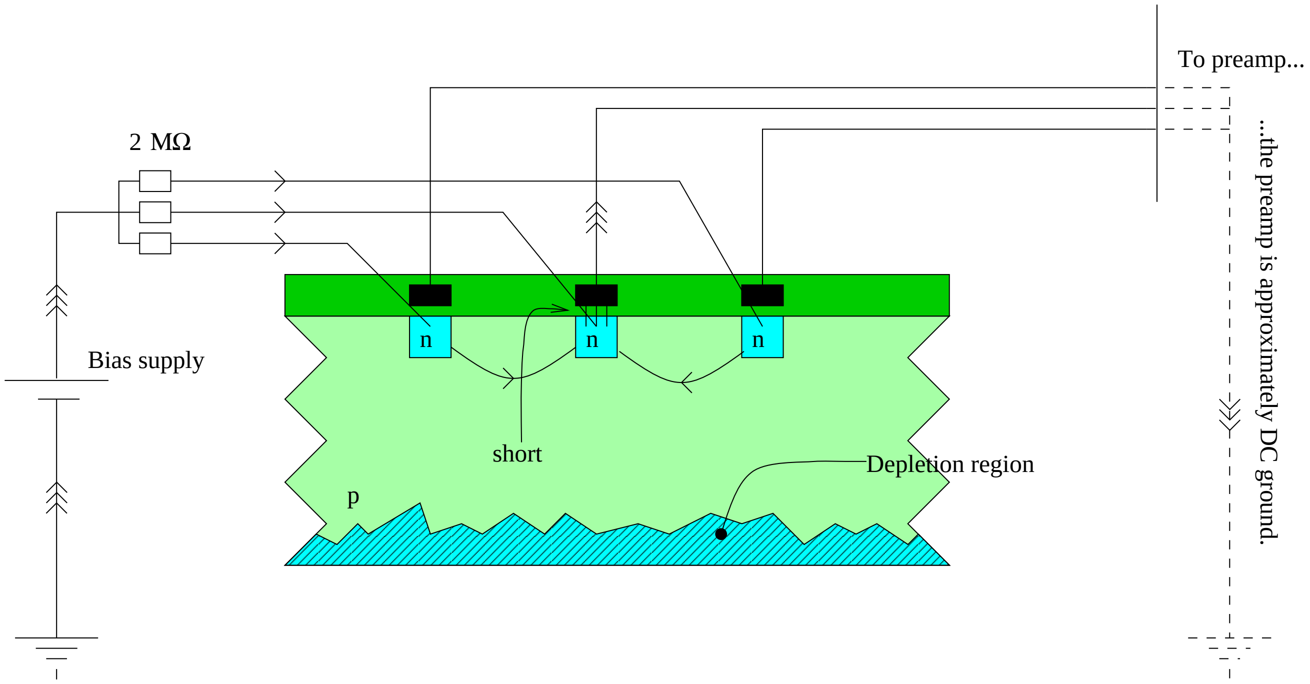

Figure 2: On the ohmic side of an undepleted sensor, strips are shorted together, illustrated by arrows through the silicon bulk. Thus a pinhole causes a much larger current to flow, as the 2 $M\Omega$ resistors are effectively connected in parallel.

an ohmic side, only they have reversed roles, the p-side now being the ohmic side, the n-side the junction side. Type inversion in SVX II will occur after about 500 kRad exposure.

2.5 Bias current Measurements

2.5.1 With Pinholes

When pinholes are present they will result in very high currents, as is illustrated in figure 3. Since the pinholes represent connections to ground, incoming current (on the n side) will not equal outgoing current (on the p side). The measurements of currents "with pinholes" presented here refer to a situation with one pinhole on the p-side, one on the n-side. But far from the apparent symmetry, the n-side pinhole completely dominates, as explained in figure 2.

One result of this extra ground reference is that it matters for the currents how the bias voltage is split. The p-side current will be largely controlled by the p-side voltage, and the other way around. This effect (for the p-side current) is seen in figure 3. The overall current is smaller when voltage is applied from the p-side, where the pn-junction is. The reason for that is that the depletion zone starts forming from that side. As the bias region extends far towards the n-side it will start cutting off the current between n-side strips shown in figure 2. This will reduce the number of 2 $M\Omega$ resistors connected in parallel.

Without taking advantage of this cutting-off effect it proved impossible to deplete the test sensor. If the n-side bias voltage is turned on without any voltage applied from the p-side, it will just max out the available current from the bias supply,[2] without contributing much to the bias voltage across the silicon. Even if the p-side bias is then also turned up until it reaches its current limit, the result will just be lots of current and little depletion. If on the other hand the p-side bias is turned on *first*, then the current drained into the n-side pinhole is limited and the sensor can be successfully depleted with the help of both bias supplies. Any particular combination of p-side- and n-side voltage will always give the same current. In this sense the detector does not possess a memory. But the order in which you do things does matter for which voltage combinations are achievable.

[2] Both bias supplies are limited in this case, allowing a current of about 600 μA

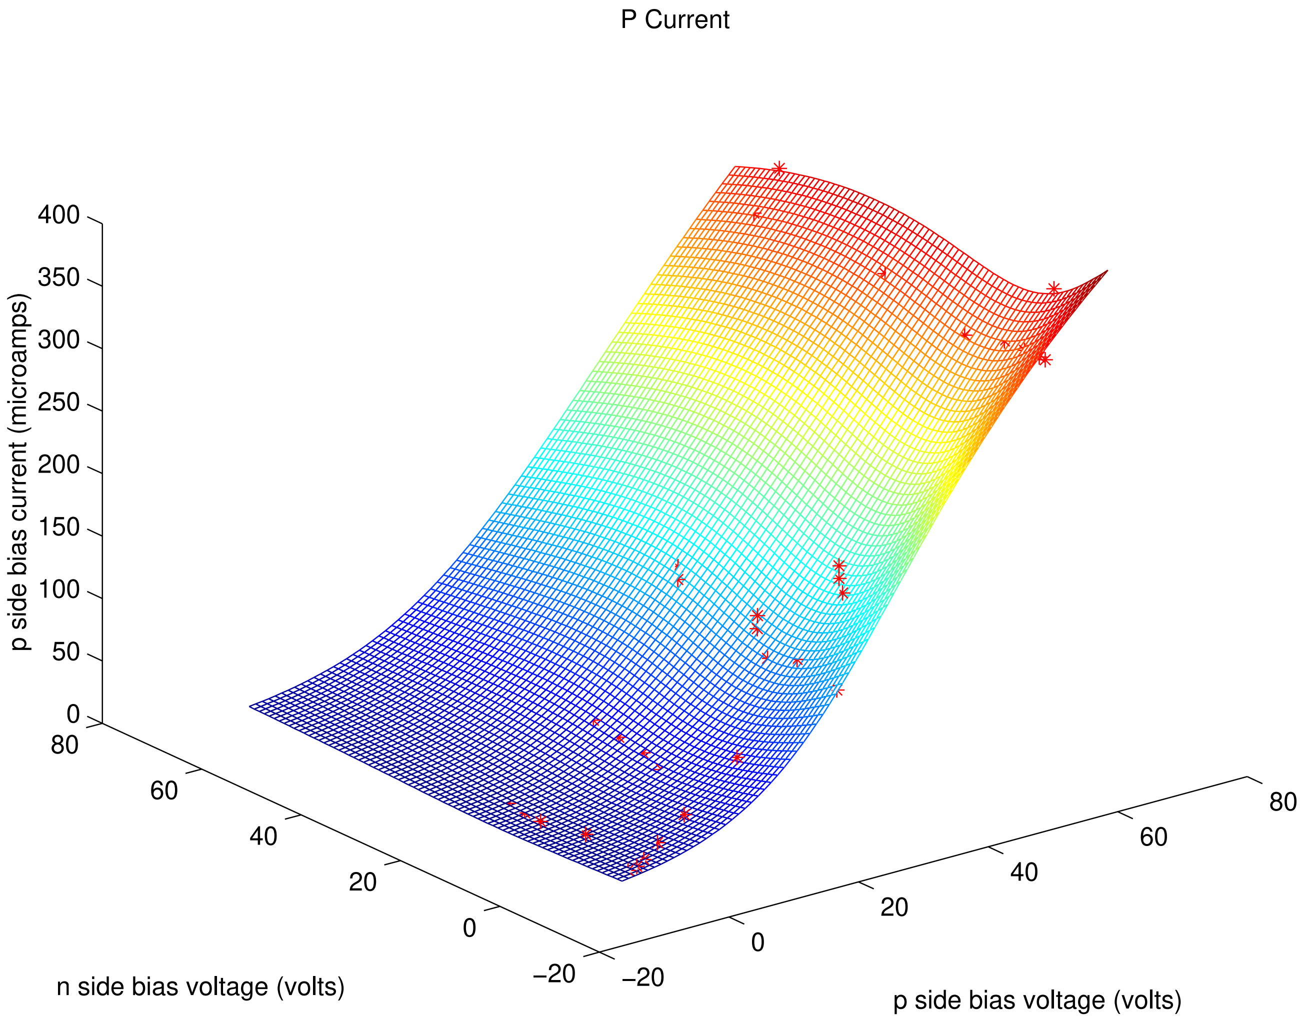

Figure 3: The current entering the p-side bias connection, when there was one pinhole on the p-side and one on the n-side. Most of this current is drained through the pinholes, meaning the n-side voltage has little effect on the p-side current. The stars (some are partly hidden below the mesh) are measured points. The mesh is an interpolation using a two-layer perceptron method.

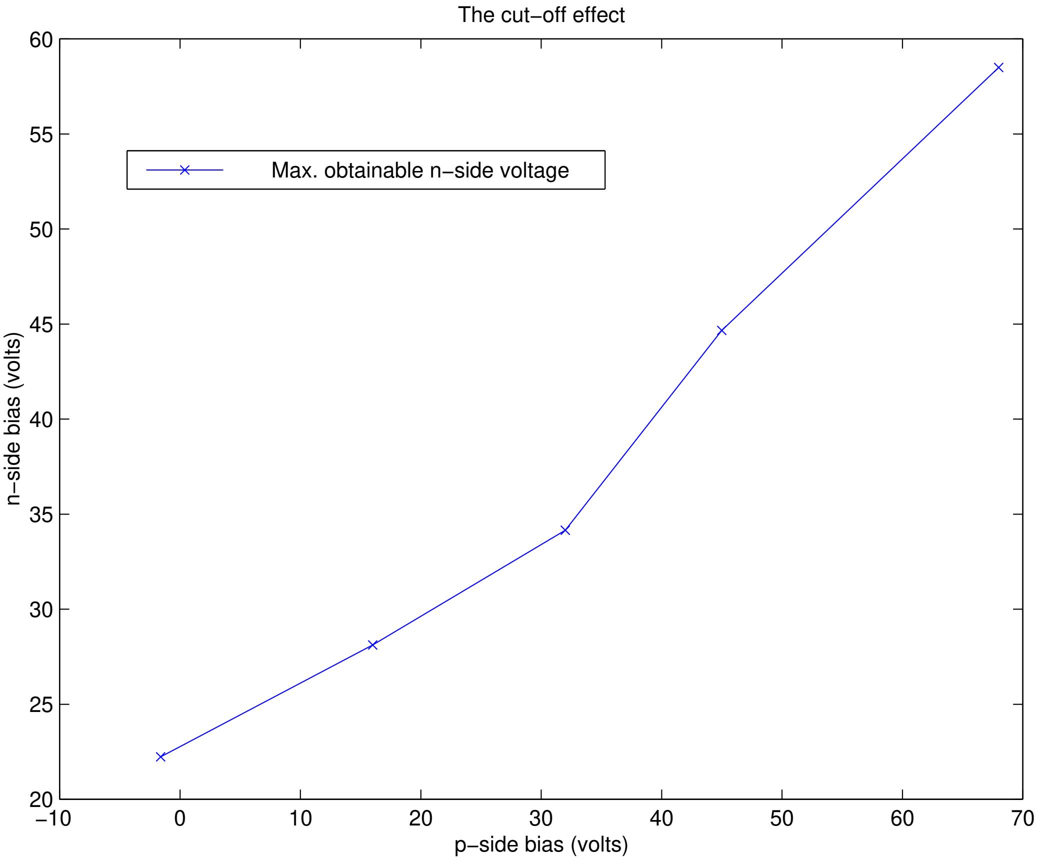

Figure 4: Cutting off the current through the n-side pinhole. The more the sensor is depleted from the p-side, the higher n-side bias can be applied before reaching the current limit of $600\,\mu A$. There is one pinhole on each side of the sensor.

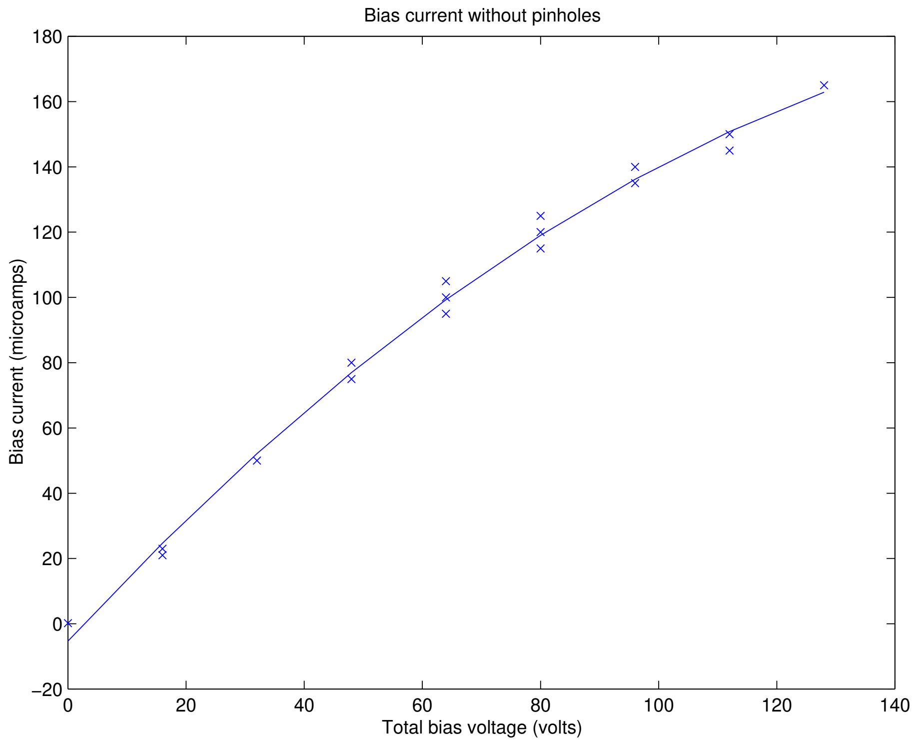

Figure 5: Bias current as a function of total bias voltage (p-voltage + n-voltage). The in-going current was found equal to the outgoing up to instrument precision ($10\,\mu A$). The crosses are measured points. When several crosses occur at one total bias reading, they originate from different combinations of p-voltage and n-voltage.

Figure 4 illustrates this effect. A related effect is that, for some voltage ranges, a larger p-side voltage will result in a *lower* n-side current, other things being equal. In this case the sensor acts like a field effect transistor, with the p-side as the gate, the n-pinhole as the drain, and the remaining n-side strips as the source.

2.5.2 Without Pinholes

When the n-side pinhole was removed from the test readout module, the bias current dropped dramatically. The subsequent removal of the p-side pinhole reduced the current by just another 10%, arriving at figure 5. The dependence of current on from which side the voltage was applied disappeared, as did the difference between incoming and outgoing current. Thus figure 5 shows the complete voltage-current behavior in *one* curve, rather than two 3D-plots.

Despite the improvements that removal of pinholes brought, the currents are still several times higher than those of a perfect sensor. Because of scarcity of good components, it is necessary to use more or less flawed materials for tests. The actual mechanism leading to high currents was not identified.

2.5.3 Conclusion

When pinholes are present on the ohmic side, special techniques are needed to deplete the sensor. This involves applying the full depletion voltage from the junction side. At this time the ohmic side is effectively grounded by the pinholes. This means that *all of the depletion voltage will appear across the junction side readout capacitors.* To put it another way, one single pinhole on the ohmic side threatens to defeat the split bias scheme.

This is not likely to be a problem initially, since all known pinholes will be pulled before the detector is installed. But new pinholes may develop during the operation of the detector. The effects of radiation damage make biasing progressively more difficult. If it ultimately turns out to be impossible to fully deplete the sensor,

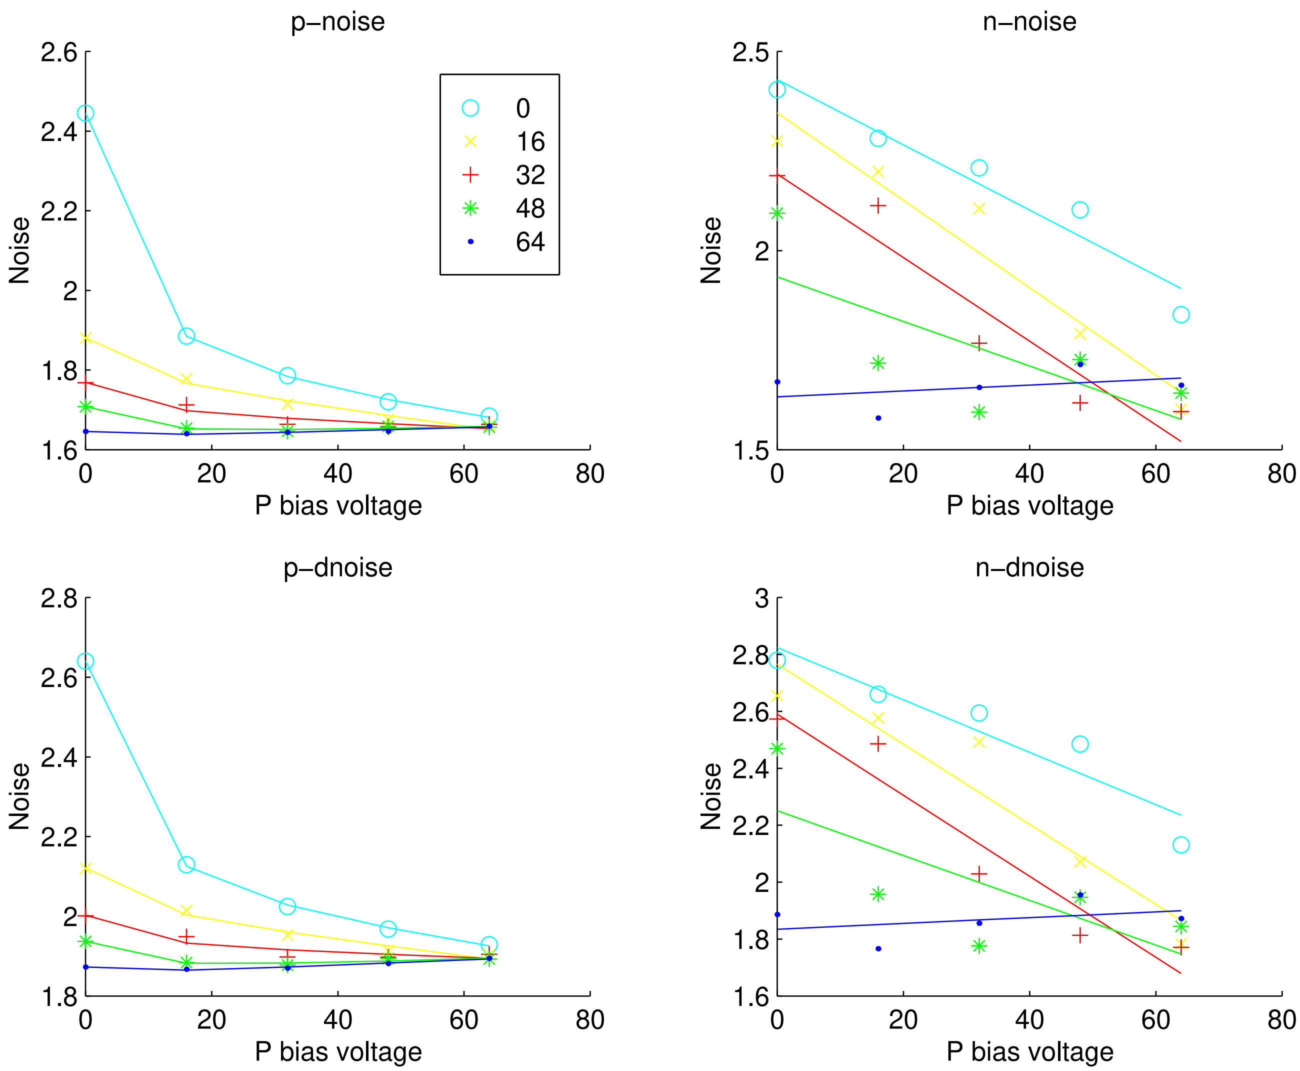

Figure 6: Noise vs. bias measurements. Each diagram shows five sets of data, each corresponding to a different ϕ-side bias voltage, as indicated in the legend box. The dots, crosses etc. are data points, the curves are fitted to the points. In the ϕ-side diagrams, to the left, a $\frac{1}{x}$ term was allowed in the fit. For explanation see section **??**.

the ohmic side will be useless. With this possibility in mind, the junction side of the sensors are designed to serve the most important function, namely to track the ϕ dimension of particle paths. But late in the life span of the detector, type inversion will have occurred and the ϕ side will be at risk.

2.6 Bias and Noise

Throughout this paper the word "noise" will refer to the RMS of the digitized output of a single channel for a sample of events of order 100 taken in rapid succession. The units used are ADC counts. One ADC count corresponds to approximately 750 electrons of input charge for the ADC ramp resistor (100 kΩ) and digitization clock frequency used (40 MHz). The term differential noise (dnoise) refers to the the RMS of the event-by-event difference between the digitized output of two adjacent channels, divided by $\sqrt{2}$.

The p-side noise is affected in a dramatic way when a small bias voltage is applied. In figure 6 it seems to drop like a $\frac{1}{V}$-curve for low bias voltage. This is in fact what happens. The depletion zones surrounding the strips act like capacitors. At low bias voltages, the depletion zones will look like semi-cylinders surrounding the strips on

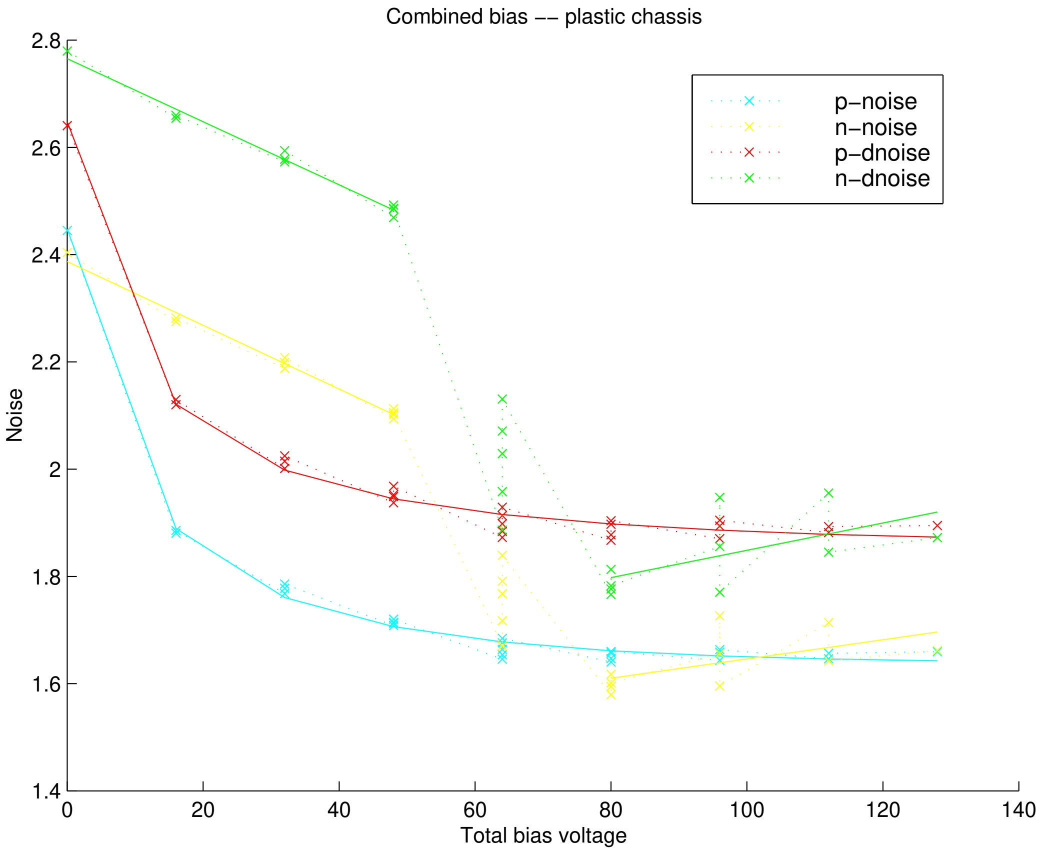

Figure 7: Noise as a function of total bias voltage. Measured data points, fitted curves.

the p-side. The capacitance of such a depletion zone is inversely proportional to the radius of the semi-cylinder, which in turn is proportional to the applied voltage, explaining the $\frac{1}{V}$ relationship.

The level of noise is seen to depend only on total bias and not on splitting configuration. Figure 7 shows this. Different data points at the same total bias level correspond to different bias splitting. For n-side noise there is a large spread of the measured points at the transition point (62 volts) where the sensor becomes fully biased. At this point the depletion region reaches the n-side strips. Different straight lines have been fitted above and below this point. For the p-side, curves including a $\frac{1}{V}$ term have been fitted, as above. Nothing dramatic happens here at the transition point. For both sides, noise does not change as the bias is increased beyond full depletion.

2.7 Bias and charge collection

The above tests all refer to a situation where there is no signal in the sensor. To verify that the depletion really works properly, charge has to be injected into the sensor. One way to do this is to shine laser light onto the sensor surface. A laser with a wavelength of 890 to 910 nm was used. Light of this frequency cannot penetrate the silicon, but generates charge carriers at the surface of the sensor. To test that charge is properly transported one therefore has to apply the laser from one side, and look at the output from the other. The p-side has been more widely tested, and is known to work well with the laser [2]. This note focuses on results with the laser applied to the p-side and data read out read out from the n-side .

When the sensor is not depleted, the charge will spread out all over the n-side strips. The undepleted region in the silicon is a conductor, meaning there is no electric field to drive the charges. Meanwhile, as has already been noted, the n-side strips are connected together when there is no depletion on that side. Both effects contribute to spread out the signal across several channels. The separation between n-side strips is 141 μm. For a fully depleted sensor, this distance is larger than the diffusion distance, meaning all of the signal will fall on one single strip, if the laser pulse is positioned directly above it.

The results are presented as an array of plots in figure 8. What changes between the plots is bias voltage. Each column represents a certain p-side voltage, each row a certain n-side bias. Both bias voltages are increased from 0 volts (upper left) in increments of 16 volts. In the lower right corner the bias is $64+64$ volts. On diagonals (sloping like this: /) the total bias voltage is constant. With the exception of the main diagonal (where total bias is 64 volts), the signal distribution is independent of how the bias is split. This is to be expected, since the sensor is not aware of the ground level, and therefore sees only the total bias voltage. The exception may be explained either by some small remaining ground connection (pinhole) or by a voltage drop between the voltmeters and the sensor, meaning the total voltage experienced by the sensor is not exactly the same as was measured.

There is a transition at total bias 64 volts from a signal spread out across more than one strip, to one sharply peaking at one single channel. When the transition is complete, at a bias of 80 volts, the sensor can safely be considered fully depleted. The signal does not change as the bias level is further increased. The sensor is found to work well for bias voltages of 80 volts, and above.

This technique was also used to check how much bias voltage was needed to make the sensor work when pinholes were present. Figure 9 shows the results.

3 Grounding and Shielding

3.1 Pickup Sources

In the environment where the SVX II readout modules are to operate there will be several sources of electromagnetic *pickup*. It is usually the sensor itself that picks up unwanted signals. The reason for this is that noise that enters the signal path before any amplification has taken place, will benefit from the greatest amplification, along with the signal. Also, the sensor is a large surface, not unlike an antenna when it comes to picking up noise. Potential noise sources could be the voltage lines powering the SVX3 chip, the data bus activity, and control signals in the barrels where the modules are mounted.

3.2 Physical Layout

The test readout module, mounted in its metal assembly fixture (chassis), turned out to suffer from a large noise problem. Interactive inspection of the readout showed that the noise originated from random pedestal shifts common to all channels, *common mode noise*. The effect was larger on the n-side of the sensor.

Since one of the things that all channels share is grounding, it was suspected that this might be involved. Could there be a bad ground connection somewhere? A series of measurements aiming to improve the grounding were carried out. The results, shown in figure 11, were surprising. Most attempts to add ground connections *increased* the noise. The added noise followed the pattern of the original problem. It was common mode, and most prominent on the n-side. The worst configuration (number eight and nine in the figure) was one where the readout module was put inside a metal box, that was connected to the chassis. The reasoning behind doing this was to build a Faraday's cage around the readout module, shielding it from any electromagnetic noise. The flaw in the arrangement is that it is not possible to shield the readout module in this way. The electrical wires connecting to the outside world provide a reference potential. Relative to this potential the chassis and box will not shield, but act as an antenna. The more material is added as "shielding", the more efficient an antenna will be created.

3.3 Capacitive coupling

To fully understand how the "antenna" works, it has to be figured out exactly how it transfers noise into the signal. A clue was that the ratio of n-side to p-side common mode noise was similar to the ratio of silicon strip lengths in direct overlap with the metal fixture. This suggested a capacitive coupling between the chassis and the sensor. Figure 12 shows the coupling geometry. The capacitive coupling occurs between the metal strips in the sensor and the metal fixture overlapping them. On the n-side, one channel corresponds to several strips, as illustrated in the bottom panel of figure 12. However, the metal fixture overlaps exactly one strip per channel for the channels in question, so the length is the effective width of the sensor.

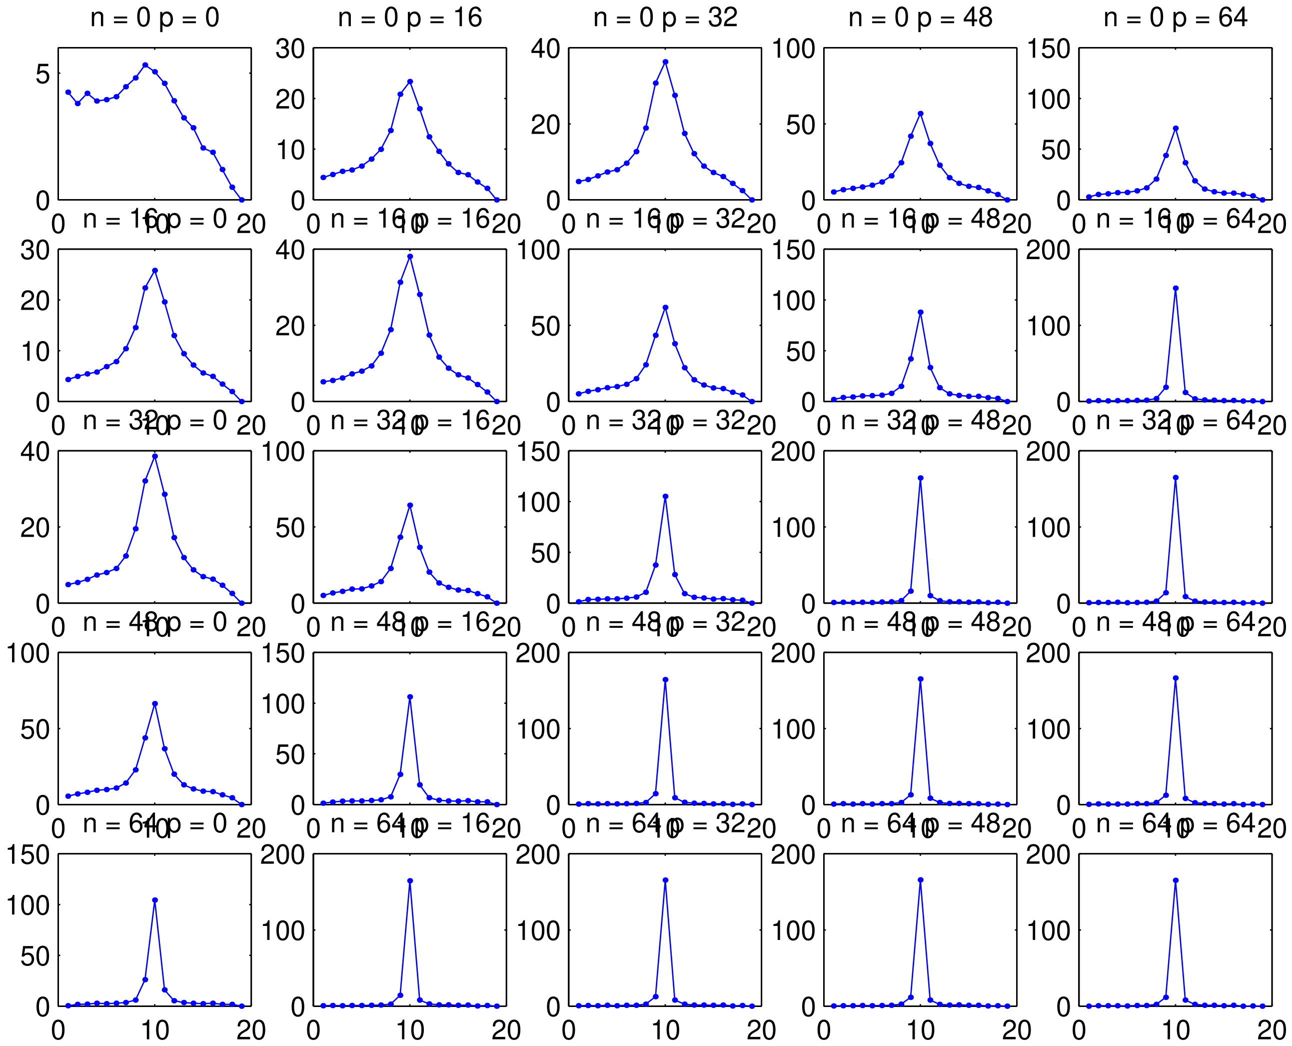

Figure 8: A laser was applied to the p-side of the sensor. Shown here are the signal distributions in the n-side strips near where the laser was pointed. The x-axis in each plot thus represents relative channel number, with the pulse centered on channel 10. The y-axis represents average signal in ADC counts. Note that y-axis scale varies, but is the same on a diagonal (/). The results are discussed in section 2.7.

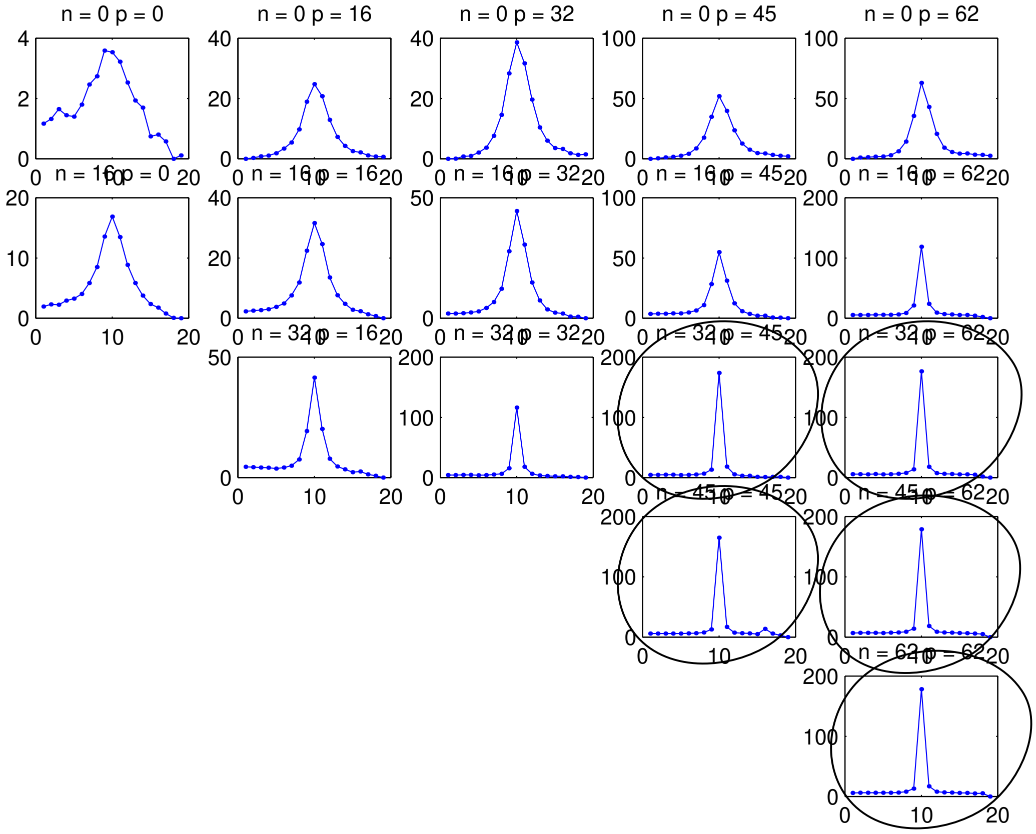

Figure 9: Test of sensor depletion with pinholes. The sensor is depleted when the laser pulse results in a narrow peak, as it does for the circled cases. Each column of plots represents a certain p-side voltage, each row a certain n-side bias. See figure 8 and section 2.7 for a further discussion. The cutoff effect was taken advantage of to achieve the bias voltage combinations shown here. Still, the combinations corresponding to the missing plots in the lower left corner could not be reached.

The strip pitch is 140 μm on the n-side and 60 μm on the p-side. The separation between sensor and fixture is roughly the same on both sides. Therefore, the ratio of capacitances will just be determined by the ratio of the length of the overlap times the strip pitch. For all the p-side channels the length of the overlap is the width of the metal fixture in contact with the sensor on that side, $w = 10.5\,mm$. For the n-side it is the effective width of the sensor ($h = 15.3\,mm$) for the set of strips over the solid part of the fixture, and zero for all other channels. From this we would expect the noise ratio to be approximately 3:1, with the affected n-side channels having the larger noise. Note that the hybrid used on the n-side of the present readout module has only one functioning chip, and all the channels on that chip happen to be connected to strips overlapping the fixture- in other words, it was not possible to look at the unaffected channels on the n-side.

To compare to the observed noise ratio, noise from other sources has to be subtracted from the curves plotted in figure 11. From figure 6 we have the noise level without the metal fixture, conditions otherwise being the same. P-side noise is here read off as 1.65, n-side noise as 1.75. Assuming that this noise is independent from the noise associated with the metal fixture, the curves plotted in figure 11 should be the sum in quadrature of the two noise types, $n^2 = n^2_{fixture} + n^2_{other}$. Extracting $n_{fixture}$ produces figure 13. From that figure, the noise ratio is found to be 2.6:1, which seems consistent with the above estimate.

3.4 Reducing the noise

There are two ways to reduce noise from capacitive coupling to a conductor. One would be to connect the box and chassis to a ground point very close to the ground of the readout module. This corresponds to a connection (**d-e**) in figure 10. As seen in configuration 9 of figure 11, there is no point in trying to connect to a ground point far away. The grounding must be achieved in an AC sense. More specifically, the impedance between chassis and sensor ground must be much lower than that between chassis and the noise signal coupling to the sensor.

The other way to get rid of the noise is to remove the antenna itself. For the test readout module, the metal chassis was replaced with a plastic one. That completely removed the large pickup effects. See figure 6 for the resulting noise level. In section 4 some more measurements of grounding schemes are presented.

3.5 Conclusions

Capacitive coupling between the sensor and an electrical conductor may introduce common mode noise. If such coupling is unavoidable, added noise can be minimized by a low impedance connection between the hybrid and the conductor in question. In the final SVX II design there is a capacitive coupling between the sensors and the metal bulkhead that the ladders are mounted on, as well as between sensors themselves on adjacent layers. This study suggests that to minimize common mode noise the SVX II hybrids should be grounded to the bulkheads.

4 Dead-timeless Operation

4.1 Additional Noise Sources

The measurements of readout module performance so far relate to a somewhat idealized mode of operation, where the SVX3 readout chip is idle at the time data is collected. To take full advantage of the high event rate of the upgraded Tevatron, the CDF system must be able to collect data at all times. It cannot afford the luxury to wait for the SVX3 chips to finish reading out the previous event before recording a new one. The activity in the chip may disturb new measurements being taken. This is known as *pickup*; chip activity, external control signals, or activity on the readout bus feeds back into the signal path.

The readout module used in this study is more sensitive to pickup than the production SVX II ladders will be for two reasons. First, the flex cable connecting the module to the upstream electronics is a prototype version that lacks an integral shield layer. Second, said flex cable comes out directly over and parallel to the p-side sensor strips, rather than folding back over the ϕ hybrid as will be the case for ladders installed in barrels. Therefore, the results of this section show a sort of conservative limit on the pickup effects to be expected with the final components (the effect of the routing of the flex cable is explicitly investigated to the extent permitted by the test setup).

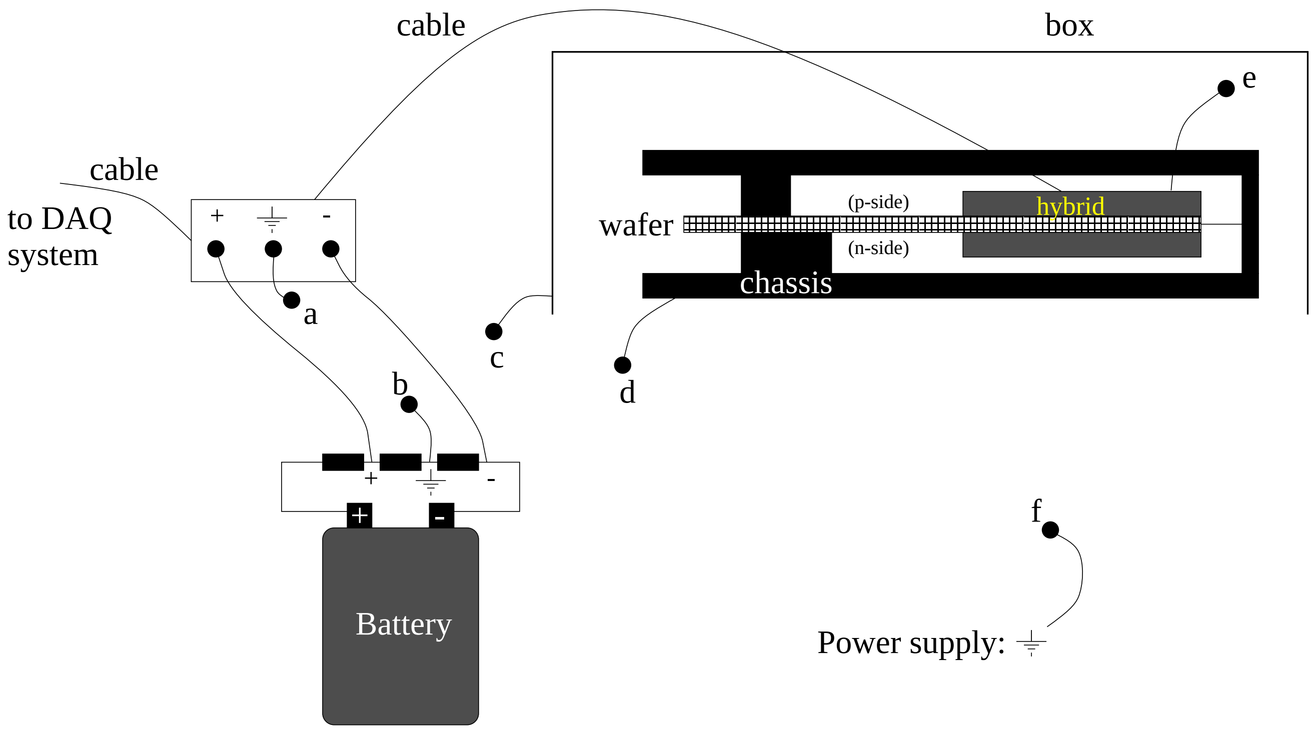

Figure 10: The test readout module mounted in a metal chassis. A voltage divider connected to a battery provides a positive and a negative bias voltage referring only to the midpoint of the voltage divider. The points **a**–**g** are all in some sense "ground". They can be connected in different combinations. The **box** can be put on top of the module like in the figure, or it can be removed.

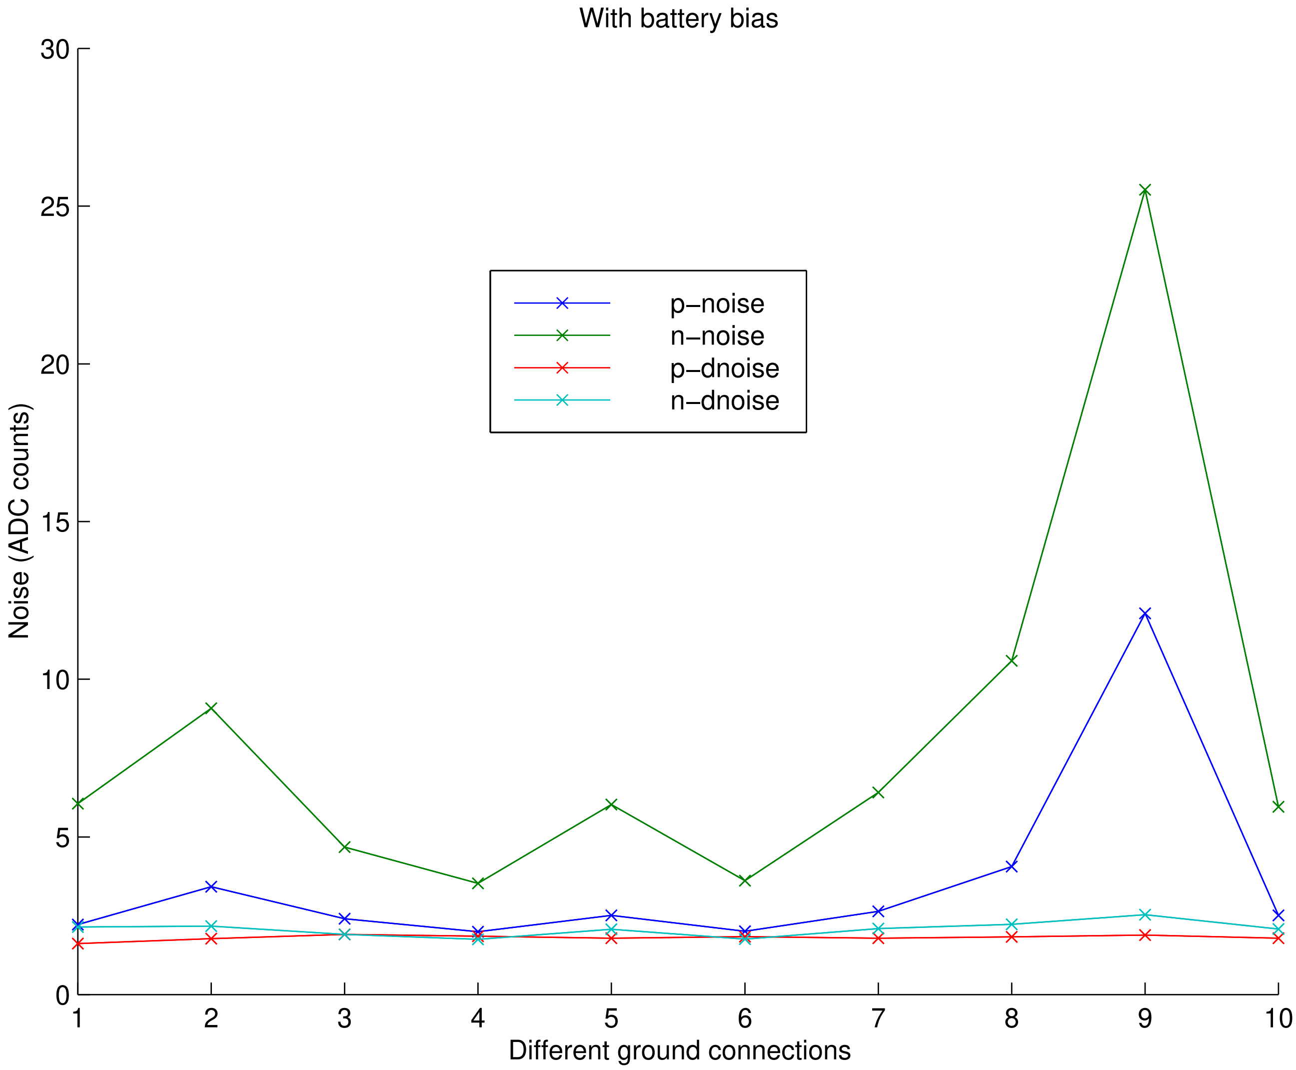

Figure 11: The effects of grounding and shielding on noise and differential noise. Refer to appendix ?? for definition of the noise types. Each number on the x-axis represents a different grounding configuration. Using figure 10 as a key, the below table lists how the grounding was done. Each dash in the "Connections" column represents a wire. A check ($\surd$) in the "Metal Box?" column means that the box was used:

#	Connections	Metal Box?
1	**a-b**	–
2	**b-d**	–
3	**a-d-b**	–
4	**b-a-b**	–
5	(none)	–
6	**a-d**	–
7	(none)	$\surd$
8	**c-d**	$\surd$
9	**a-d-c**	$\surd$
10	**b-c**	$\surd$

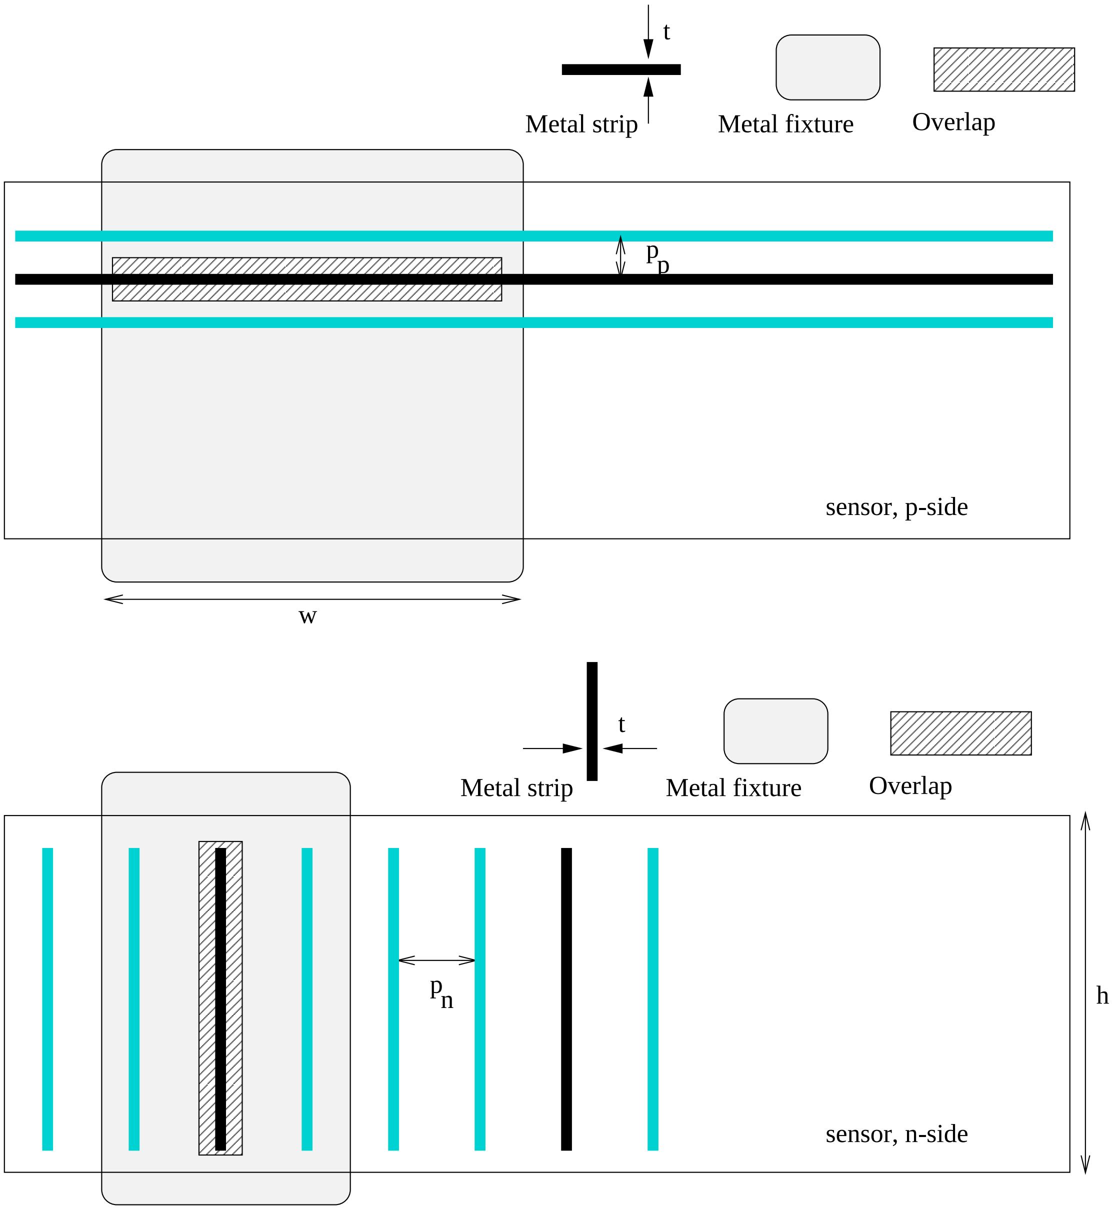

Figure 12: The sensor and the metal chassis. Does the coupling geometry explain the relative noise magnitudes? This figure defines the relevant overlap areas for the two sides of the sensor. See section 3.3 for the calculations.

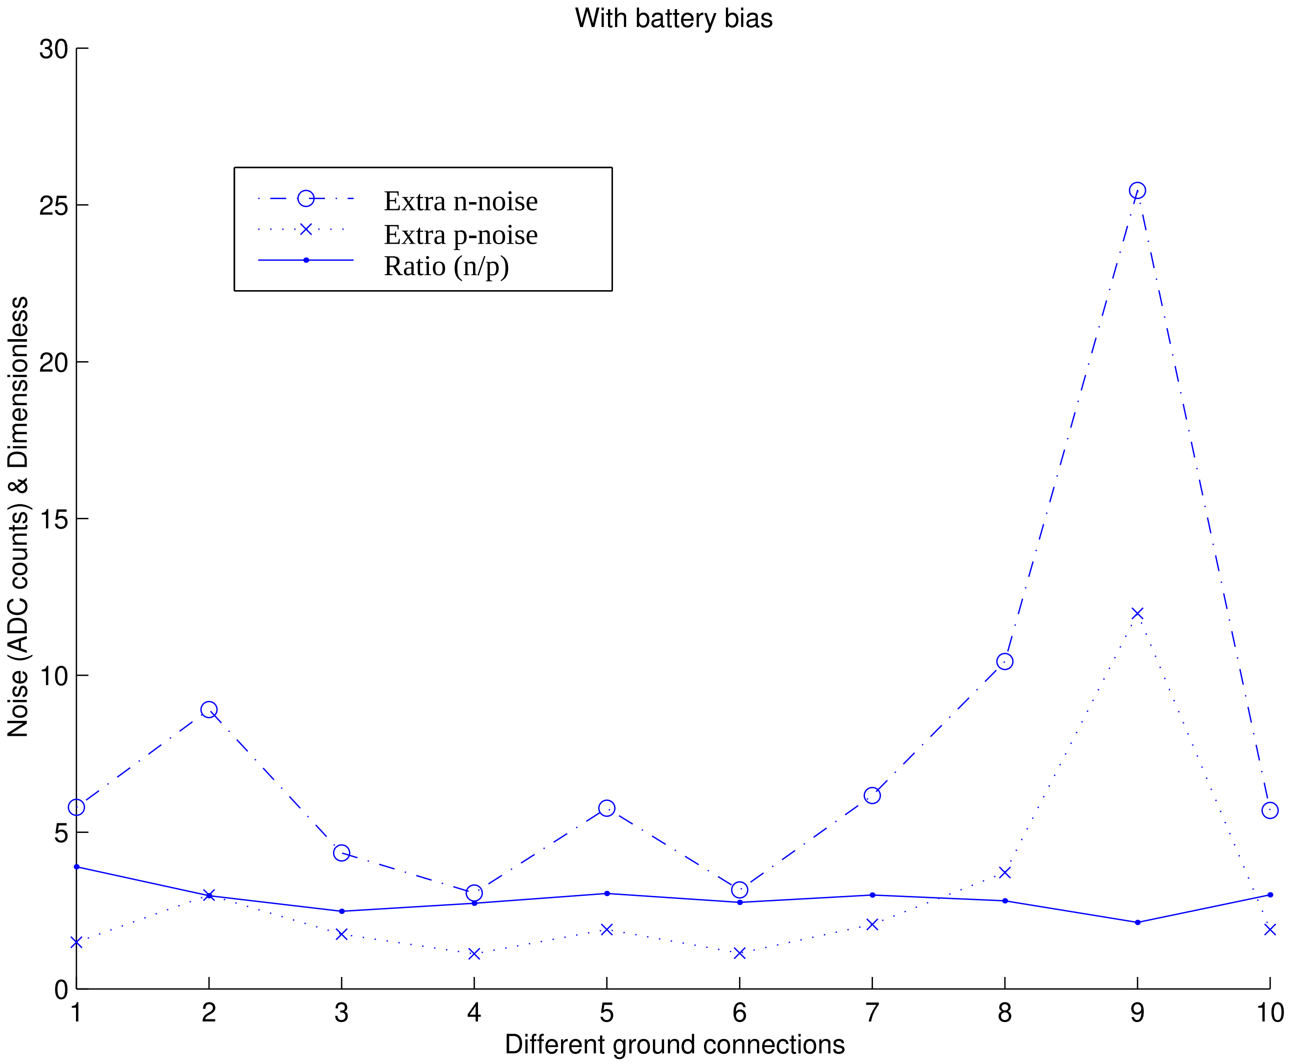

Figure 13: "Extra" noise originating from the metal chassis. Supposing an independent background noise of 1.75 for the n-side, 1.65 for the p-side, the extra noise is plotted for the n– and p-side. The ratio between the n-side and p-side additional noise is almost constant, with a mean of 2.6. This was the observation leading to the identification of the coupling mechanism for the noise.

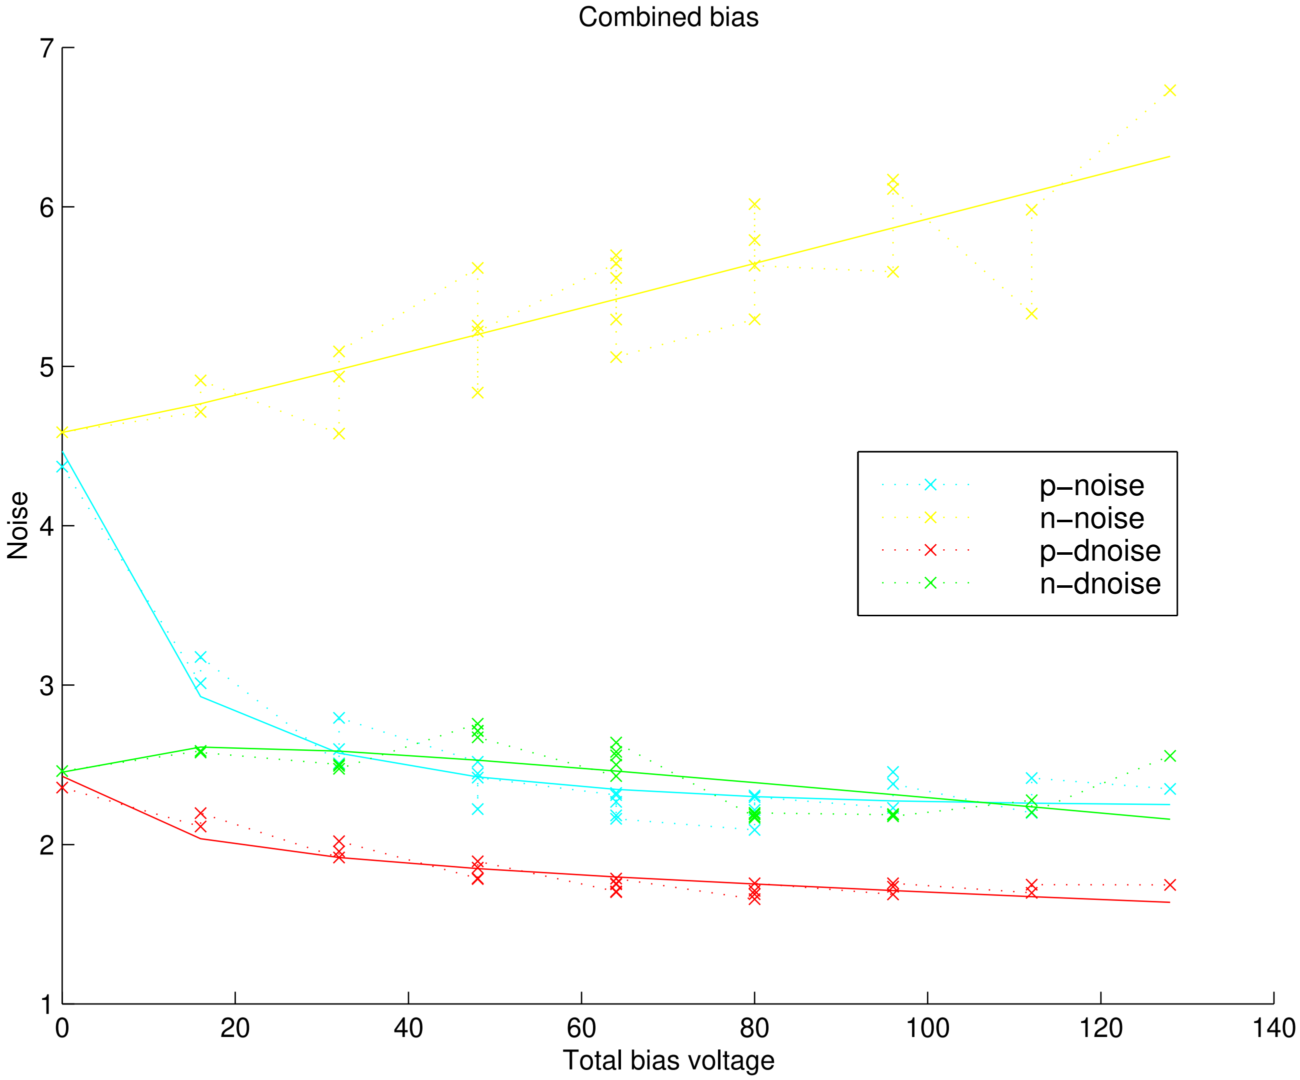

Figure 14: Noise as a function of total bias voltage, using the metal chassis.

4.2 Principle of Measurements

In the SVX3 chip measurements are continuously being stored in the analog pipeline [3]. With each tick of the front end clock, the pipeline stores a measurement in the current cell and advances to the next storage cell. The chip can be instructed to digitize and read out a previously stored measurement at any time. The cell containing the measurement in question is then taken out of the write sequence to prevent it from being overwritten, but new measurements are still being stored in the remaining cells.

To determine the effect of chip activity on the quality of the recorded signal, we look at what happens to measurements that were recorded and stored in the pipeline while the readout module was performing different operations. This is accomplished by a special program, known as *map*. The map program performs a dummy measurement that makes the SVX3 chip go through its sequence of states – idle, digitize, readout, idle. At the same time a second, real, measurement is taken. By placing the trigger signal for this second measurement at different points in time, it can be investigated how data quality is affected by chip activity. Using the map program, the dead-timeless performance of the test readout module was tested and its dependence on bias voltage, grounding configurations and shielding was determined. This latter aspect of the measurements was similar to what was described in sections 2 and 3.

4.3 Data analysis

Measurements performed at a specific point in the operational sequence of the SVX3 chip are said to belong to the same *bucket*. The immediate output of the map program is a number of plots showing the signal distribution in individual buckets, like those in figure 15. Each bucket contains pickup from the chip activity that was going on at that particular time. The sensor does not carry any signal in these measurements, only the pedestal level. Therefore the output level will be synonymously referred to as “pedestal level”.

To the left in figure 15 is shown the pedestal levels of chip 1, one of the two ϕ-side chips, in a bucket corresponding to a time when the readout chip was idle. The solid curve is the average pedestal level, the dashed curve the RMS noise in the pedestal level. The bias voltage is $48 + 48$ volts, provided by the bias supply. The channels with an average pedestal level above 35 ADC counts in the plots are the pulled (unconnected) channels. In data analysis those channels along with their neighbors (which show a pedestal shifted slightly downwards) are disregarded. The remaining channels (“good channels”) behave nicely for this bucket in three ways: 1) The overall pedestal level (about 33 counts) is not shifted from its undisturbed value. 2) The pedestal level is flat across the chip. 3) The noise is low. Compare this to the right hand panel of the figure, where there is severe pickup.

4.3.1 Performance measures in a bucket

Each bucket represents a number of readouts of all channels in the readout module pertaining to a certain point in time. To determine the performance corresponding to such a bucket, five indicators were calculated; pedestal shift, maximum pedestal deviation, RMS pedestal deviation, maximum noise and average noise. Below is described how each of those are defined. The plots illustrating the measures were taken from chip 1, with a bias voltage of $48 + 48$ volts, just like in figure 15.

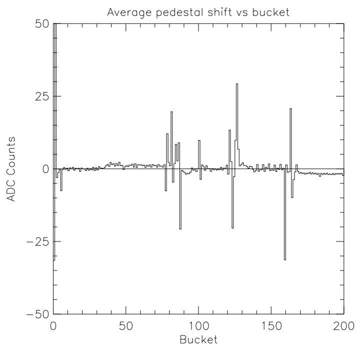

Pedestal shift

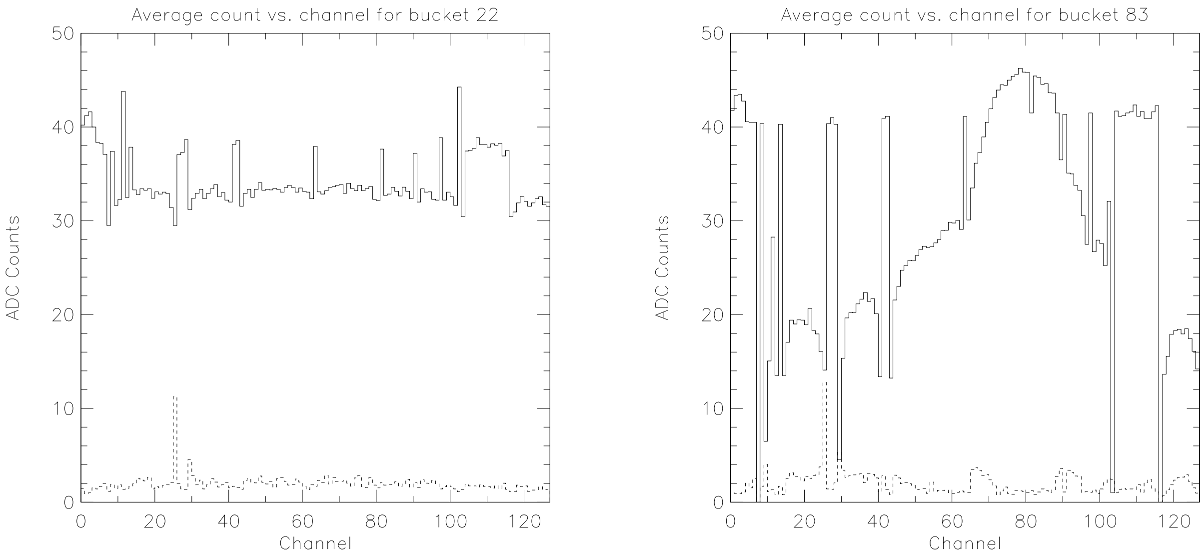

Bucket 22. A nice bucket. Bucket 83. A bucket with pickup.

Figure 15: Pedestal of chip 1 in quiet mode (left) and digitize mode (right). The solid curve is the average pedestal level, the dashed curve is the RMS noise fluctuations in that level. In the “bad” bucket to the right the pedestal is severely distorted. Both the average level shift and the deviations from that level are large.

The first measure on data quality is *pedestal shift*. To find the pedestal shift for a bucket, the pedestal level average taken over all good channels in that bucket is calculated. From that number is subtracted the average pedestal level for all good channels in *all* buckets. The difference is the pedestal shift for that particular bucket.

A pure pedestal shift, one which affects all channels evenly, may (within certain limits) be subtracted away using the dynamic pedestal subtraction mechanism described later. In that case some rather large level of pedestal shift can be tolerated. But if the DPS feature for some reason cannot be used in the SVX detector, the pedestal shift would be of great concern. An upward pedestal shift is more serious than a downward shift, at least if it affects many chips simultaneously. If this happens, there may be more signal-carrying channels than there is time to read out, and the event will be lost. In the case of a downward shift, chances are that most signal-carrying channels will still make the sparsification threshold and be read out; little information will be lost.

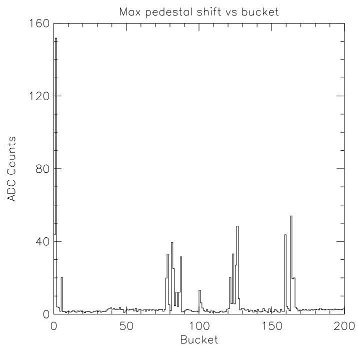

Maximum pedestal deviation

The next quality measure tells about deviation from the pure pedestal shift. *Maximum pedestal deviation* is defined as the greatest deviation from the pure pedestal shift for any one channel. To find maximum pedestal deviation for a bucket, the pedestal shift for each channel in the bucket, relative to the overall average level for that channel is calculated. This tells how much each individual channel has been shifted. Then the overall pedestal shift for the bucket, described above, is subtracted from the pedestal shift for the channel. What remains is a channel-specific shift that is in addition to the collective shift. The maximum absolute pedestal deviation is defined as the largest channel-specific shift for the bucket. Pedestal deviation cannot be subtracted away by the DPS system. For that reason it is regarded as a kind of noise.

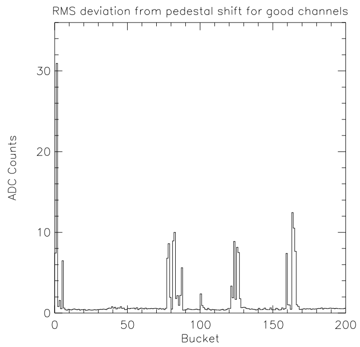

RMS pedestal deviation

The *RMS pedestal deviation* is calculated in a way similar to the maximum pedestal deviation. Instead of taking the largest absolute shift, one calculates the RMS over the individual shifts of the channels in a bucket. Together with the maximum pedestal deviation the RMS pedestal deviation gives a statistical description of how the pedestal shifts look. This could be refined by adding more statistical moments.

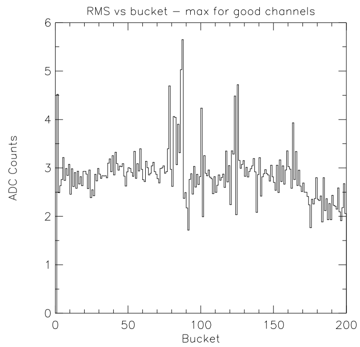

Maximum noise

In contrast to the above measures which all deal with averages for a bucket, the noise level describes a *variation* in the pedestal level for a certain channel in a certain bucket. This is simply RMS noise in the same sense as it was used for the measurements that are not dead-timeless. *Maximum noise* is the noise level of the noisiest channel in a bucket.

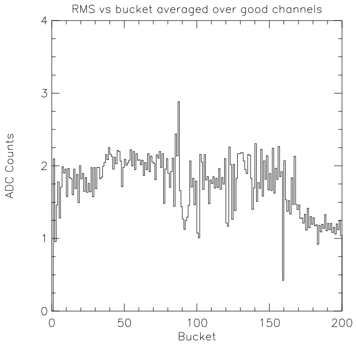

Average noise

Average noise is the average RMS noise in the pedestal level for good channels. Together with the RMS pedestal deviation it gives an overall measure on noise; unwanted signal that cannot be subtracted away.

4.3.2 Performance measures for a configuration

To determine how well the readout module works in a certain situation – for a given bias combination, in a grounding or shielding configuration – the above bucket-wise performance measures are combined into overall ratings. One set of ratings is calculated for each chip mode; Quiet1, Digitize, Readout and Quiet2.

The aggregate rating corresponding to bucket-wise pedestal shift is obtained by calculating the RMS of that measure over each bucket in a given mode. The ratings corresponding to the four remaining measures are taken as averages. For instance, the "Deviation in shift" rating is simply the average of the RMS pedestal deviation taken over all buckets belonging to the mode in question. Another measure of performance is the number of "good" buckets in a certain mode. In the plots 17–22 the fulfillment of all of the following criteria have been required to classify a bucket as good:

1. An upward pedestal shift of no more than 5 ADC counts.
2. A downward pedestal shift of no more than 10 ADC counts.
3. A sum in quadrature of RMS pedestal deviation and average noise of no more than 7 ADC counts.

In figure 23 the first two conditions have been relaxed to

1. An upward pedestal shift of no more than 5, 15, 25, 35, 45, 55 and 65 ADC counts, respectively.
2. A downward pedestal shift of no more than 10, 20, 30, 40, 50, 60 and 70 ADC counts, respectively.

The relaxation corresponds to dynamic pedestal subtraction. In figure 23 the calculation of the other performance ratings has also been changed, to include only the good buckets in each mode.

When counting the ratio of good buckets each chip is treated independently. This means that what is measured is the average number of bad buckets in a chip rather than in the readout module as a whole. To determine the latter, or the ratio of lost buckets in the detector as a whole, you need to know to what extent bad buckets overlap for different chips in the system, and what the tolerance for bad chips is.

4.3.3 Identification of noise sources

Figure 16 illustrates some of the things that happen as the readout module goes through different modes. The figure consists of two parts, one that shows data quality and one that shows the power consumption of the SVX3 chips. Time and bucket number run from left to right in the figure. The top panel shows how three measures of data quality change depending on what events takes place in the chip. There are events of different kinds; signals changing on the bus lines, activity in the digital part of the SVX3 chip and activity in the analog part of the chip. The topmost curve, that relates to the right hand y axis scale, shows pedestal shifts. The other curves relate to the left hand y axis scale. The solid curve is average noise, as defined in section 4.3.1. The dashed curve is the deviation from pedestal shift, as defined in section 4.3.1. In this panel time is measured in buckets. In the bottom panel one time division (box) corresponds to $5\,\mu s$. The two panels are synchronized, sharing their x-axes. This is possible since one bucket corresponds to $200\,ns$ (8 ticks of the $40\,MHz$ clock). Different points in time correspond to different activity in the SVX3 chip. In the section labeled "D" in the top panel, the chip is in digitize mode, the "R" means readout mode. Before the digitize mode is a period that will be referred to as "quiet1". After the readout follows "quiet2"

The bottom panel is an oscilloscope representation of the SVX3 chip's power consumption, as measured by the voltages on its power supply lines. Curve 1 is digital voltage, curve 3 analog voltage. Curve 2 is a signal known as *pipeline reset* that was used as a trigger for the oscilloscope, as well as to synchronize the two panels. A *drop* in a voltage curve corresponds of course to an *increase* in power consumption. Both curves are recorded in AC mode, meaning the DC level has been filtered away so that only changes in voltage remains. If the curves are integrated they will therefore sum up to zero. For the digital voltage the y scale is $50\,mV$ per box, for the analog voltage it is $10\,mV$.

Quiet1 The first thing that happens in figure 16 is the preamplifier reset. That occurs in the very first bucket. Resetting the preamplifiers means bringing their output levels back to their reference voltage. The preamplifier reset shows up as a huge pedestal shift and a small drop in the analog current. The two buckets affected by this are not included in the calculation of module performance, since it will be performed during *beam abort gaps* in the accelerator, when no events take place. A beam abort gap lasts 1–2 μs which corresponds to 5–10 buckets.

Digitize The next thing that happens is that the chip enters digitize mode. That means reading out a measurement from the analog pipeline and feeding it to the A/D-converters. As all operations of the chip, this is controlled by signals on the bus. The point in time when signals on two particular bus lines change is clearly seen in the form of pickup. A little later the comparators in the ADC are seen to fire. Since all channels are at the same pedestal level, they will do that at roughly the same time. Firing of comparators in this case does not create any pedestal deviation, only pedestal shift. That means all channels are symmetrically affected.

Readout After digitization the chip is told to go into readout mode, which means sending the digital output to the readout electronics over the bus. The data bus uses differential signals, but when it is first used, it must be charged up to whatever DC level the SVX3 chip uses. Because of this it initially draws a lot of current. This also shows up as pickup. In readout, there is one bucket that will always be bad. This is the bucket where a storage cell is returned to the pipeline sequence for writing, which is known as a PRD2 operation. Similarly to the pipeline reset, the operation will be performed during beam abort gaps and can therefore be disregarded when assessing performance.

Quiet2 In the beginning of the second quiet period, there is a sequence that is not so quiet. The first six buckets correspond to a brief cycling through digitize and readout modes (mini digitize-readout). The purpose of this cycling is to properly reset the dynamic logic of the data output memories, which could go to a high current state if the chip command lines are not exercised for a "long" time after the end of a normal readout cycle (the test hardware used runs map patterns at intervals of order 1ms). However, in an actual, continuous data acquisition situation such mini digitize-readout cycles may not be necessary, or may be done in a different way affecting a smaller number of buckets. The operations in those buckets are the same as in the actual digitize and readout periods.

4.3.4 Biasing

Total bias voltage Figures 17 and 18 show the dependence of chip performance on the total bias voltage. They represent the first quiet period, and the digitize mode, respectively, as those modes are indicated in figure 16. Performance in readout mode is very similar to that in digitize mode, and the second quiet period is similar to the first. The performance measures plotted in these and following figures are those defined in section 4.3.2. For all three operational modes it can bee seen that increasing the bias voltage improves performance. From the laser pulse tests it was found that the transition to full depletion occurs at 64 volts. That is the case in situation 3 in the plots. In situation 4, with a total bias of 96 volts the detector is overdepleted.

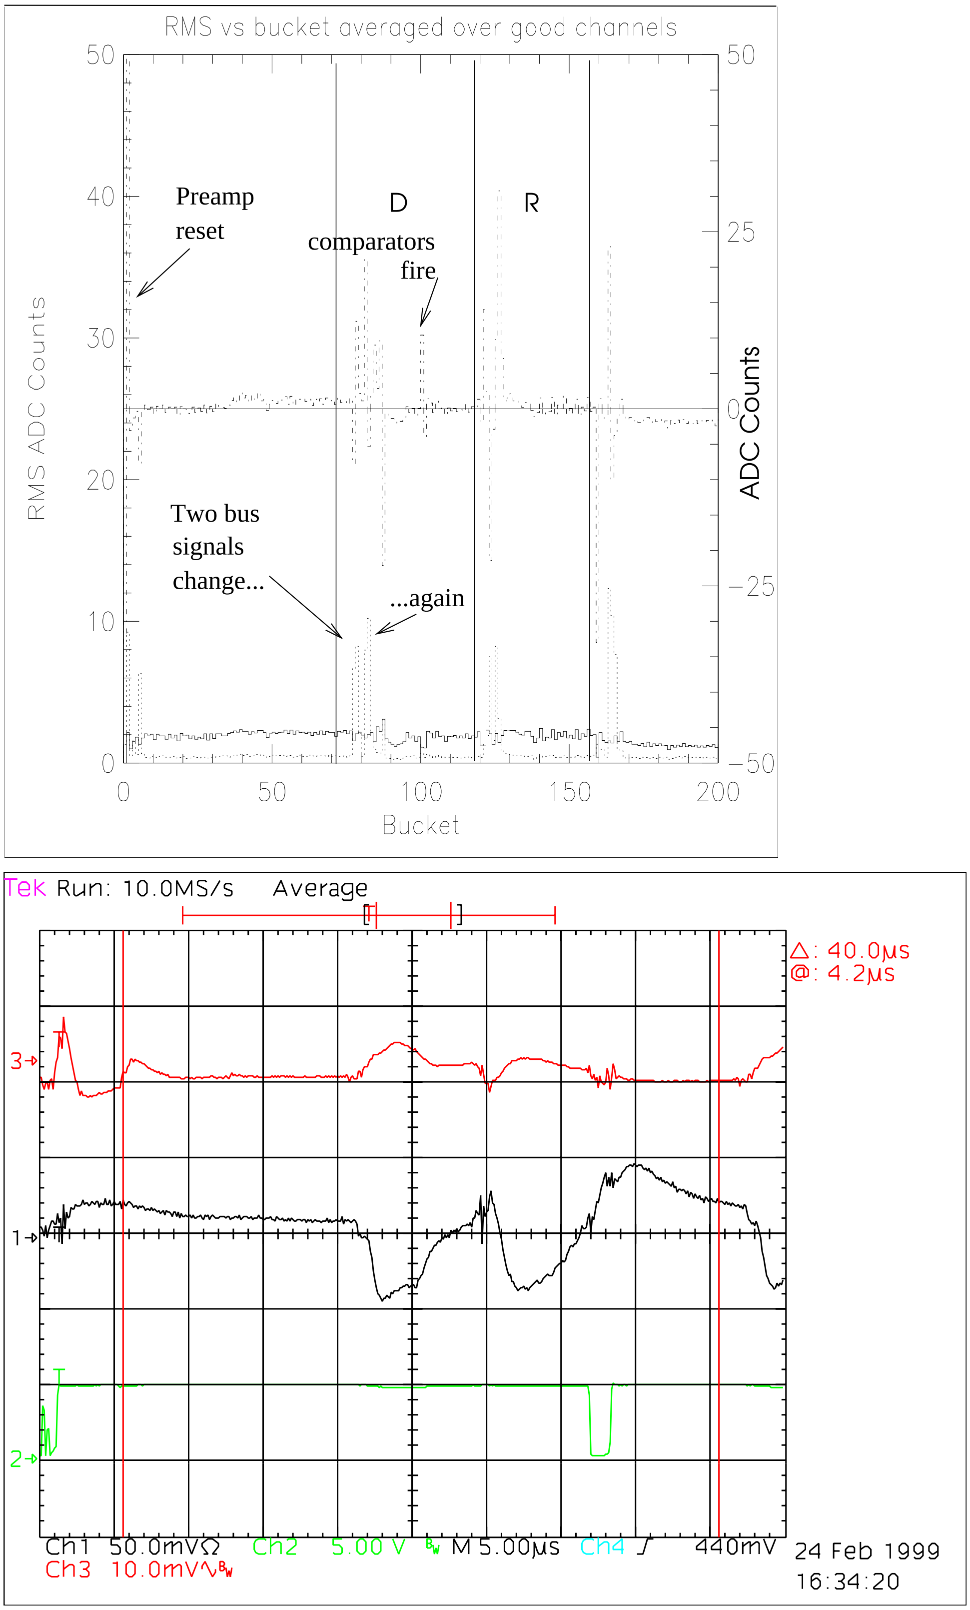

Figure 16: Overall performance in different modes. Top: Pedestal shift, pedestal deviation and noise. Bottom: Oscilloscope measurements of voltage supplies. The two panels are synchronized. The sensor is biased from the bias supply, using $48+48$ volts. A plastic chassis is used. For an explanation of what happens in this figure, see section 4.3.3.

Bias splitting The splitting of bias voltage has not been found to have any effect on the performance of the readout module. Figure 19 seems to show an increase in bad buckets with increasing p-side voltage. But on closer inspection of why buckets were classified as bad this is found to be a statistical coincidence. There may be some small effects of bias splitting on data, but none that affects overall performance in any significant way.

4.4 Grounding and Shielding

4.4.1 Setup

One source of electromagnetic pickup is activity in the data cable connecting the readout module to the readout electronics. In the test setup this problem is accentuated, since the cable lacks shielding. The cable also runs close to the silicon sensor, as was schematically shown in figure 10. In the finished detector, the cables will have built in shielding and will be folded back over the hybrid, rather than pass over the bare silicon. Results from attempts at both shielding and "folding" of the data cable are described below.

4.4.2 Results

Grounding and shielding schemes should be important for pickup noise, since they determine the coupling between noise source and signal path. In addition to the different grounding schemes presented in section 3, shielding of the data cable from the readout module was attempted here. For this purpose an aluminum foil was used. Also some different ways of grounding the shield were tried.

By standing the readout module on its edge it was possible to have the cable run at right angles away from the readout module. This substantially increases the separation between cable and sensor. The effect on a bucket in which there is activity on the data bus is seen in figure 21. Figure 20 shows some combinations of grounding, shielding and bending of the data cable. In contrast to figure 11, where a metal chassis was used, no extreme noise levels are observed. By comparing the situations within the sequences 1–7 (battery) and 8–11 (bias supply) it is evident that the presence of a shield improves performance, and the presence of a grounded shield even more so. In situations 6–7 and 10–11 the module was standing on its edge, permitting the data cable to be bent away from the readout module. The pickup reduction in those situation is therefore certain to pertain to the cable. As was expected, pickup reduction was greatest in the readout phase, when the data bus is active. For no combination of grounding/shielding was data quality found to depend on how the bias voltage was split between the two sides of the sensor. A difference was found between battery bias and supply bias. This can be attributed to a discharged battery, providing a voltage of only 50 volts, compared to the 85 volts for previous tests with battery bias.

4.5 Capacitive coupling

In section 3 it was found that noise entering through capacitive coupling to the sensor surface can cause severe problems. To further investigate this under controlled circumstances, tests have been carried out where a piece of metal was placed on top of the sensor. From these tests it was clear that noise was indeed introduced by this mechanism – added noise only appeared on the channels where the metal piece touched. To get an unambiguous noise signal the metal piece had to be pressed to the sensor surface, to make good capacitive contact. Because of the difficulties in doing that in a consistent way, the metal pieces were not pressed in the data represented in the figures below. Unfortunately this resulted in very poor coupling. The metal pieces used were a 23x36x1 mm aluminum piece and a 76x76x0.5 mm copper piece. The latter was expected to couple better because of its larger weight creating a better contact.

By picking up electromagnetic noise and capacitively coupling it to the sensor, the presence of a metal piece on top of the silicon was expected to increase pickup. Different ways of shielding from this increased pickup were then tried. Figure 22 compares otherwise identical situations with and without a metal piece in order to visualize how much extra pickup is generated by its presence. Without a grounded shield, the pickup mediated by the the aluminum piece can be seen as an increased pedestal shift (situations 2 and 4). With a grounded shield, the pickup can no longer be seen, not even that from the copper piece.

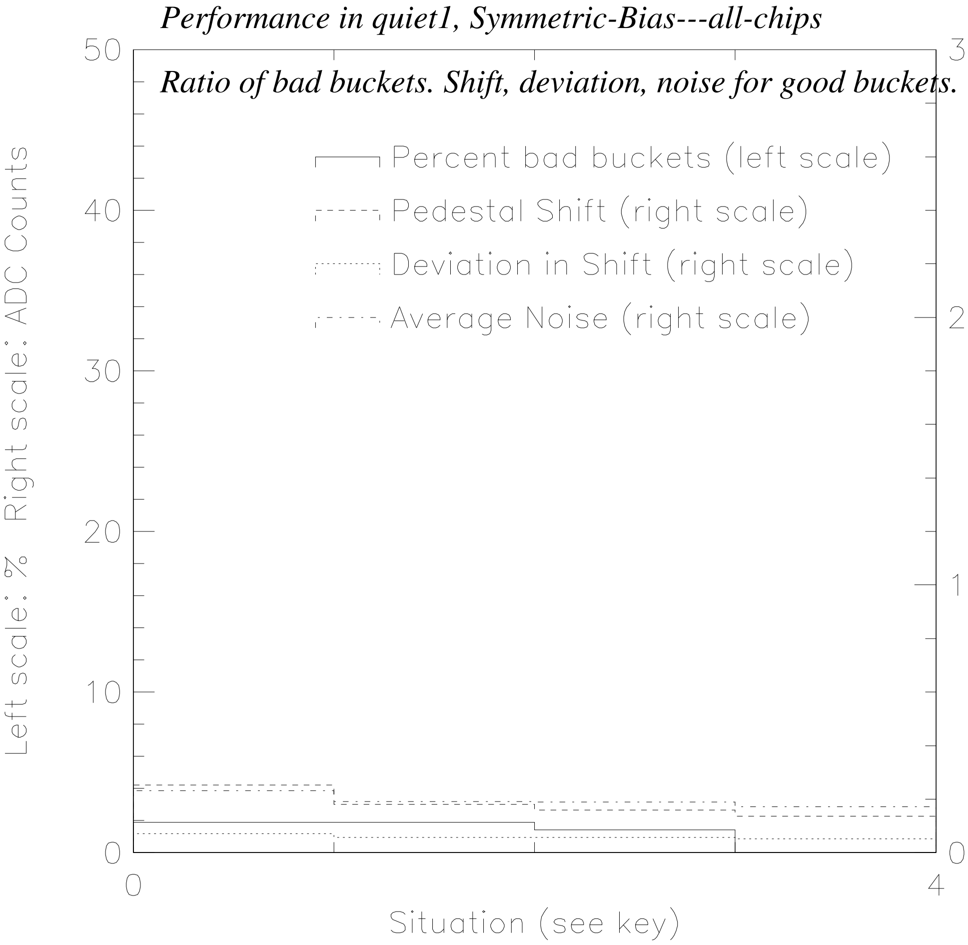

Figure 17: Performance of the test readout module in four bias situations. From left to right: Bias voltage $0 + 0$, $16 + 16$, $32 + 32$ and $48 + 48$ volts. The solid line is the percentage of rejected buckets. The other curves are the aggregate performance indicators. Note that the solid line refers to the left-hand y-scale, the others to the right-hand y-scale. This plot takes both the ϕ-side and the z-side into account. The plot shows performance for the first quiet period as defined in figure 16. The first two buckets (preamplifier reset) are not included.

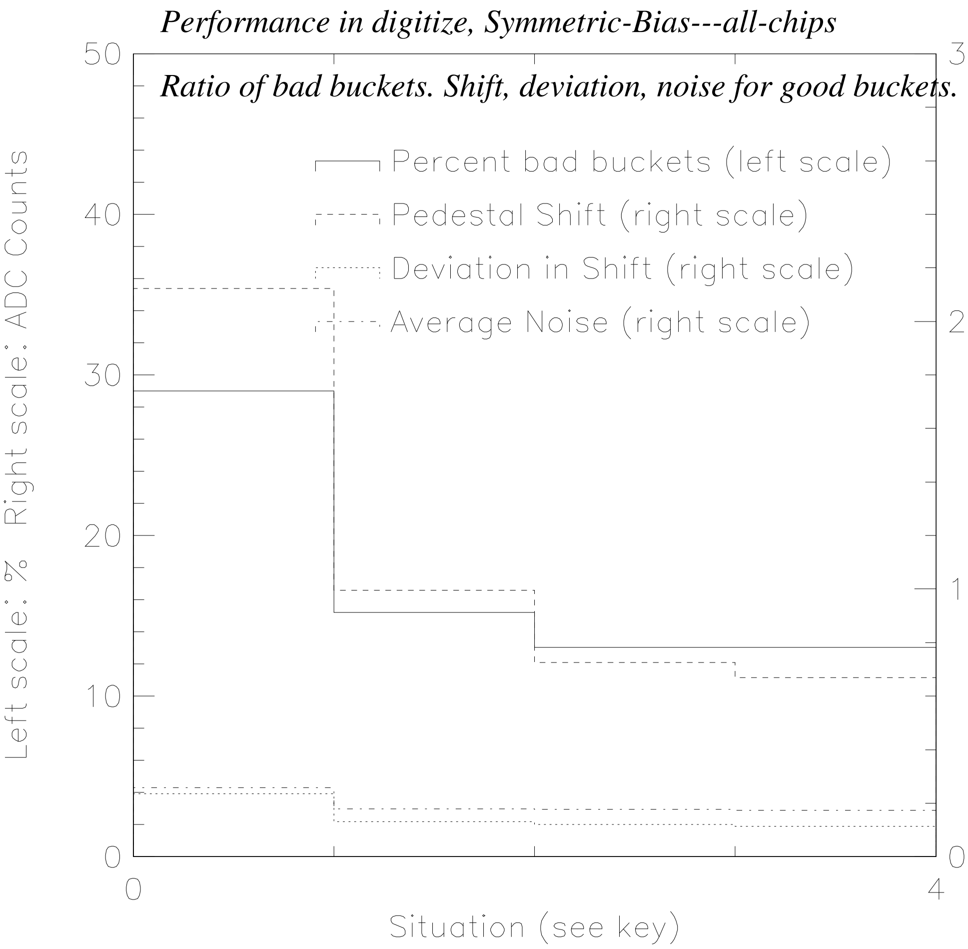

Figure 18: Same as figure 17, but for data taken while the chip was in digitize mode. Refer again to figure 16 for the definition of the modes. Note that performance is much worse compared to data taken with a quiet chip.

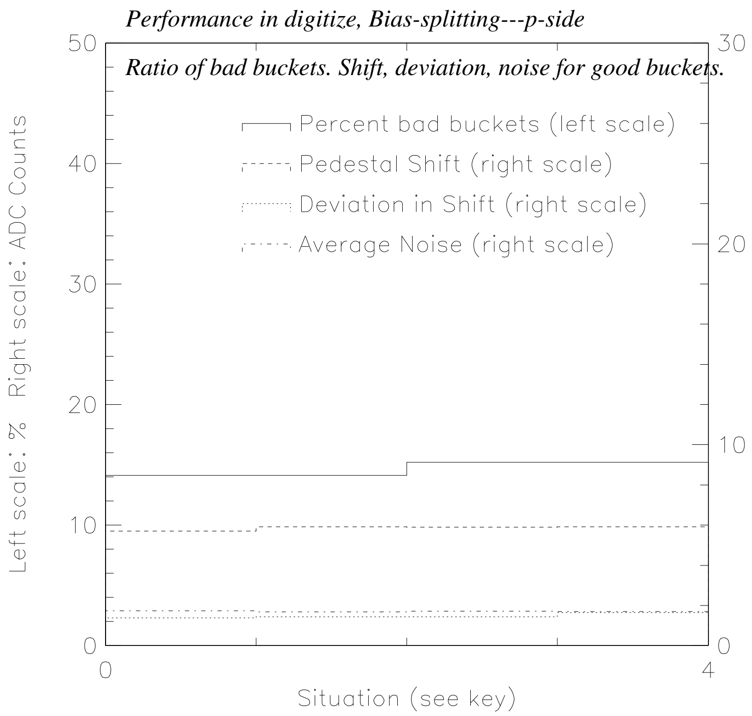

Figure 19: Performance in digitize mode. This time we are only looking on the ϕ-side. Unlike what was the case in the previous figures, total bias voltage is held constant at 80 volts. Instead it is the splitting of those 80 volts between the two sides of the sensor that is varied. For the four situations on the x axis in the plot, the bias voltages are:

Situation #	1	2	3	4
n-side	64	48	32	16
p-side	16	32	48	64

The apparent small rise of the number of bad buckets is an illusion – it is due to statistical effects. Thus, no dependence of performance on bias splitting was found in this test.

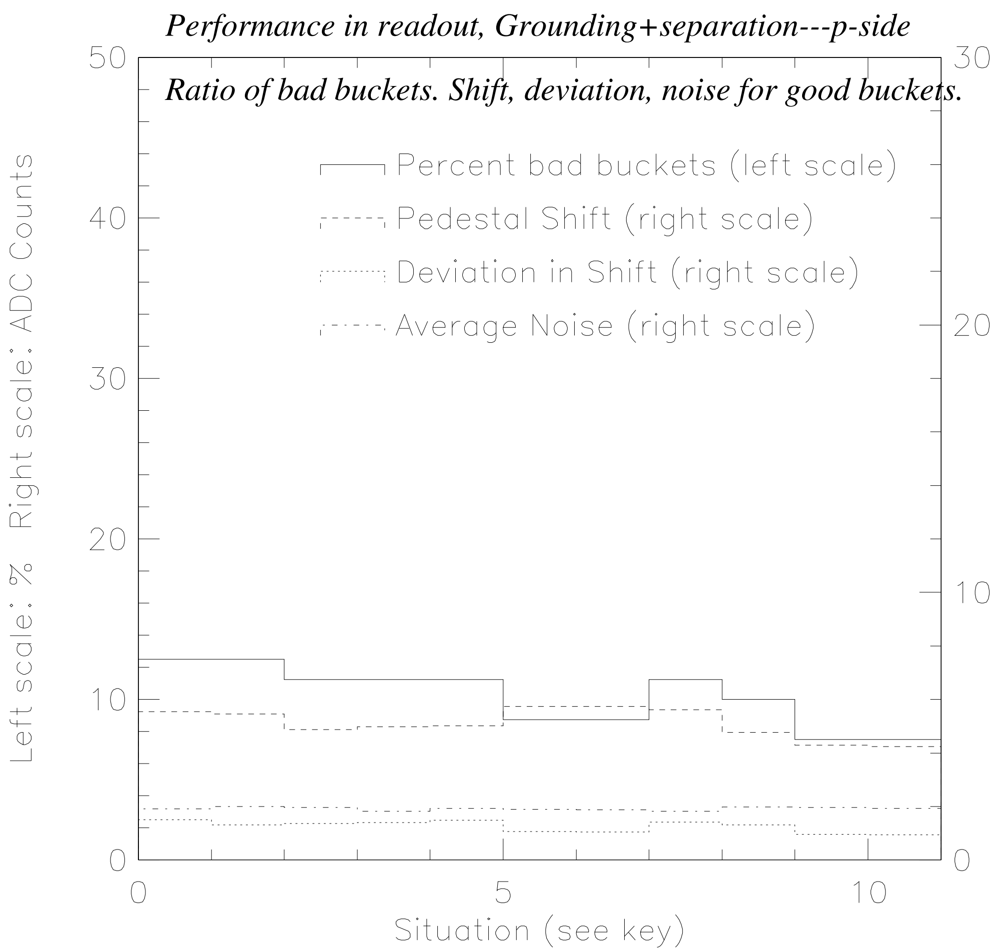

Figure 20: Chip performance for different ground configurations. Key to situations on the x axis:

Situation #	1	2	3	4	5	6	7	8	9	10	11
Source: battery/supply	b	b	b	b	b	b	b	s	s	s	s
Shield?	–	√	√	√	√	√	√	–	√	√	√
Shield grounded?	–	–	√	√	√	√	√	–	√	√	√
Bias ground to shield:	–	–	–	near	far	–	near	–	–	–	near
Cable close to silicon?	√	√	√	√	√	–	–	√	√	–	–

The shield referred to in the table is a piece of aluminum foil wrapped around the data cable. The shield was in most cases grounded to the readout module ground at the end closest to the readout module. In addition, the bias ground was connected to the shield in some cases, either to the end closest to the readout module, or furthest from it. This figure shows performance in readout mode for the ϕ-side of the sensor.

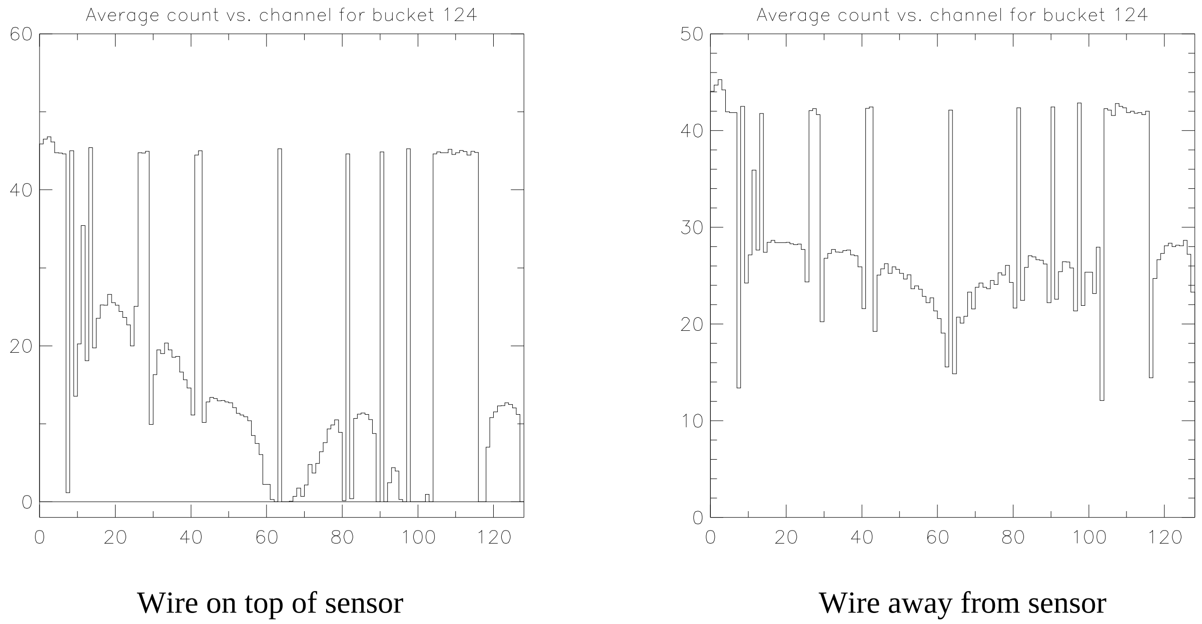

Figure 21: Bucket 124. On the left hand side with the data cable close to the sensor, on the right hand side with the cable away from the readout module. Bias 48 + 48 volts from the bias supply. Data relating to one ϕ-side chip.

4.6 Dynamic pedestal subtraction

The dynamic pedestal subtraction is a function of the SVX3 chip that removes common mode pedestal shifts. When the mechanism is working properly such shifts are completely eliminated. This means that data quality improves for all buckets, and that buckets that were previously rejected for having too large a pedestal shift can now be used. There are some limits on how large a pedestal shift can be taken care of by the dynamic pedestal subtraction. The mechanism works by starting a counter at different times inside a certain interval, and an excessive shift will move the desired starting time outside the available interval.

The DPS circuit was not correctly implemented in the SVX3C prototype chip and is therefore not available in the readout module tested here. However, the performance of the DPS circuit in the SVX3D production chip has been well measured [4], and it is straightforward to re-analyze the data presented so far for the case that DPS had been enabled.

In figure 23 the combined benefits of DPS and of separating the data cable from the readout module are shown. The situations on the x-axes in the plots relate to performance in cases where the DPS is assumed to be able to deal with shifts of between 0 and 60 ADC counts in any direction. Also shown in the figure are the usual performance ratings, but with a different definition than previously. Since this figure is intended to illustrate the actual performance of the readout module, only the accepted buckets have been included in the calculation of performance. One therefore would expect performance to deteriorate as more buckets are classified as good. This is clearly the case for the pedestal shift. Note however that this shift exists only inside the chip; the DPS removes it from the digital measurements. The other indicators do not change noticeably as more buckets are included. This would indicate that the buckets that are reclassified as good, as a consequence of allowing larger pedestal shifts, do not bring with them significant increases in noise or pedestal deviation. There *is* such an effect, but it is quite small, and it is diluted when averaging over many buckets as is done in the figure.

The number of bad buckets for two situations: a non-working DPS and a DPS capable of dealing with shifts of 30 ADC counts, are compiled from the figure into tables 1 and 2. The percentage of bad buckets is also the dead time for a chip in the readout module. For each chip, there is a total of 200 buckets. 8 of those are not

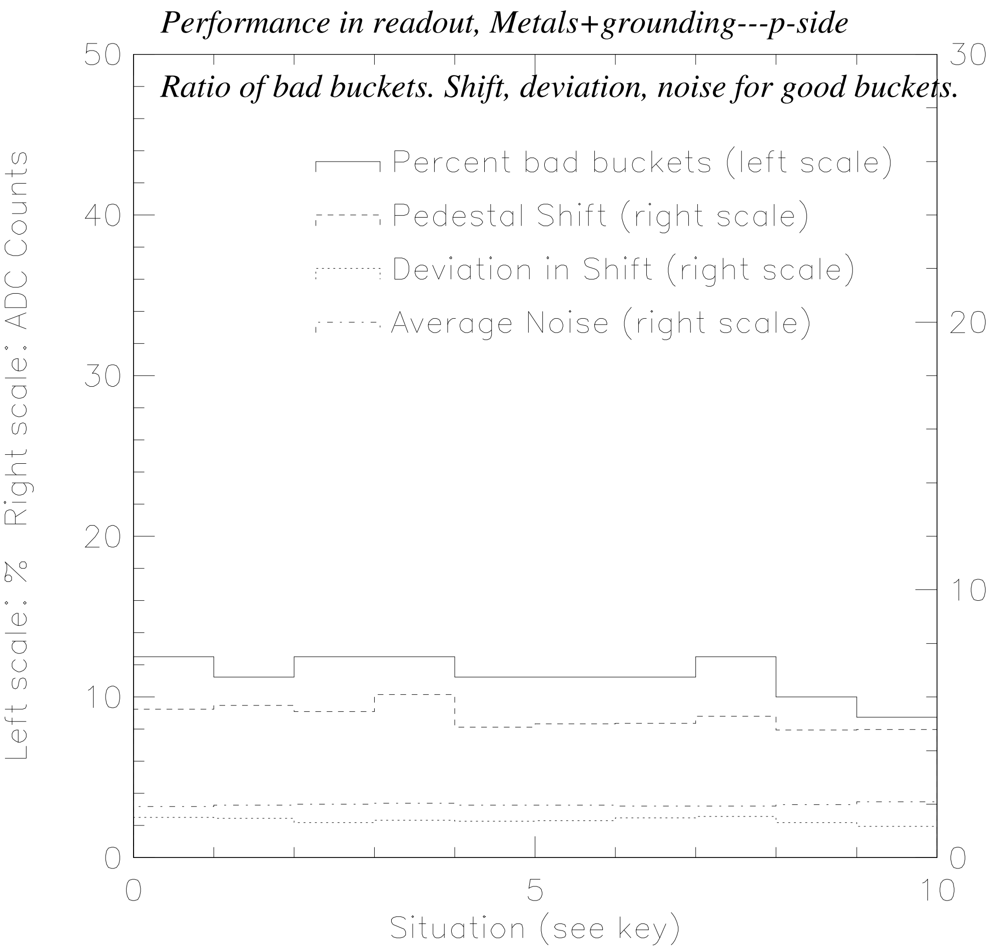

Figure 22: Capacitive pickup. Comparison of performance in digitize mode with and without metal pieces.

Situation #	1	2	3	4	5	6	7	8	9	10
Metal piece:	–	Al	–	Al	–	Al	–	Al	–	Cu
Source: battery/supply	b	b	b	b	b	b	b	b	s	s
Shield?	–	–	√	√	√	√	√	√	√	√
Shield grounded?	–	–	–	–	√	√	√	√	√	√
Bias ground to shield:	–	–	–	–	–	–	far	far	–	–

Mode	Quiet1	Digitize	Readout	Quiet2	Total
Number of bad buckets	1	15	10	1	27
Total buckets in mode	213	138	120	99	573
Percentage bad buckets	0.5%	10.9%	8.3%	1.0%	4.7%

Table 1: Lost buckets without DPS.

Mode	Quiet1	Digitize	Readout	Quiet2	Total
Number of bad buckets	0	1	2	0	3
Total buckets in mode	213	138	120	99	573
Percentage bad buckets	0%	0.7%	1.7%	0%	0.5%

Table 2: Lost buckets with DPS capable of removing a shift of 30 ADC counts.

included; they correspond to preamplifier reset and mini digitize-readout operations.[3] Since each chip is treated individually, the total number of buckets is $3 \times (200 - 8) = 573$. The dead time calculated here is the average dead time for a *chip*. This means that the dead time for the *readout module* as a unit will be slightly higher.

4.7 Conclusions

The pickup from chip and bus activity remains the most serious problem. The resulting pedestal deviations are by far the most powerful source of noise for the affected buckets. For some buckets the problem may be so large that they have to be discarded, introducing dead time into the whole detector system. It is of great importance to find shielding– or other techniques to reduce pickup.

In theory, since the largest part of the pedestal deviation was found not to fluctuate between measurements, it could be subtracted away. If pickup patterns could be linked to activity on the bus, for instance, the off line system could model the bus activity and apply the appropriate corrections. This scheme has three difficulties. One is that not all activity can be perfectly modeled (the firing times of converters in the ADC are not known). The second is that it may be necessary to investigate the effect of every bus line on every channel in the entire detector to build the model. The third is that off line analysis can never correct for incorrect sparsification. Channels that are not read out cannot be recreated and incorrectly read out channels will still contribute to readout time.

The Dynamic Pedestal Subtraction (DPS) recovers buckets that are bad due to pedestal shifts. When the mechanism works properly, such shifts are completely eliminated. There are some limits on how large shifts can be eliminated set by the timing of the digitization. With DPS able to remove pedestal shifts of 30 ADC counts the best case performance of the chips in the test readout module was found to be 99.5% good buckets. That is a 90% reduction of the number of bad buckets compared to no DPS. Performance will be worse if DPS cannot be used, if the data cable runs close to the sensor or if there are metal parts capacitively coupled to the sensor.

References

[1] Martin Rehn, "Tests of double-sided SVX II sensors", **TRITA-FYS 9035**, Royal Institute of Technology, Stockholm, Sweden (1999).

[2] Christopher Tollnert, "Testing of Readout Electronics for the Silicon Vertex Tracker of CDF", **TRITA-FYS 9026**, Royal Institute of Technology, Stockholm, Sweden (1998).

[3] T. Zimmerman *et. al.*, "SVX3: A Dead-Timeless Readout Chip for Silicon Strip Detectors", Nucl. Instr. and Meth. **A409**:369, 1998.

[3]The first 2 buckets in quiet1 are taken out in quiet1 due to pipeline reset, The first 6 in quiet2 due to the mini digitize-readout cycle.

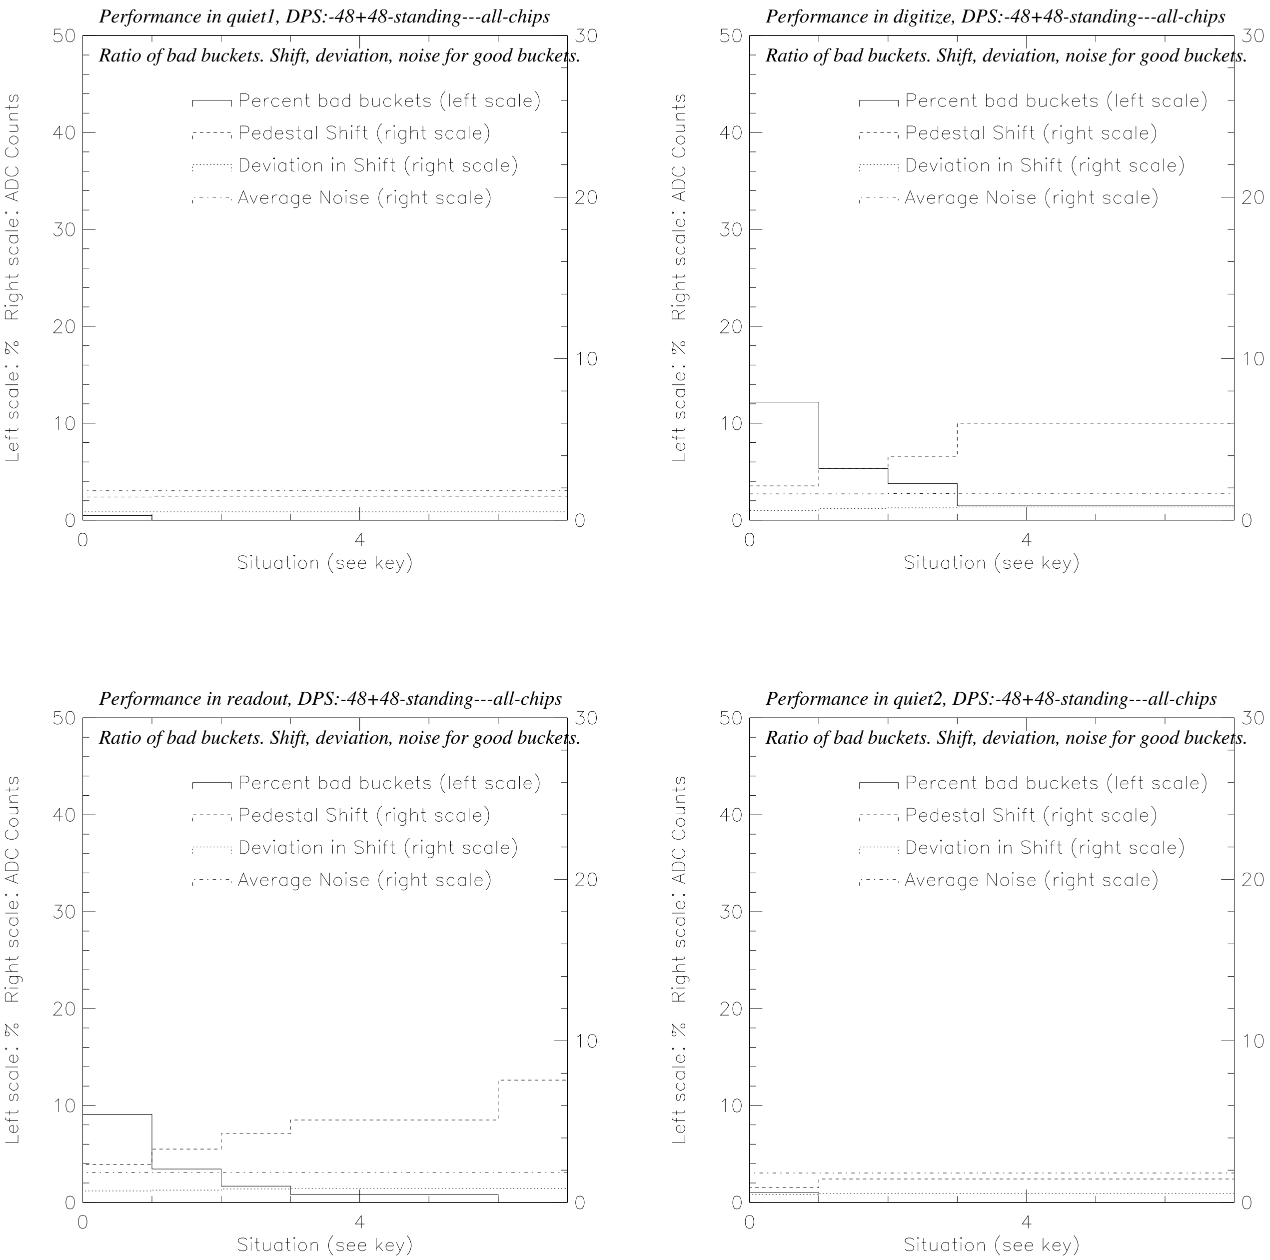

Figure 23: Here is overall performance in all four chip modes. The test module is fully biased, and standing on its edge, in order to move the data cable away from the silicon. The DPS is supposed to be capable of removing a pedestal shift of 0, 10, 20, 30, 40, 50 and 60 ADC counts in any direction in the respective situations. **Note:** The performance measures (pedestal shift, pedestal deviation and noise) have a different definition in this figure. Here they relate only to the accepted buckets. This is why the pedestal shift rises as more buckets are accepted. Otherwise those measures would of course be independent on how buckets are classified, as was the case in previous figures.

[4] Martin Rehn *et. al.*, "Statistical Study of SVX3D Chip Dynamic Pedestal Subtraction Threshold Level", CDF Note **4852**, 1999.