Tests of double-sided SVX II sensors

Master's thesis

Martin Rehn

May 1999

Abstract: Tests of the SVX II readout module in deadtimeless and non-deadtimeless modes are described. The effects of pinholes, biasing, grounding, capacative coupling and electrical shielding on signal quality are investigated. Signal quality is related to detector performance. The potential of dynamic pedestal subtraction to improve performance is simulated.

Royal Institute of Technology, Stockholm
Lawrence Berkeley National Laboratory, University of California at Berkeley

TRITA-FYS 9035

Thanks to my supervisor Maurice Garcia-Sciveres, Tony Affolder, Igor Volobouev and the other people with whom I worked at Lawrence Berekeley Lab, for all your help, and for making my stay in Berkeley a pleasant one. Thanks to Mats Danielsson and my examiner Per Carlson, both at KTH, for making this possible.

Contents

Chapter 1

Introduction

1.1 Goals

The goals for this work are

1. To test the double sided SVX II readout module
2. To determine the effects of biasing schemes on performance
3. To investigate deadtimeless operation
4. To simulate dynamic pedestal subtraction

1.2 The CDF Experiment

The *Collider Detector at Fermilab* (*CDF*) is one of two detectors installed on the *Tevatron* particle accelerator at Fermilab. The Tevatron became operational in 1984, the CDF the following year. Both are currently undergoing major upgrades in preparation for *Run II,* scheduled to begin in the year 2000. The Tevatron is an accelerator for protons and anti-protons. Its name derives from its ability to accelerate those particles to energies approaching $1\,TeV$, or one trillion electron volts.[1] When the particles collide head on the total energy available for physics reactions is twice that, or $2\,TeV$.

The Tevatron accelerates particles grouped together in *bunches.* For Run II there will initially be 36 bunches of protons traveling clockwise in the accelerator ring, and an equal number of anti-proton bunches going in the opposite direction. A pair of bunches will meet inside the collider detector every 132 nanoseconds. A measure of particle density in the accelerator is its *luminosity.* The luminosity for Run II will initially be $5 \cdot 10^{31}\, cm^{-2}s^{-1}$. This is eventually to be quadrupled. At full luminosity each bunch crossing will on average result in 5.8 collisions between a proton and an anti-proton.[2] The resulting reaction frequency is a demanding specification for the CDF detector systems.

The task of detectors in high energy physics experiments is to track and identify particles emerging from a collision, allowing physicists to reconstruct what reactions took place. To accomplish that task, detectors have evolved into huge structures surrounding the interaction point. From the inside out, silicon detectors track particles' paths with great precision, scintillators detect them by light emitted as they blaze through certain materials, drift chambers measure their movement in a magnetic field and calorimeters measure the remaining kinetic energy. The information is used to determine the mass, charge and point of origin for each particle. Everything contributes to reconstructing the physics event. In figure 1.2 the components of the CDF are shown together with two humans.[3]

[1] The Tevatron will not fully live up to its name until Run II. Currently, the most powerful in its kind in the world today, it produces only about 0.9 TeV.

[2] Data according to the *Run II Handbook.* These are the current design specifications for Run II. Actual performance will vary with time.

[3] I apologize for the failed attempt at political correctness in this figure. The woman's clothing was obviously not chosen to be practical in the industrial environment. Also the man is shown in movement, the woman in a passive pose.

Figure 1.1: The accelerators of Fermilab. The *Tevatron* is shown together with the *Collider Detector*.

Figure 1.2: A schematic showing the size of the CDF detector. The SVX II system is right in the center, next to the beam pipe inside which the particles collide.

1.3 The SVX II Detector

The *Silicon Vertex Detector* (*SVX II*) is the innermost part of the CDF tracking system. It operates together with the similar *Intermediate Silicon Layers* (*ISL*). See figure 1.2. The name, *Vertex Detector*, means that SVX II is capable of distinguishing *secondary vertices* from the *primary vertex*, the point where the proton and the anti-proton collided. This ability was critical when the *top quark* was discovered in the CDF experiment. The software in the *Data Acquisition System* of SVX II and ISL first reconstructs *traces*, the paths of particles. Then it follows those traces backwards to determine a possible common origin, or vertex. The SVX II will consist of five *layers* of detectors that accurately trace particles in three dimensions. The layers are numbered from 0 to 4, starting from the inside. The sensors are mounted in three concentric *barrels* as seen in figure 1.3.

1.4 Physics Goals

The most famous discovery so far at the CDF experiment is of course that of the top quark. The top quark was the last missing particle in the generally accepted Standard Model of particle physics. In Run II, investigations of that newly discovered particle will continue with precision measurements of top mass, the top quark's weak interactions and the cross section for top production. The second heaviest quark, the bottom quark, is being investigated as well. Those experiments may answer questions about *CP*-violation, a phenomenon which means that a mirror image of nature, where all charges are reversed, would not work exactly the same as the universe we know.

Theorists and experimentalists participate in several groups preparing questions to ask to the Run II experiments. One group will look for supersymmetric particles, particles predicted to exist by a group of theories that promise to solve some disturbing problems in the Standard Model. They will also look for the much-wanted Higgs boson, which is predicted to exist by supersymmetric models as well as by the Standard Model. The Higgs mass, and therefore the difficulty in creating and detecting it (or ruling it out), depends on which supersymmetric model you believe in. Already the CDF has put some restrictions on the existence of both Higgs bosons and supersymmetric particles.

Another group looks into the dynamics of the strong force, exploring alternatives to the standard "chromodynamics" that are believed to rule the world of quarks, such as a family of theories known as technicolor models. Some of the issues listed here may be answered by Run II, some will not. The latter may have to wait for the Large Hadron Collider to become operational at CERN, at which time even higher energies will become available than the Tevatron has to offer.

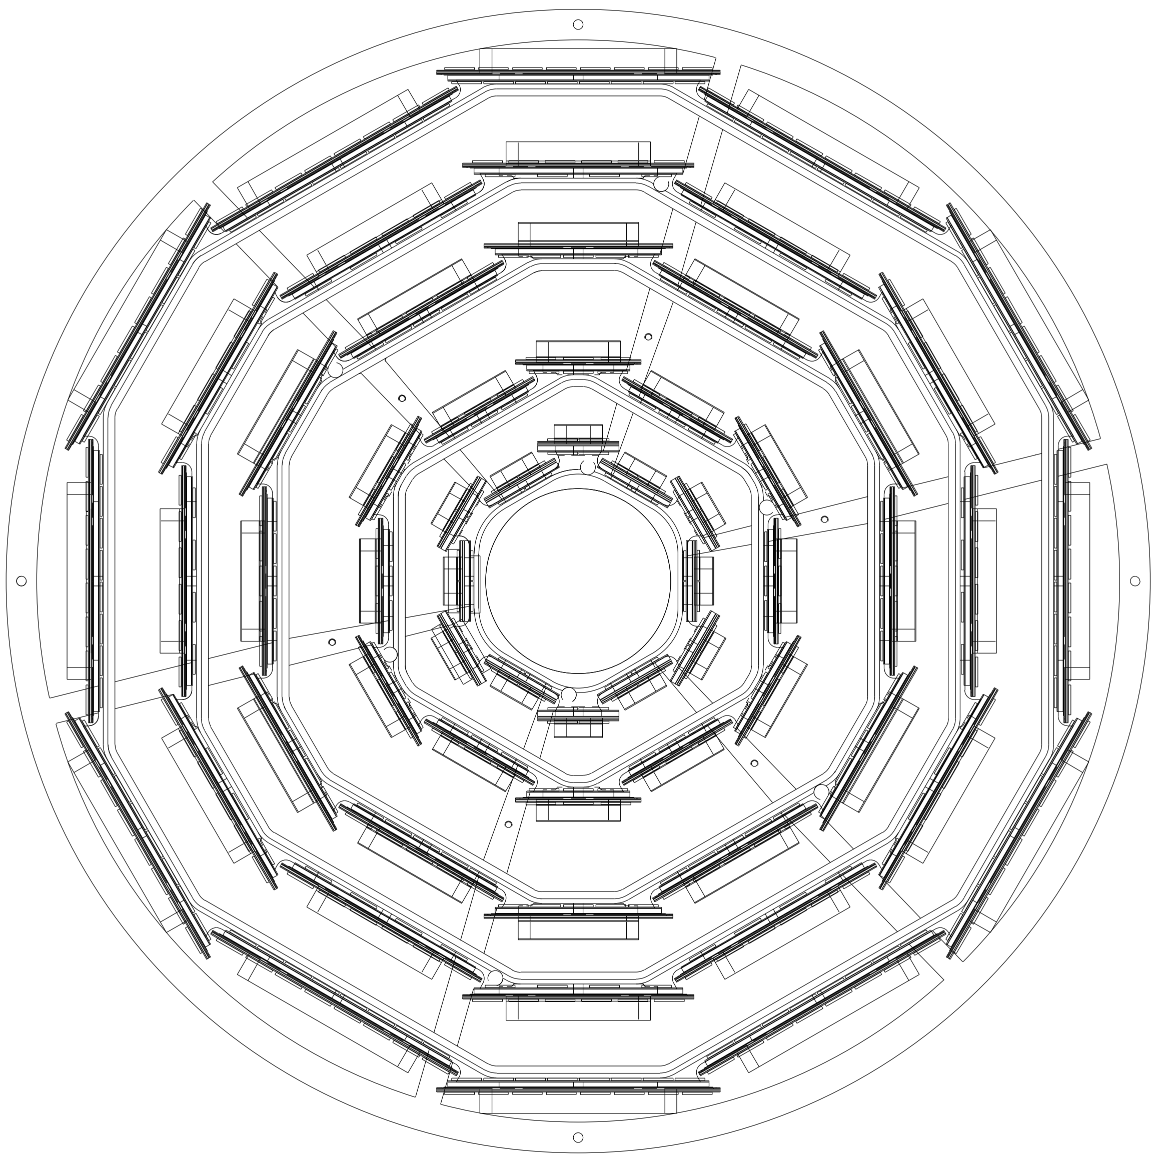

Figure 1.3: End view (looking along the beam pipe) of the five layers of sensors in the SVX II system. The diameter of the innermost sensor layer is 4.8 cm, for the outermost layer the diameter is 17.4 cm.

Chapter 2

Silicon Detectors

2.1 Characteristics of the SVX II Sensors

The SVX II sensors belong to the popular class of *silicon strip sensors.* There is also another class called *pixel* sensors. The sensors are *double-sided.* The sensors are read out by *integrated chips* containing *both analog and digital parts.* The readout chips permit *deadtimeless operation*, which means they can collect new measurements without having to wait for the old ones to be read out. This reduces the number of lost events. Data is read out in *sparse mode,* meaning that only channels that contain signal are read out. The use of a dead-timeless detector and the necessity for sparse readout arise from the high luminosity and short crossing time of the Run II Tevatron. Sparse readout creates certain problems that are addressed by a feature known as *dynamic pedestal subtraction.*

2.2 Operation of Silicon Strip Sensors

2.2.1 Basic Principles

Semiconductors

Silicon is a semiconductor, meaning the conductivity of the pure element is in between that of insulators and good conductors (such as metals). The reason is that semiconductors have a very limited number of *charge carriers,* mobile charged particles that conduct electricity. Unlike the situation in a metal, the vast majority of electrons in silicon are not mobile, but bound to fixed positions.

There are two kinds of charge carriers in a semiconductor. One kind is *free electrons*, which carry negative charge. The other is *holes*, positions in the crystal where an electron is missing. The hole can be regarded as a positively charged particle. It can "move" by means of a neighboring fixed electron falling into the available space. For instance, if the right-hand neighbor electron is pulled into the hole, the net effect will not be that the hole is destroyed, but that it moves one atom to the right. Holes and free electrons *can* be destroyed, if they run into each other. The free electron will fill the hole and both charge carriers are lost. This process is called *recombination.*

If you keep your pure silicon crystal in the dark near absolute zero, there will be neither free electrons nor holes. Heat, light, or –in the case of particle sensors– highly energetic particles, will add enough energy to a few fixed electrons to promote them to free charge carriers — generating an equal number of holes in the process. There is a way to have more electrons than holes, or vice versa, namely by *doping* the silicon with small amounts of other elements. Silicon that is doped to have a surplus of electrons is called *n-type.* The opposite, a surplus of holes, is called *p-type.* Doped silicon is a much better conductor than is pure silicon.

Silicon sensors

A silicon strip sensor is a diode array. All we need to know about silicon diodes in order to understand sensors is the following. A diode contains a junction between p-type and n-type silicon, a *pn-junction.* Both types are conductors in themselves, but a diode will only conduct electricity in the *forward* direction, as shown in figure 2.1.

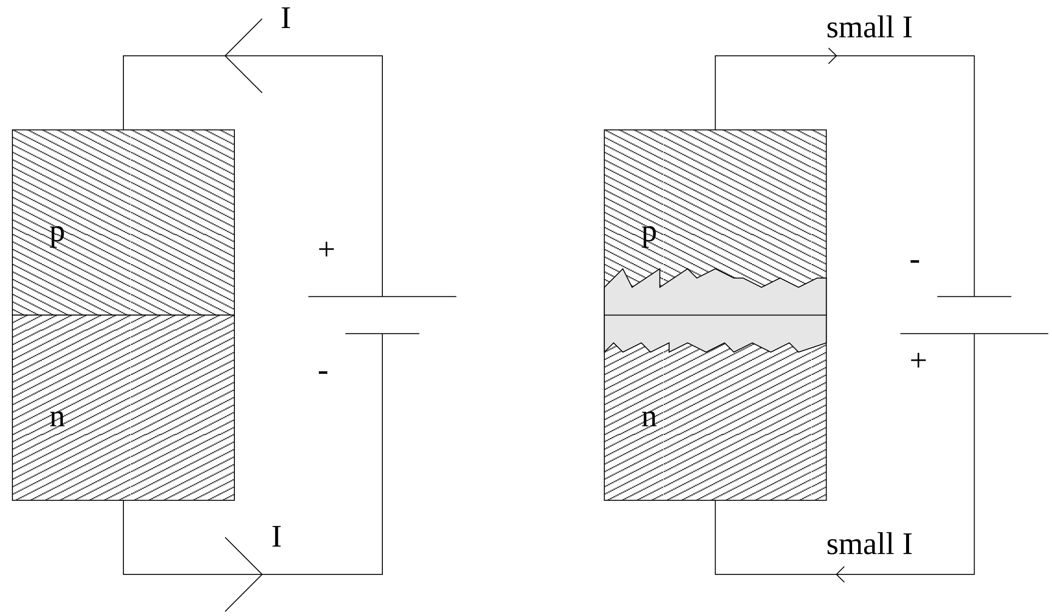

Figure 2.1: The left hand diode is *forward biased* and conducts a large current. The right hand diode is *reverse biased* and conducts only a very small *leakage* current. At its pn-junction a *depletion region* (gray) has developed.

In the *reverse* configuration there is a *depletion region* that is devoid of charge carriers. It acts as an insulator, preventing current from flowing. The depletion region gets larger the more reverse voltage is applied.

Silicon sensors operate in this reverse bias mode. The principle of operation is to detect particles by the charge carriers they create in the depletion region. The charge thus created is free to move, and propelled by the electric field from the bias voltage it creates a brief current pulse.

The SVX II sensors are constructed from n-type silicon. All of the silicon bulk is depleted by applying bias voltage between the strip implants – "n+" on the bottom and "p+" on the top in figure 2.2 (the plus signs represent a large amount of doping). The bias voltage also creates the electric field, shown as "E". The small "p+" strips on the bottom are known as p-stops. Not connected to anything, their function is to insulate the n-strips from each other.

When an ionizing particle passes through the silicon it may interact one or more times with the silicon atoms, ripping electrons out of their positions. One such interaction is symbolically represented by the big star in the figure. The charge carriers will drift in the electric field – the positively charged holes in the direction of the field, the negatively charged electrons in the opposite direction. Some will meet with an unfortunate fate, an electron will come upon a hole or vice versa. Then *recombination* will take place, resulting in a loss of signal (the small star in the figure). Recombination is the reason why silicon strip sensors are not made arbitrarily thick. If the thickness is increased beyond the mean distance that a charge carrier will travel before recombination occurs, losses will increase rapidly. It turns out that the mean distance before recombination is actually the optimal choice for sensor thickness. For silicon this is $300\,\mu m$. Therefore the SVX II sensors are $300\,\mu m$ thick.

The charge carriers that don't recombine will safely arrive at an implant strip. Remember that the electric field that propels them originates from the bias voltage applied to the strips. Therefore they will be directed towards one of the strips at the end of their journey.

By diffusion, the charge carriers created by one particle spread out and may arrive at more than one strip (as is shown in the figure). This is to be desired. The signal distribution among the strips can then be fitted to a model of diffusion, and the position of origin be determined, to greater precision than the strip spacing.

2.2.2 Double sided sensors

A double sided sensor is one where charge is collected and measured on both sides of the silicon plate. This can be used to measure two coordinates with just one sensor. Layers 0, 1 and 3 in the SVX II detector have the strips on the two sides of their sensors oriented perpendicular to each other to do just that. In layers 2 and 4

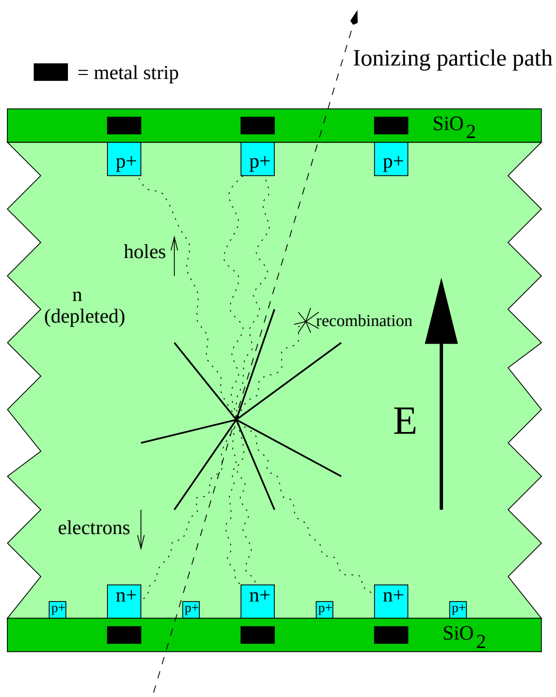

Figure 2.2: A cut through an SVX II silicon sensor. Uppermost is the p- or ϕ-side. At the bottom is the n- or z-side. On both sides the implant strips run into the figure, which tells you that this has to be a sensor from either layer 2 or layer 4 in the detector. Also shown are the processes of charge generation and -collection, which are explained in section 2.2.1. Not to scale. The silicon is 300 μm thick, the implants are 1.5 μm deep, and the SiO_2 layers are 3 μm thick. The separation between strips varies between 60 μm and 141 μm.

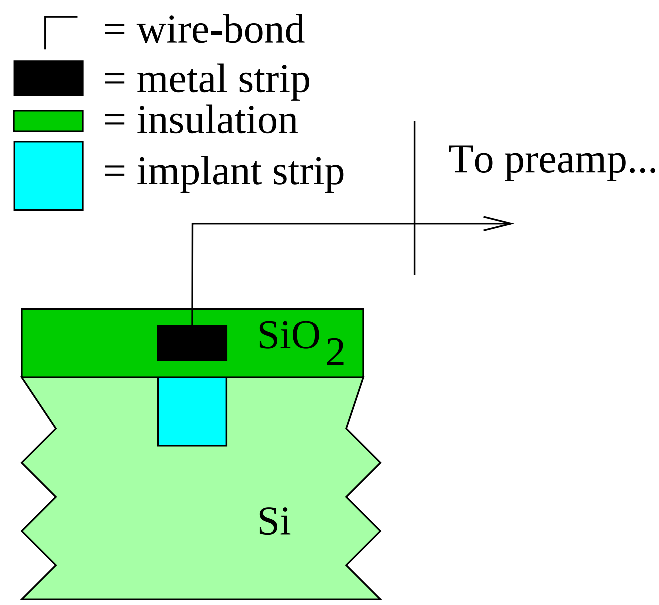

Figure 2.3: A close-up on an implant strip together with its metal strip in the sensor. The implants are separated by a very thin layer of insulating material from the metal strips. The latter are connected to the readout electronics. Together the two strips form a capacitor, coupling AC signal from sensor to electronics, but separating their DC levels.

the strips are put at a small angle (1.2°) relative to each other. In this way, both sides of the sensor contribute to the precision measurement of one coordinate, while some information is still obtained about the other. This configuration is called *small angle stereo* because of its similarity to stereophonic perception.

For SVX II, the ϕ-coordinate, the angle in a plane perpendicular to the beam pipe, is the most important coordinate. For this reason ϕ is given priority over the other coordinate in the plane of the sensors, z, which is the position along the beam pipe. The remaining cylindrical coordinate, distance from the beam pipe (r) is fixed in a layer. The two sides of the sensor are referred to as ϕ-side and z-side, depending on which coordinate they measure. The p-side of any sensor measures ϕ, the n-side measures z. p-side and ϕ-side are therefore used synonymously, and the same goes for n-side and z-side.

2.3 Readout Electronics

2.3.1 The SVX3D chip

A readout chip has been developed specifically for the silicon sensors in the upgraded CDF detector. Known as the SVX3D chip, it is used in the ISL detector, as well as in the SVX II. The "3D" suffix has nothing to do with spatial dimensions, but is a revision number.

One chip handles 128 channels. Depending on the size of a sensor, 4–14 chips are connected to it. The smallest sensors are found in the innermost layer, the largest in the outermost. The chip consists of an analog *front end* and a digital *back end*. The silicon sensor itself (the *sensor*) is mounted along with the chips on a specially designed circuit board, a *hybrid*. Two silicon parts are electrically connected to one set of chips to form an *SVX II readout module*.

2.3.2 Charge collection

The first step in converting the charge collected on the implant strips into useful data takes place inside the sensor itself. As can be seen in the a close-up of the sensor –figure 2.3– a capacitive coupling transmits AC signal from the sensor to the readout electronics. In fact, as will be seen in chapter 3, there will be a large DC voltage across this capacitor. The larger the capacitance, the better the charge coupling between sensor and electronics. The readout capacitor competes with other capacitances, such as those between strips for the collected charge.

The metal strips on the sensor are directly connected to the SVX3 chip by means of wire bonds. In the chip's analog front end, the signal is amplified in a charge integrating pre-amplifier. As the name indicates, this is a type of amplifier that deals with charge rather than with current or voltage. It is designed as an operational amplifier with a $0.2\,pF$ capacitive negative feedback coupling. If the operational amplifier is considered ideal, its inputs will remain at a fixed voltage and no charge will enter there. Therefore, all charge will be absorbed by the capacitor, and all of the resulting voltage drop, $U = \frac{Q}{C}$, will appear on the output side. The gain of the preamplifier is then found to be be $5\,mV/fC$.

2.3.3 Deadtimeless Analog Pipeline

For each channel in the chip, there is a *pipeline*, 46 analog memory positions, in each of which a measurement can be stored, while the CDF trigger systems decide if that measurement should be used or not. The pipeline is implemented as 46 capacitors per channel.

The preamplifier output is connected to a special *write* voltage amplifier, the output of which can be switched between the capacitors in the pipeline. Another *read* amplifier can independently have its input connected to any other of the capacitors. By using two amplifiers, the pipeline has been made *deadtimeless*, reading and writing can take place at the same time.

The write amplifier continuously loops over the capacitors in the pipeline. When it is decided that a previous measurement is of interest the chip is issued a trigger signal, known as *level 1 accept*. It will then look a certain number of positions back in the pipeline (set by the programmable number *pipeline delay*), remove that capacitor from the write loop and start reading it out. Simultaneously new measurements are still being stored in the pipeline.

2.3.4 The A/D converter and Dynamic Pedestal Subtraction

The back end of the SVX3D chip includes an analog/digital converter (ADC). It converts voltages into digital signal (ADC counts) by comparing them to a voltage ramp, and clocking the time when the signal- and ramp voltages are equal. By shifting the starting point of the clock, it is also able to subtract a pedestal level from all channels. For more details, see appendix B.

2.4 Noise

2.4.1 Types of Noise

For the mathematical basics on how noise is measured, see appendix A. There are a few different sources for noise in the detector system. One is thermal noise from resistors. This is caused by the random thermal motion of electrons in the resistor. Another type of noise is specific to semiconductors, arising from the statistics of charge carriers crossing the pn-junction in a silicon sensor. A third type of noise is externally generated noise, such as that from a time-varying external electric field.

A noise signal can enter the signal path of the system in three ways, depending on the location of the noise source. It can take the form of *charge* added to the measured charge, as *voltage* superposed onto the voltage of the readout capacitor or it can be created *inside* the amplifiers.

Capacitance hooked up to the preamplifier input is the main source of noise for SVX II. Capacitance converts voltage fluctuations into charge fluctuations. The more capacitance the larger the charge fluctuations and therefore the larger the noise. Noise is measured in terms of equivalent noise charge (ENC) in units of electrons.

2.4.2 The SVX3 chip and noise

Noise has been measured for the SVX3 chip in the idealized case when it is not connected to a sensor, but to a capacitor. Results are as follows, where noise is quoted as an Equivalent Noise Charge, the amount of charge on the sensor that would give rise to a signal of the same magnitude as the noise level. It has also been converted into "ADC counts" so that it can be compared to results obtained with the test readout module.

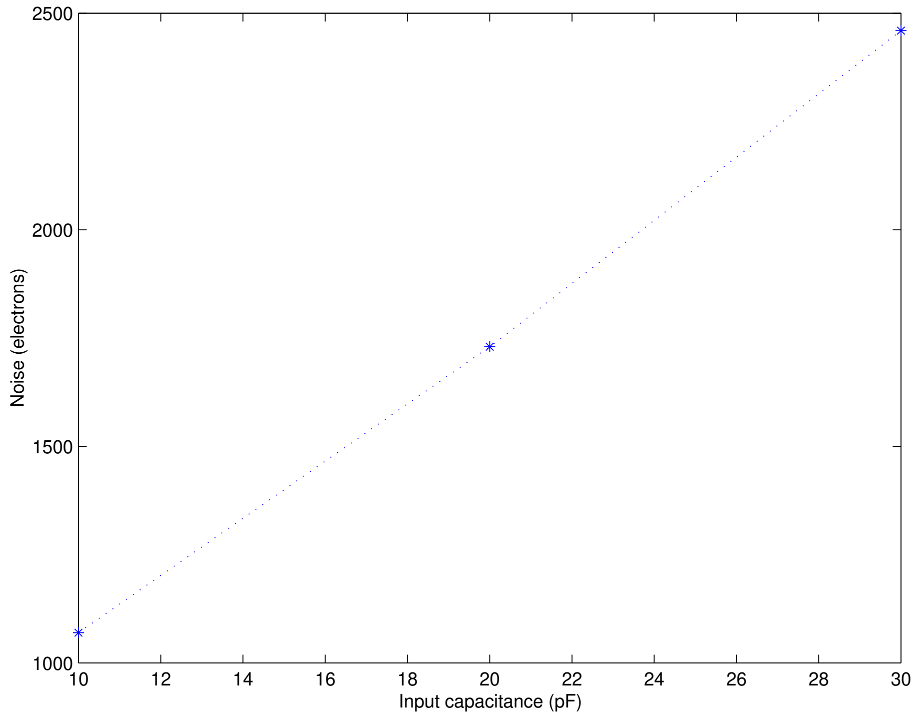

Figure 2.4: Illustration of the input capacitance–noise relationship for the SVX3 chip.

C_{in}	Bandwidth setting	Noise (ENC, eV)	Test module gain	Noise (counts)
10	4	1070	1/545	1.96
20	2	1730	1/545	3.17
30	1	2460	1/545	4.51
See figure 2.4 for a plot of this relationship.				

To adapt the preamplifier to input load the readout chip has a programmable bandwidth setting. In the table above this setting was chosen to achieve a constant bandwidth. The inverse of the bandwidth is the rise-time. For these measurements the 0 to 90% rise-time was 70 ns. That means that if you present a charge pulse at the preamplifier input, the output responds to 90% of full response in 70 ns. The integration time was 105 ns. As can be seen from the plot, this results in a linear relationship between noise and capacitance. The reason for this is that the dominant noise sources are those that generate noise in the form of voltage. The preamplifier on the other hand is sensitive to charge, and the connection between the two is the definition of capacitance, $Q = C \cdot V$.

2.5 The Test Readout Module

While doing this work, the author worked on a few different sensors –including ISL readout modules. But with one exception, all measurements presented in this report pertain to one single readout module, namely an SVX II Layer 0 prototype, with one sensor attached to it. As was mentioned before, in the finished detector two sensors will be bonded to one hybrid to form a readout module. Layer 0 is the innermost layer, so this is the smallest readout module, with 256 channels for each side of the sensor and a total of four SVX3 chips. Only three of those chips were working on the test readout module. Layer 0 has a 90° stereo angle, a 60 μm ϕ-side separation between strips (*pitch*) and a 141 μm z-side strip pitch. This module will be referred to as the *Test readout module* throughout this report.

It is a prototype in two ways. The first is just that some connections on the readout module are not done in the same way as they will be on the final version. The second is that it uses SVX3C chips instead of SVX3D. The difference is that the "C" version of the chip is lacking the Dynamic Pedestal Subtraction feature. The exception mentioned above is therefore everything in this text that refers to this feature. In this case a very similar readout module was used, but with SVX3D chips on it.

Measurements have been carried out using a $40\,MHz$ clock, and with a bandwidth setting of 4. This is the appropriate bandwidth setting for sensor capacitance in question, as defined in the table above. With these settings one ADC count corresponds to 545 electrons (the gain is $\frac{1}{545} ADC\,counts/e^{-}$).

Chapter 3

Biasing

3.1 Basics

A silicon sensor only works when biased. If the silicon is not depleted of charge carriers, the additional electron-hole pairs created by a passing particle would not be noticeable. And without an electric field to collect the charges, nothing would be read out anyway.

A lower limit on the bias voltage is therefore the voltage needed to fully deplete the sensor. In the SVX II design there is also an upper limit on the bias voltage, set by the readout capacitors (see figure 2.3). If the voltage drop across these capacitors exceed approximately 100 volts,[1] they will break down and be destroyed.

3.2 One Sided vs. Two Sided Sensors

In a one-sided 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.

3.3 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 readout capacitors of the first side. The highest available bias voltage will then be set by the voltage limitation of those capacitors alone. That is an unnecessarily low limit. Instead, the 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 $2\,M\Omega$ 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 readout 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.

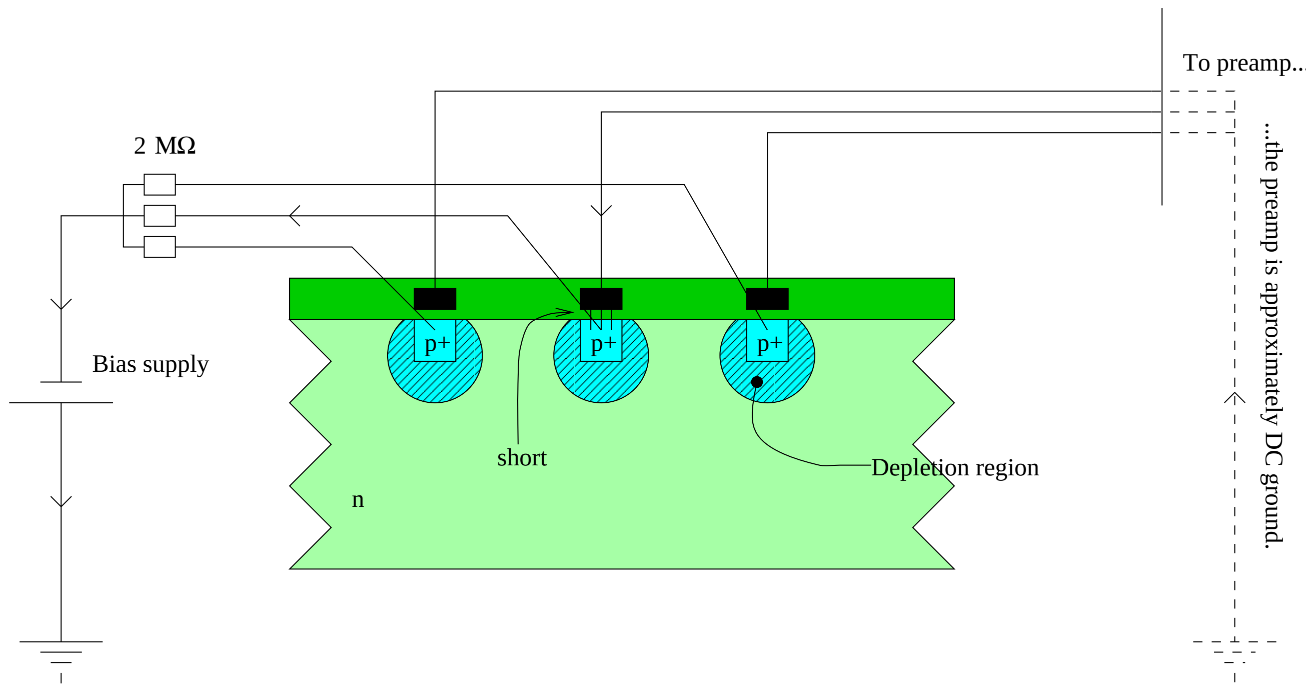

Figure 3.1: A pinhole on the junction side (currently p-side) of a sensor. The current flowing through the short is limited by the 2 $M\Omega$ resistor.

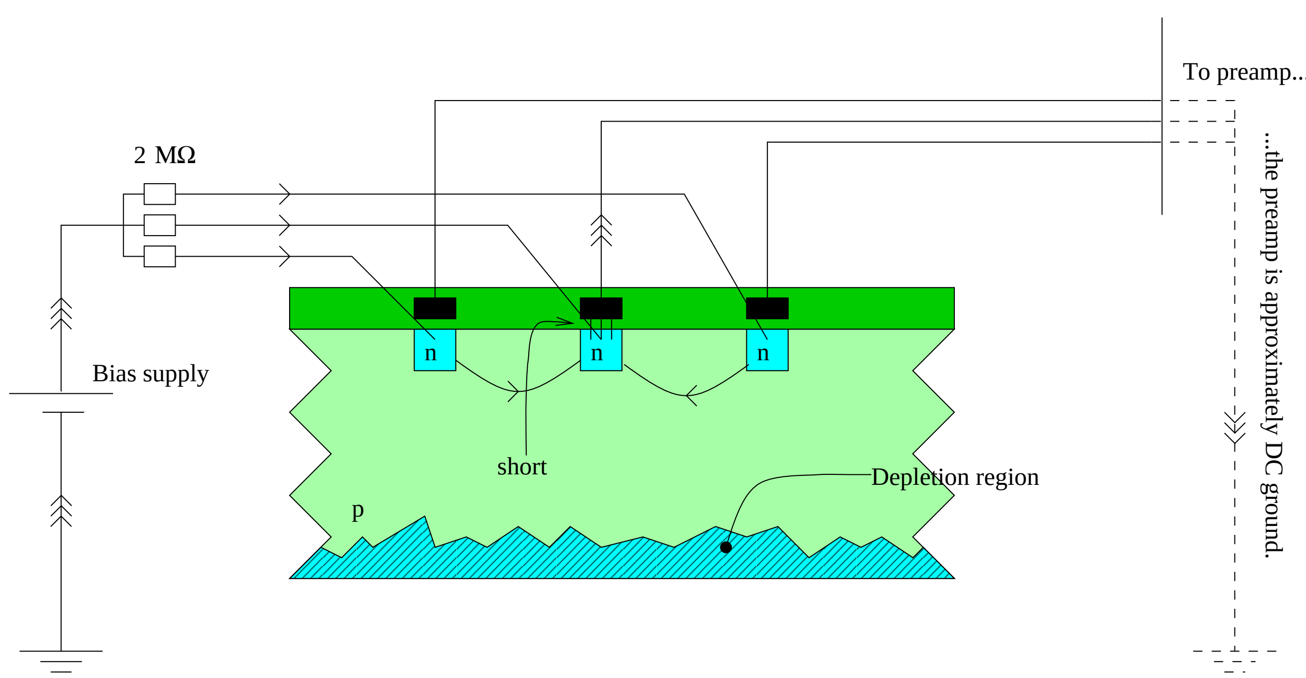

Figure 3.2: On the ohmic side of an undepleted sensor, strips are shorted together, illustrated by arrows through the silicon bulk. Thus a much larger current flows, as the 2 $M\Omega$ resistors are effectively connected in parallel.

3.4 Pinholes

A pinhole is a short through a readout capacitor. It can be caused 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. For the bias voltage, a pinhole therefore looks like a connection to ground, as illustrated in figure 3.1. To remedy the situation, the wire-bonds connecting strips suffering from pinholes to the SVX3 chip are removed, *pulled*.

Microdischarges

Apart from pinholes, sensors may suffer from another defect, microdischarges. A microdischarge is a localized discharge caused by a non-destructive electrical breakdown of the silicon. They occur where there is a strong electric field. Microdischarges will not only increase the bias current, but because of their stochastic occurrence, they will increase the noise levels.

3.5 Effects of Radiation

A layer in the SVX II detector will be subjected to a radiation dose on the order of $1\,MRad$ during its lifespan. See figure 3.3. The radiation will have several effects as far as biasing is concerned. Firstly, it will reduce the breakdown voltage of the readout capacitors, as was noted in the footnote on page 20. 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 an ohmic side, only they have reversed roles, the p-side now being the ohmic side, the n-side the junction side.

3.6 Bias current Measurements

3.6.1 With Pinholes

When pinholes are present they will result in very high currents, as is illustrated in figure 3.4. 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 3.2.

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 3.4. 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 3.2. 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.[3] Figure 3.5 illustrates this effect.

[1] 100 volts is the limitation for an irradiated sensor. A non-irradiated sensor has a limit of 120 volts.

[2] Both bias supplies are power limitated in this case, allowing a current of about $600\,\mu A$

[3] One particular combination of p-side- and n-side voltage will always give the same current. In this sense the detector does not

Figure 3.3: The projected radiation doses of the different layers in the SVX II. Layer 0 will be replaced after two years. Layer 1 and both layers 0 will undergo type inversion. From [**?**].

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. Compare to the workings of a field-effect transistor shown in figure 3.6.

3.6.2 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 3.7. The dependence of current on from which side the voltage was applied disappeared, as did the difference between incoming and outgoing current. Thus figure 3.7 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 high currents may be due to imperfections (lattice dislocations, cracks etc.) or impurities in in the silicon, possibly involving microdischarges. The actual mechanism has not been verified.

3.6.3 Conclusion

When pinholes are present on the ohmic side, special techniques are needed to deplete the sensor. This involves to 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, as it was motivated in section 3.3.

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, 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. See section 2.2.2 for details. But

possess a memory. But the order in which you do things does matter for which voltage combinations are achievable.

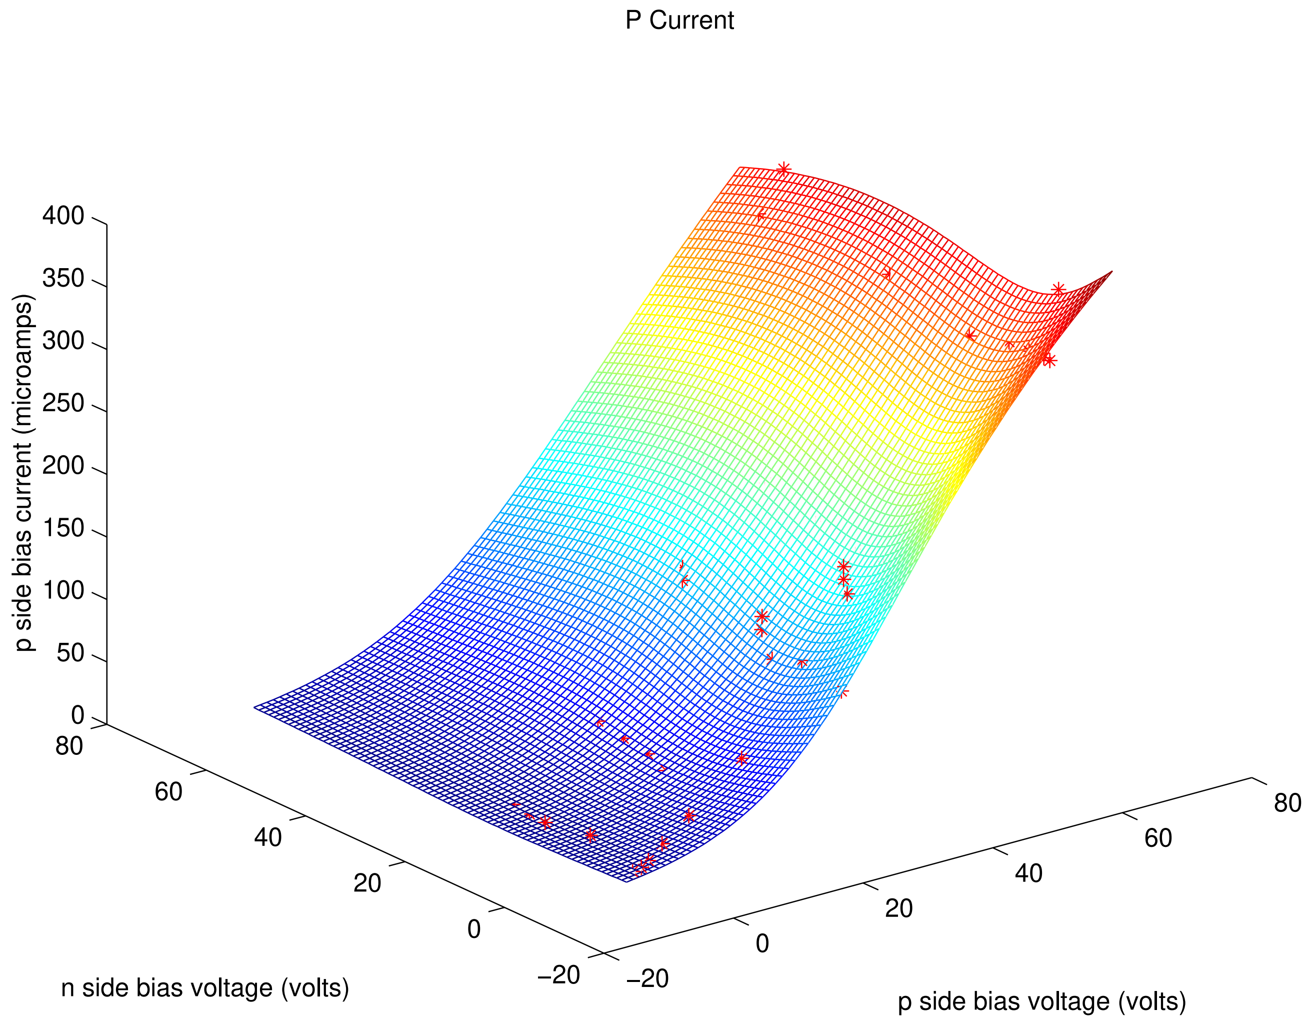

Figure 3.4: The current entering the p-side bias connection, when there was one pinhole on the p-side and one on the n-side. Most of this current is drained through the pinholes, meaning the n-side voltage has little effect on the p-side current. The stars (some are partly hidden below the mesh) are measured points. The mesh is an interpolation using a two-layer perceptron method.

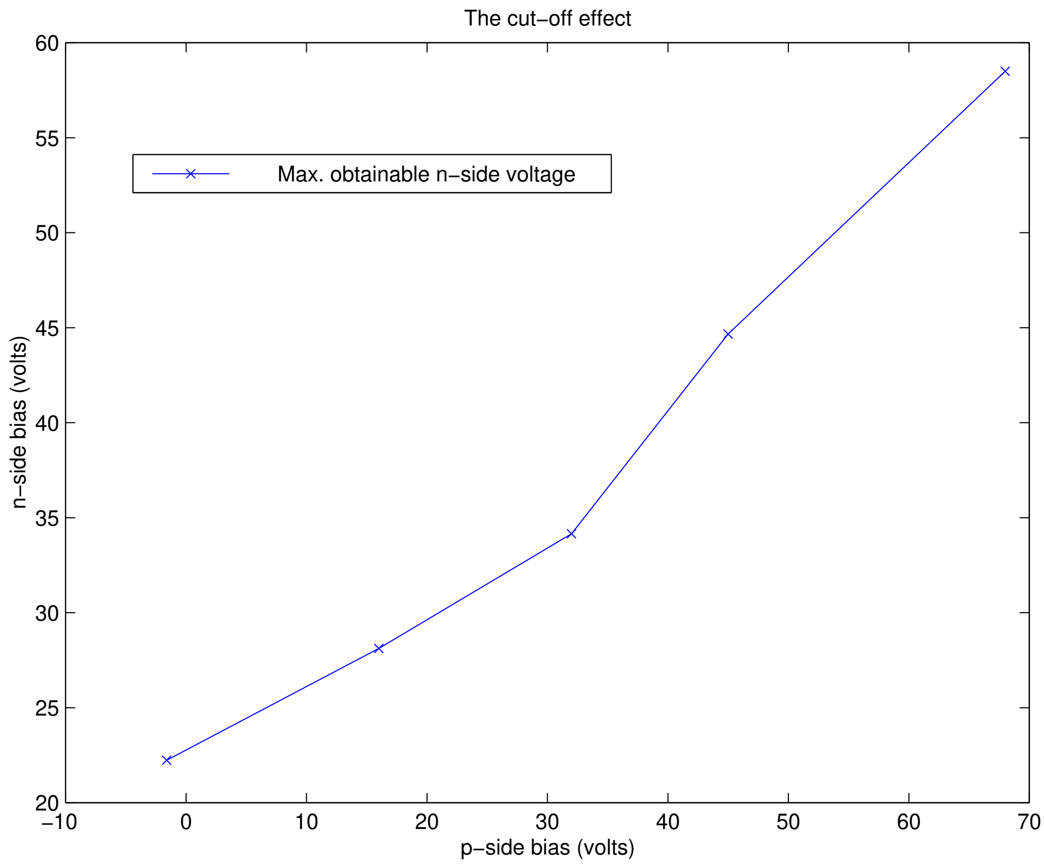

Figure 3.5: Cutting off the current through the n-side pinhole. The more the sensor is depleted from the p-side, the higher n-side bias can be applied before reaching the current limit of $600\,\mu A$. There is one pinhole on each side of the sensor.

Figure 3.6: A field effect transistor. As the depletion region grows into the conductive channel, the current from drain to source is cut off.

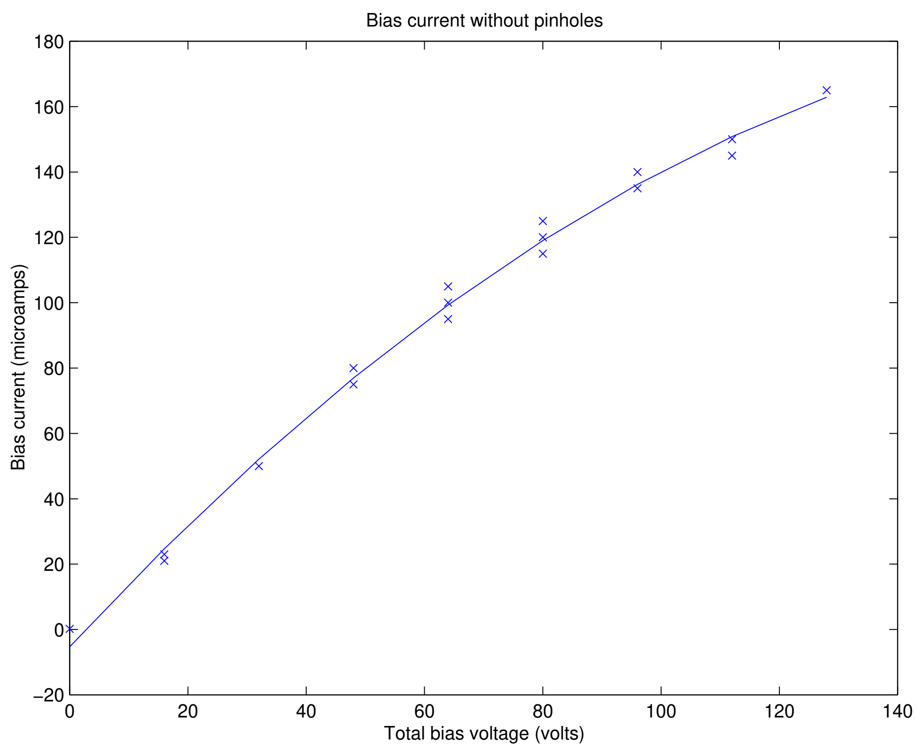

Figure 3.7: Bias current as a function of total bias voltage (p-voltage + n-voltage). The ingoing current was found equal to the outgoing up to instrument precision ($10\,\mu A$). The crosses are measured points. When several crosses occur at one total bias reading, they originate from different combinations of p-voltage and n-voltage.

at the time when biasing is most likely to be a problem, late in the life span of the detector, type inversion will have occurred. Then it will be the more important ϕ side that is lost.

3.7 Bias and Noise

The p-side noise is affected in a dramatic way when a small bias voltage is applied. In figure 3.8 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 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. This again illustrates the point that capacitance equals noise in the presence of a voltage noise source.

The levels of all noise types seem to depend only on total bias. Figure 3.9 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.

3.8 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 injected into the sensor. One way to do this is to shine laser light onto the sensor surface. The laser, with a wavelength of 890 to 910 nm 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 extensively tested, and is known to work well with the laser. Here is therefore reported 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 3.10. 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: $\diagup$) 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 3.11 shows the results.

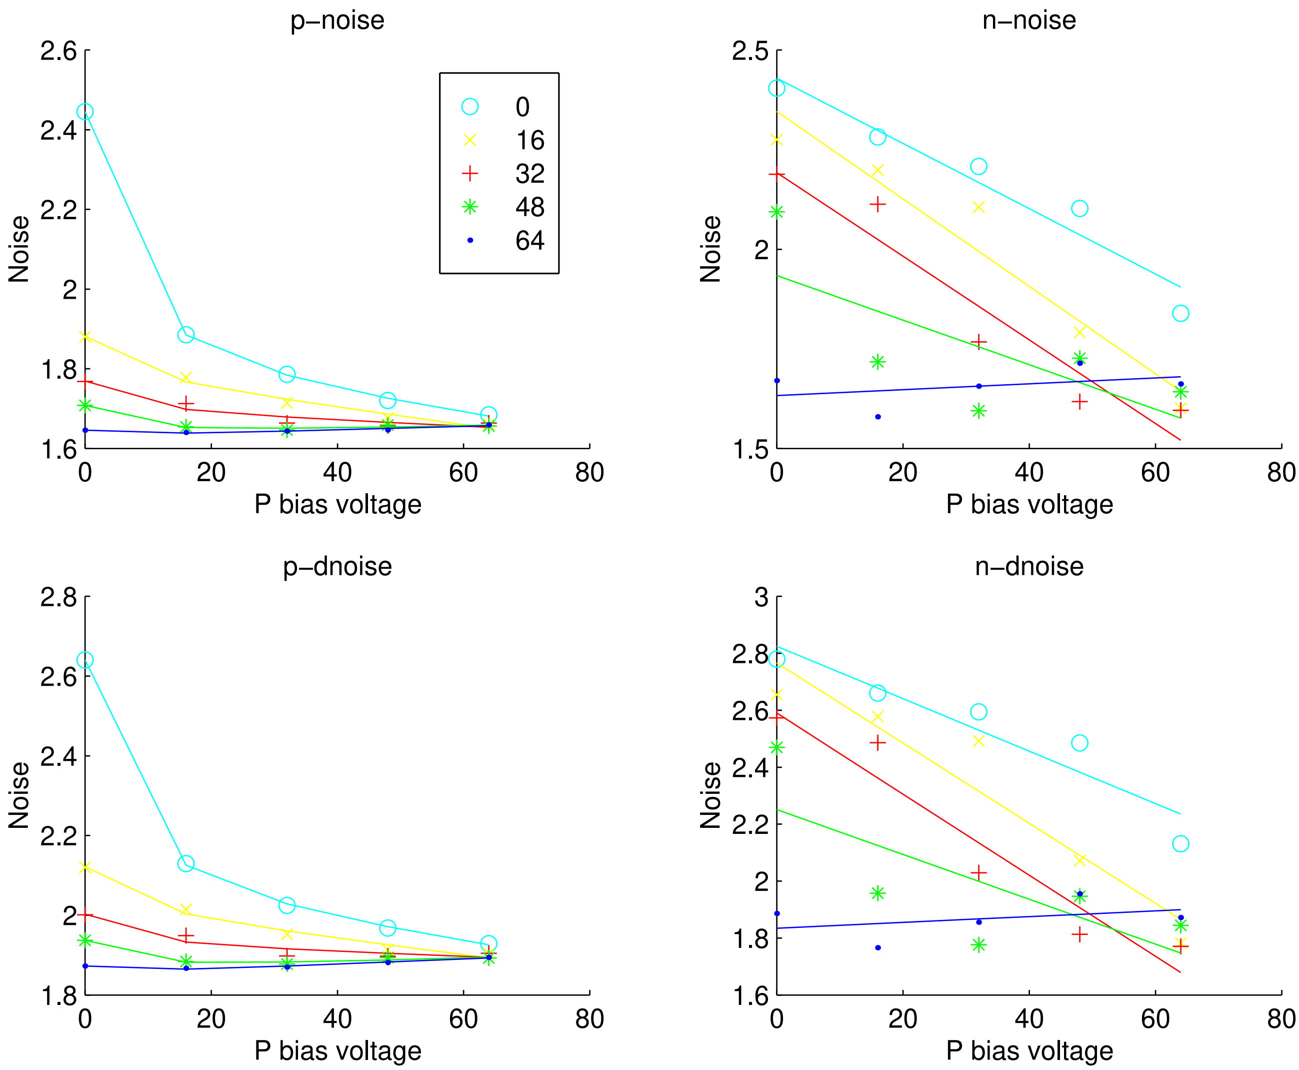

Figure 3.8: Four different types of noise in the readout module. See appendix A for an explanation of their definition. Each diagram shows five sets of data, each corresponding to a different ϕ-side bias voltage, as indicated in the legend box. The dots, crosses etc. are data points, the curves are fitted 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 3.7.

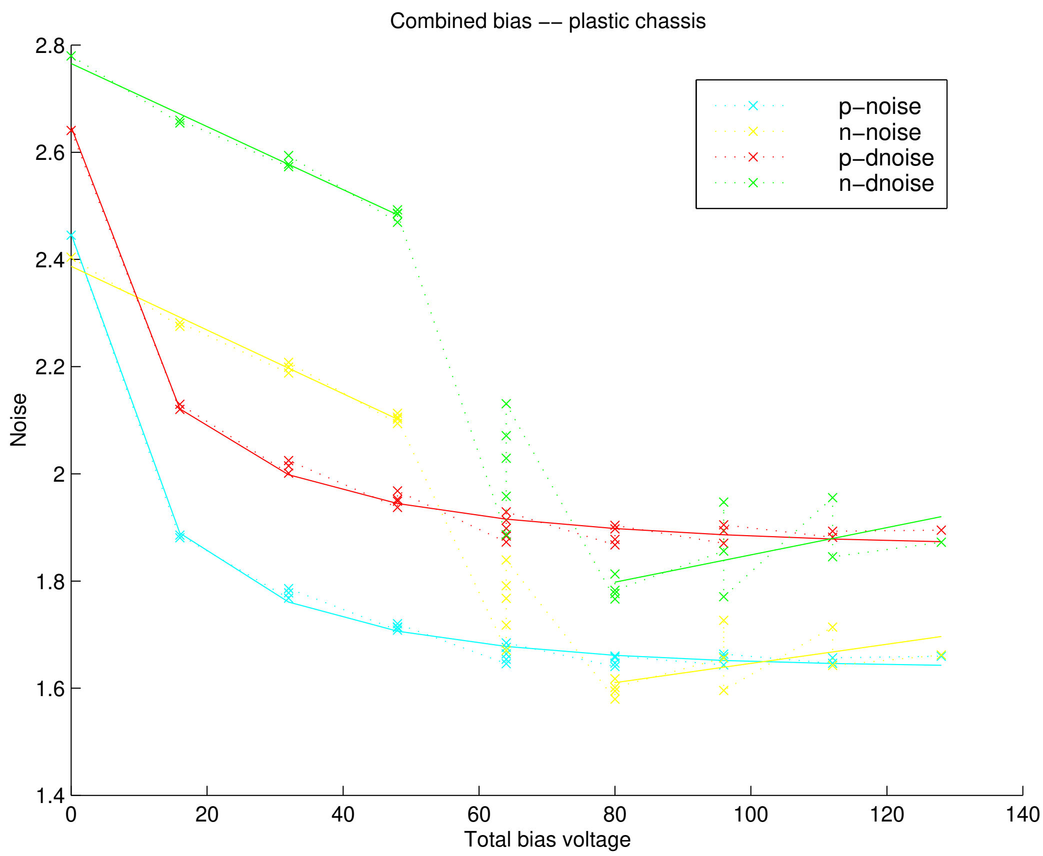

Figure 3.9: Noise as a function of total bias voltage. Measured data points, fitted curves.

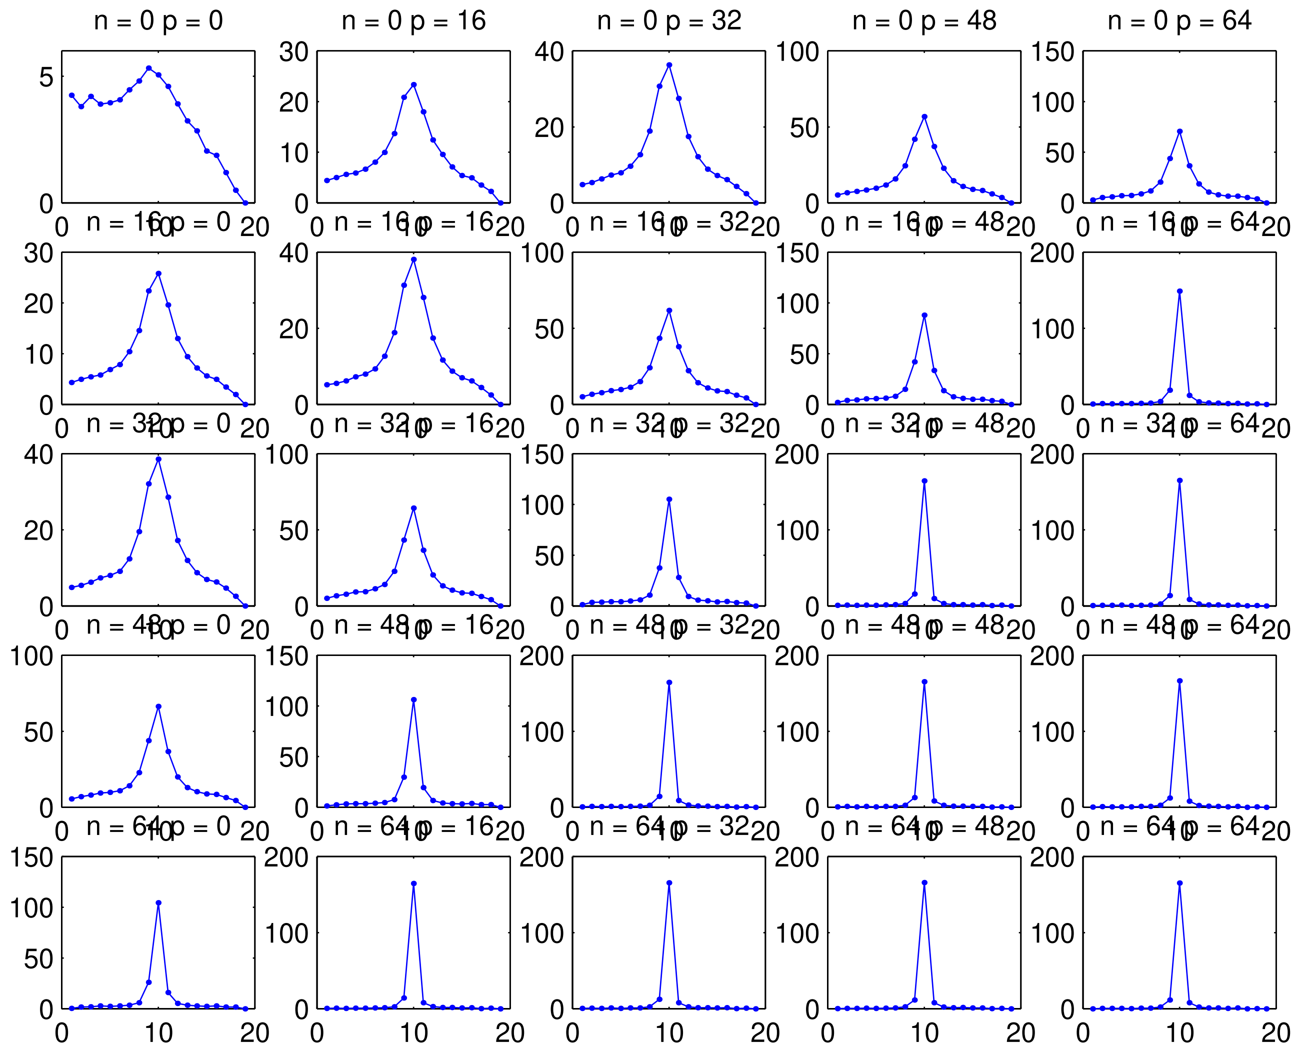

Figure 3.10: A laser was applied to the p-side of the sensor. Shown here are the signal distributions in the n-side strips near where the laser was pointed, but on the opposite side of the sensor. The x-axis in each plot thus represents relative channel number, with the pulse centered on channel 10. 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 3.8.

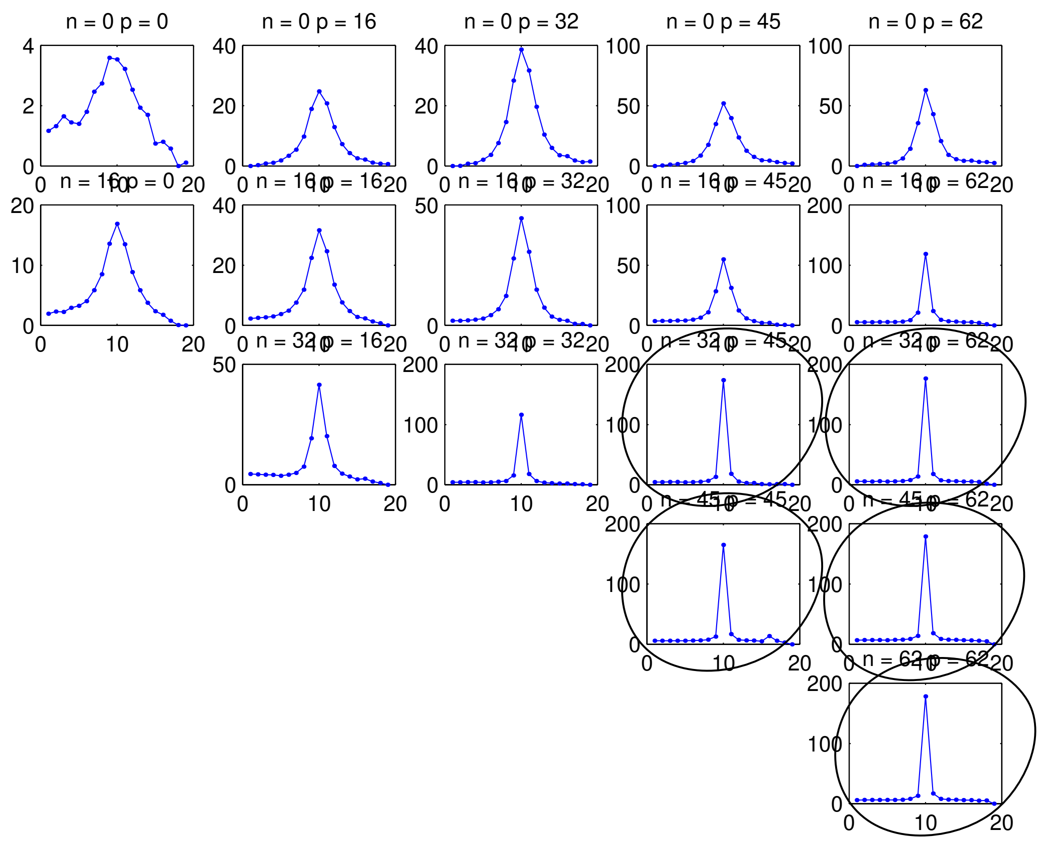

Figure 3.11: Test of sensor depletion with pinholes. The sensor is depleted when the laser pulse results in a narrow peak, as it does for the circled cases. Each column of plots represents a certain p-side voltage, each row a certain n-side bias. See figure 3.10 and section 3.8 for a further discussion. 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.

Chapter 4

Grounding

4.1 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 bias lines connecting to the sensor, the data bus and activity inside the chips themselves and signals in the barrels where the modules are mounted.

4.2 Results

The test readout module, mounted in a metal chassis, turned out to suffer from a large noise problem. Visual 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 4.3, 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.

4.2.1 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 4.4 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 4.4. 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.

The strips have the same width on both sides of the sensor, and the separation between sensor and fixture are also roughly the same (within a factor 2). Therefore, the ratio of capacitances will just be determined by the ratio of the length of the overlap. For one channel on the p-side the length in question 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 ($h = 15.3\,mm$). From this we would expect the noise ratio to be 1.5:1, with the n-side having the larger noise.

To find the actual capacitances, we also need to know the width of the metal strips, and the separation between strip and fixture. The strip width, t, is $25\,\mu m$ on both sides of the sensor. On the p-side there is a tape strip insulating the sensor surface from the fixture. On the n-side there is an insulating coating serving the same function. The tape, specified to be $25\,\mu m$ including glue, was measured to be $4\,\mu m$ thick with a micrometer. The coating could not be measured. In practice the thickness of tape and coating are not the most important factors in determining the separation of sensor and fixture. A rule of thumb says that uncoated surfaces will be separated by about $25\,\mu m$ when pressed together. Within a factor of 2 both dimensions can therefore be taken to be $25\,\mu m$. The usual formula for a parallel plate capacitor requires that the width of the plates is much larger than the separation of the plates, and can therefore not be used in this case. The electrostatic problem to be solved in order to find the capacitance is illustrated in figure 4.1. It is a two-dimensional problem since fringe effects at the ends of the strips contribute very little to the capacitance. There are several ways to solve this kind of problem. The following approximate formula was derived by Bakaglou. It is a good approximation when the width (W) of the strip is not large compared to the separation (D). In our situation, when the two are equal, it is has been shown by Campbell [?] to be accurate to within five percent, which is perfectly satisfactory.

$$C = l \cdot \frac{2\pi\epsilon}{ln\left(\frac{8D}{W} + \frac{W}{4D}\right)}$$

With $W = D$, this gives $C = l \cdot \epsilon \cdot 2.98$. Taking the dielectric constant to equal that of free space, $\epsilon = \epsilon_0 = 8.854 \cdot 10^{-12}\,C^2/(Nm^2)$ and inserting the lengths the capacitances turn out to be $C_n = 400\,fF$, $C_p = 280\,fF$. The absolute uncertainty for those values are a factor of two. The uncertainty of their ratio is also a factor of two. These very approximate values can be compared with other capacitances that occur in the detector to estimate noise from other sources.

There is an additional uncertainty in this result. Figure 4.1 shows just one strip coupled to a ground plane. In reality, there is an array of parallel strips. The neighboring strips will shield off parts of the ground plane, squeezing the field lines of the strip in consideration down to its "fair share" of the ground plane only. The situation then approaches that of an ideal parallel plate capacitor, with straight field lines. This makes capacitance more critically dependent on the separation of the plates, and thus increases uncertainty.

If the two sides of the test readout module were equivalent with respect to this noise source in every way except for the different capacitances, the pickup noise ratio would fall in the range 0.75:1—3:1. To calculate the noise ratio, noise from other sources has to be subtracted from the curves plotted in figure 4.3. From figure 3.8 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 4.3 should be the sum in quadrature of the two noise types, $n^2 = n_{fixture}^2 + n_{other}^2$. Extracting $n_{fixture}$ produces figure 4.5. From that figure, the noise ratio is found to be 2.6:1. Within the accuracy the result neither proves nor disproves the hypothesis that the ratio of capacitances corresponds to the ratio of noise levels. In other words, within a factor of two there may be other reasons for the relative noise levels.

4.2.2 Reducing the noise

There are two ways to get rid of the antenna effect. 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 4.2. As configuration number nine in figure 4.3 shows, 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 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 3.8 for the resulting noise level. In chapter 5 some more measurements of grounding schemes are presented.

4.3 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

Figure 4.1: A thin strip on top of a ground plane. (Adapted from Campbell.)

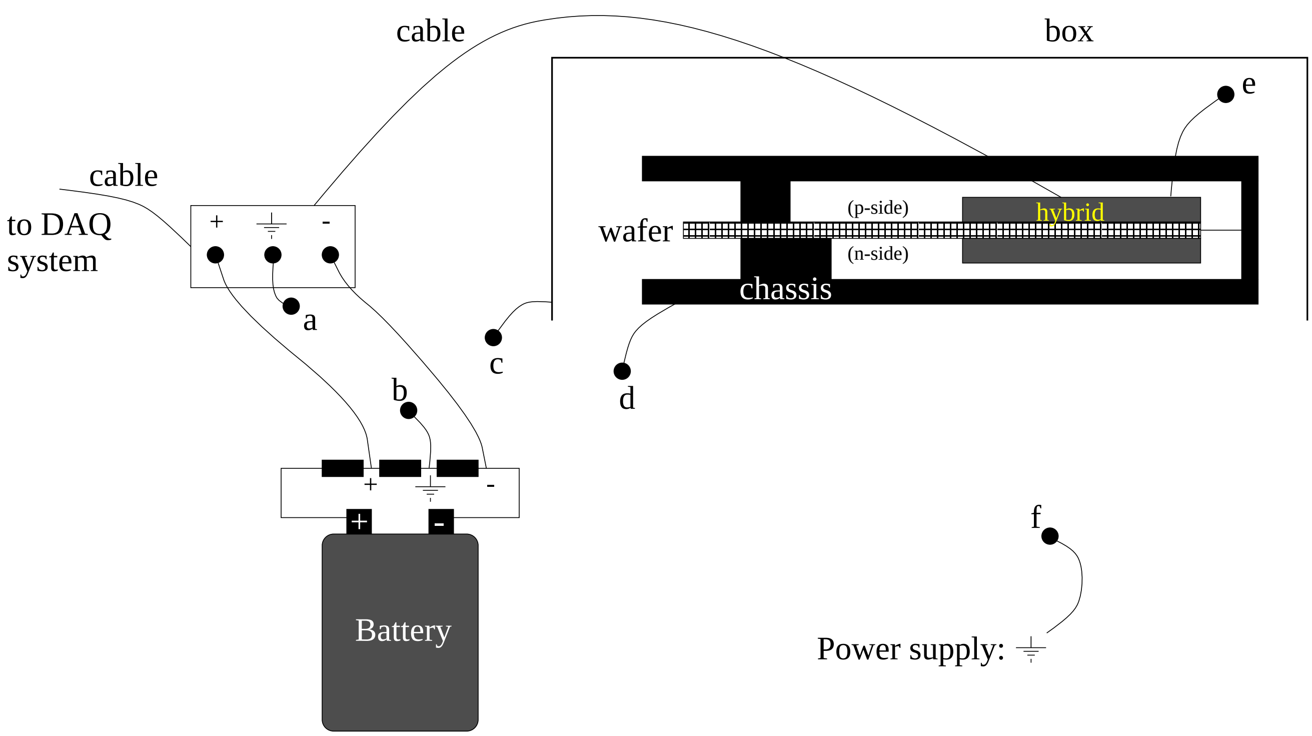

Figure 4.2: The test readout module mounted in a metal chassis. A voltage divider connected to a battery provides a positive and a negative bias voltage referring only to the midpoint of the voltage divider. The points **a**–**g** are all in some sense "ground". They can be connected in different combinations. The **box** can be put on top of the module like in the figure, or it can be removed.

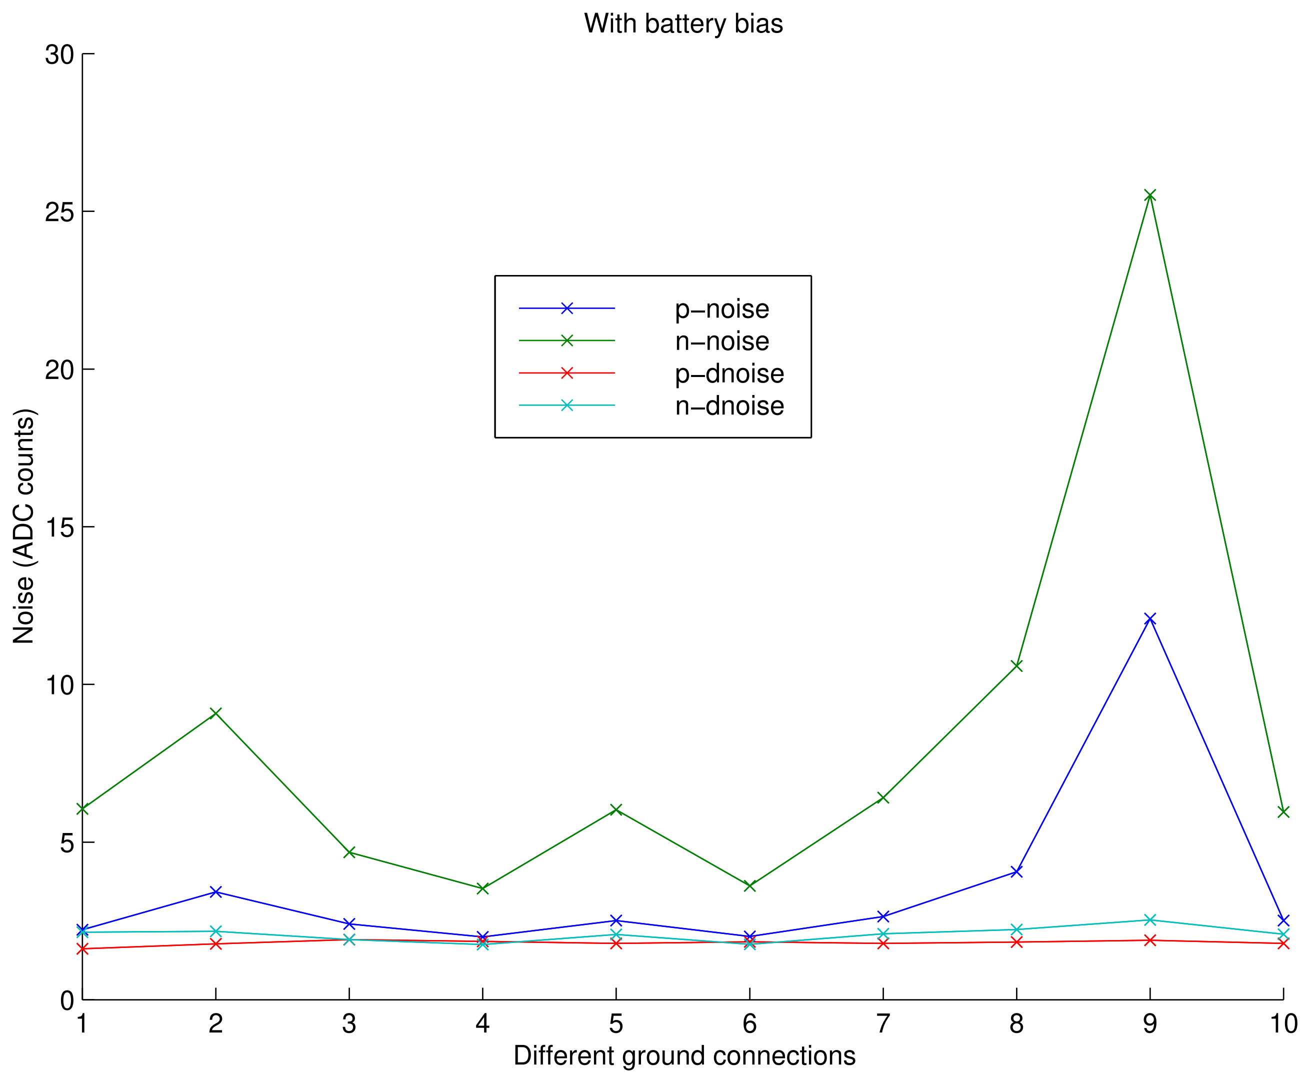

Figure 4.3: The effects of grounding and shielding on noise and differential noise. Refer to appendix A for definition of the noise types. Each number on the x-axis represents a different grounding configuration. Using figure 4.2 as a key, the below table lists how the grounding was done. Each dash in 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**	$\surd$

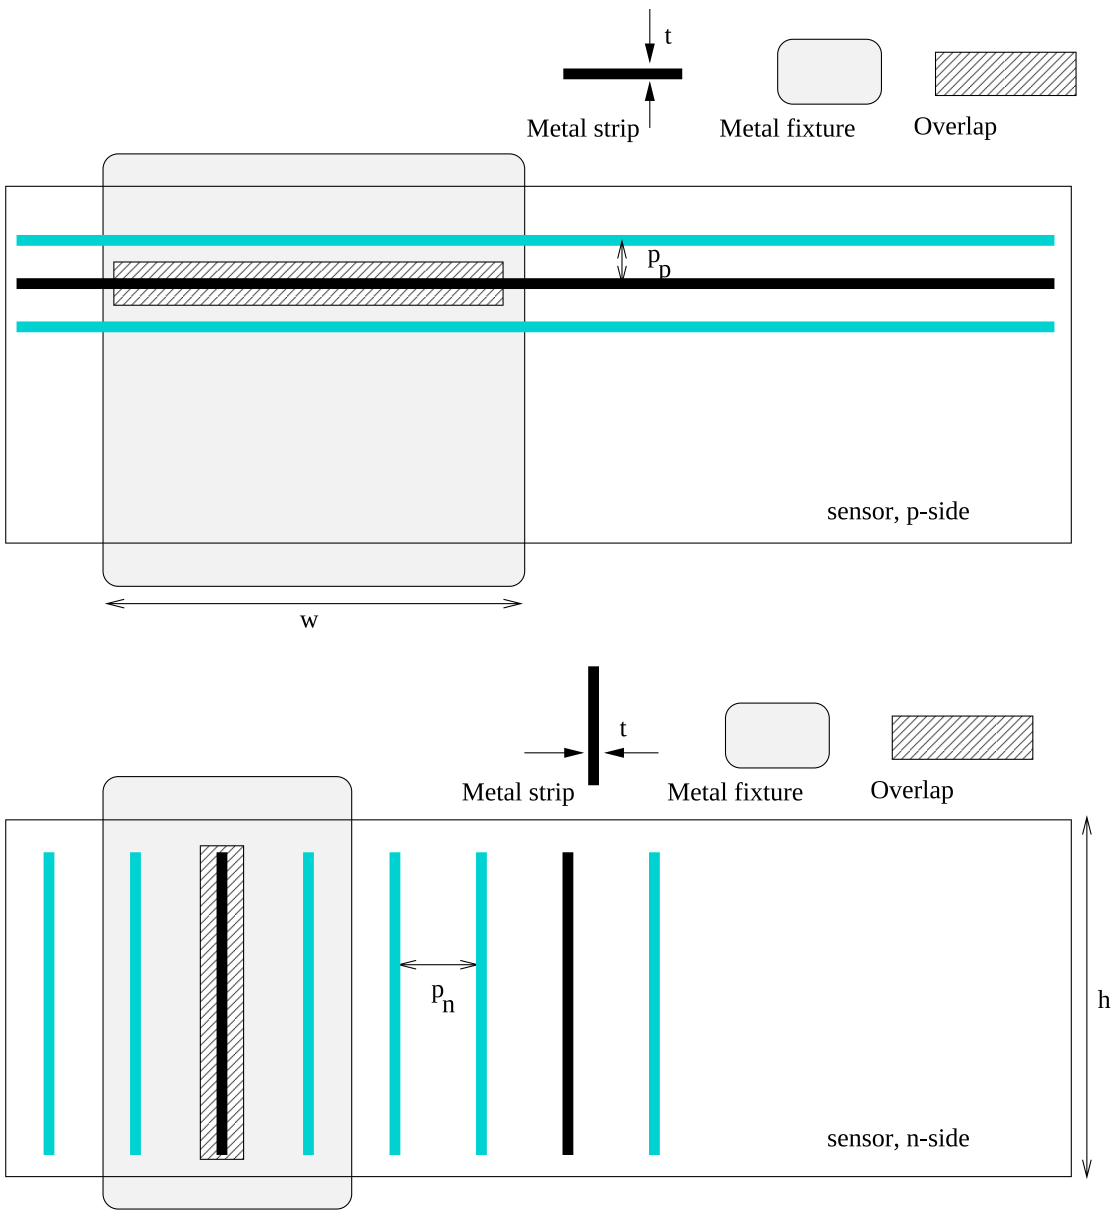

Figure 4.4: The sensor and the metal chassis. Does the coupling geometry explain the relative noise magnitudes? This figure defines the relevant overlap areas for the two sides of the sensor. See section 4.2.1 for the calculations.

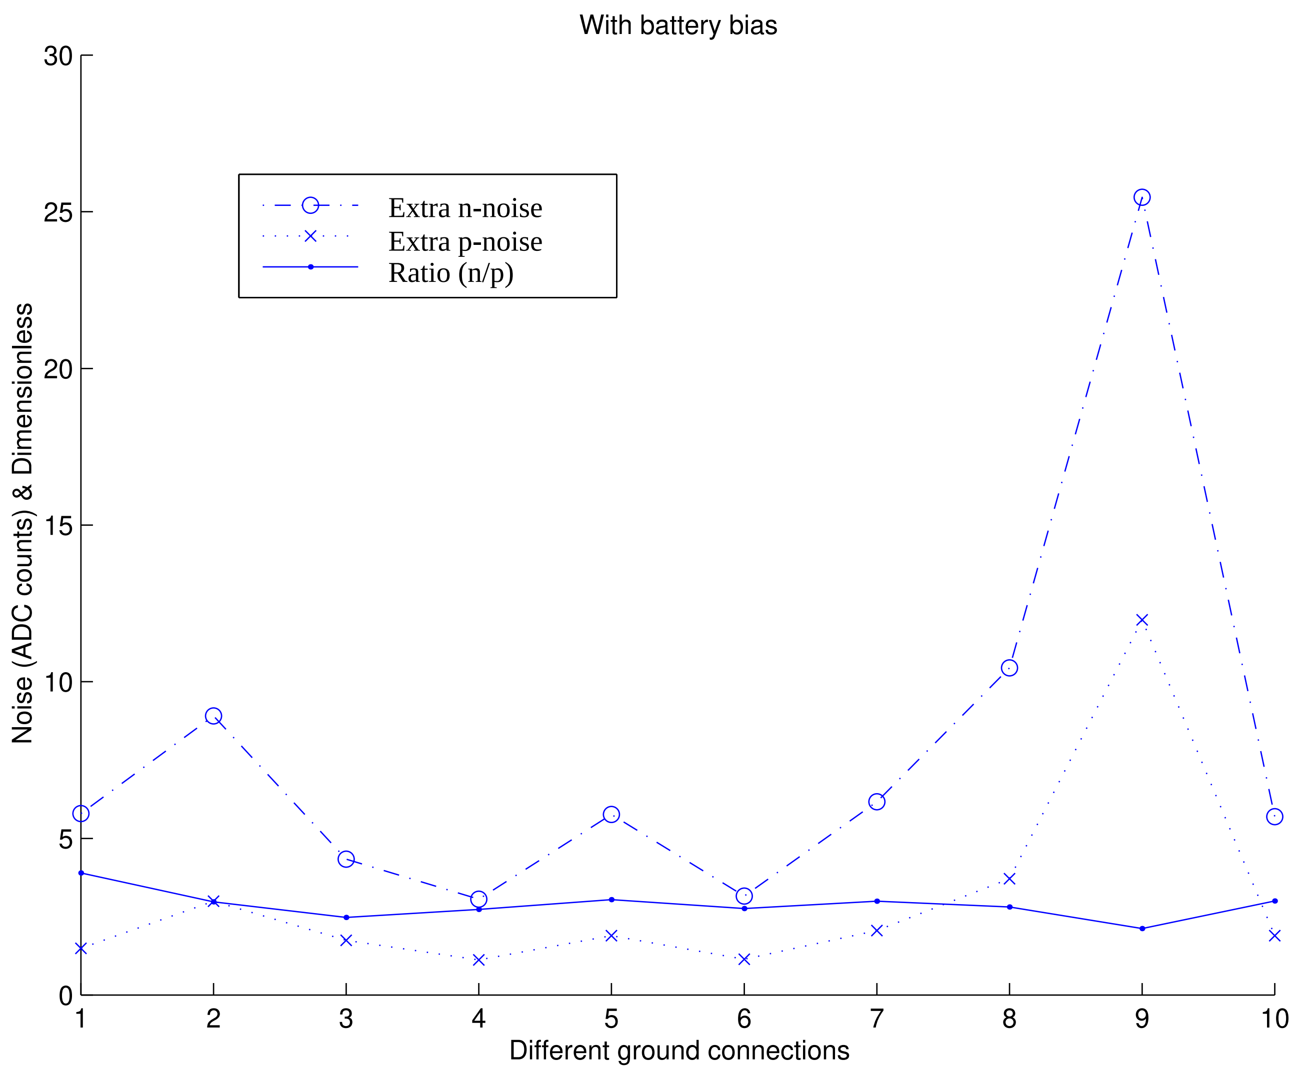

Figure 4.5: "Extra" noise originating from the metal chassis. Supposing an independent background noise of 1.75 for the n-side, 1.65 for the p-side, the extra noise is plotted for the n− and p-side. The ratio between the n-side and p-side additional noise is almost constant, with a mean of 2.6. This was the observation leading to the identification of the coupling mechanism for the noise.

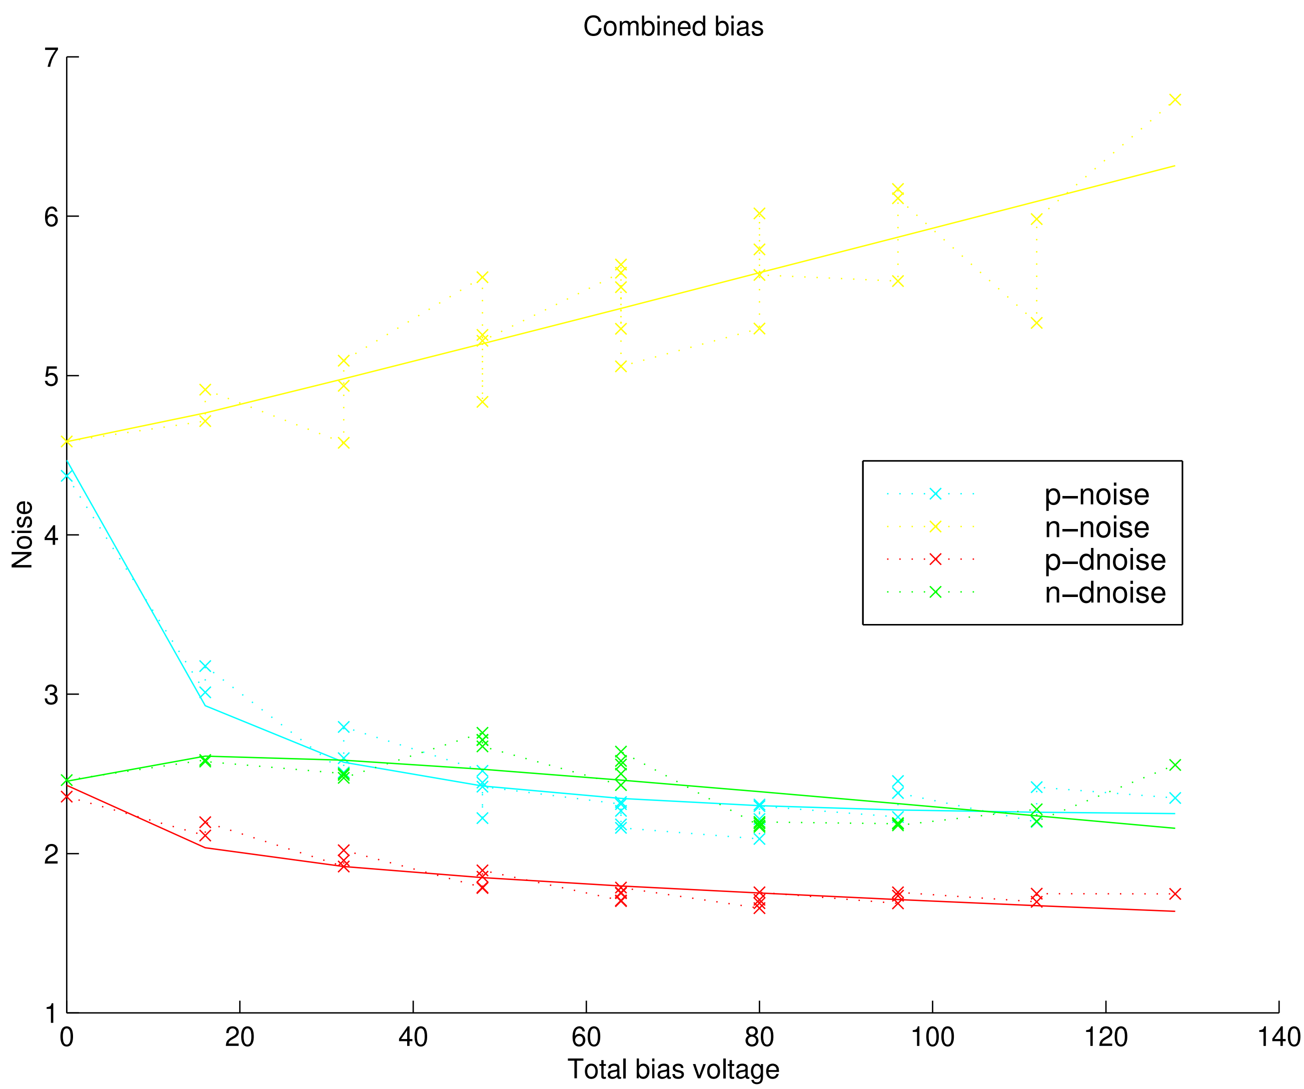

Figure 4.6: The four types of noise, as a function of total bias voltage, using the metal chassis.

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.

Using Bakaglou's formula for a strip capacitor, the factor which was found to be 2.98 for the $25\,\mu m$ separation to a metal chassis, becomes 0.65 for a separation $D = 5\,cm$. For a full length sensor on the p-side ($w = 21\,mm$) this means a capacitance of $120\,fF$. Without grounding to the bulkhead, and for the same electromagnetic noise level as in the test readout module, this would give a noise level of about half that appearing on the p-side in figure 4.5. This capacitance is also what the ground connection impedance should be compared to.

Chapter 5

Deadtimeless Operation

5.1 Additional Noise Sources

The measurements of readout module performance referred 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 or activity on the readout bus feeds back into the signal path.

5.2 Principle of Measurements

As was described in section 2.3.3, measurements are continuously being stored in the analog pipeline. With each tick of the front end clock, the pipeline stores a measurement in the current position in the pipeline 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 deadtimeless 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 chapters 3 and 4.

5.3 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 5.1. 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 5.1 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

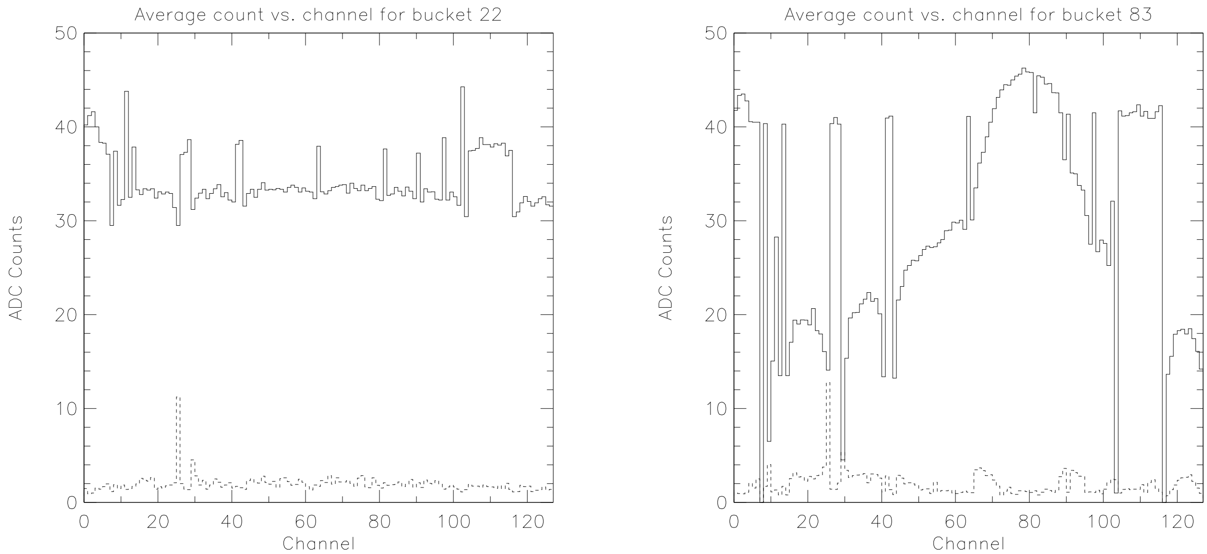

Bucket 22. A nice bucket. Bucket 83. A bucket with pickup.

Figure 5.1: 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.

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 from the digitization process in the SVX3 chip.

5.3.1 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 5.1.

Pedestal shift

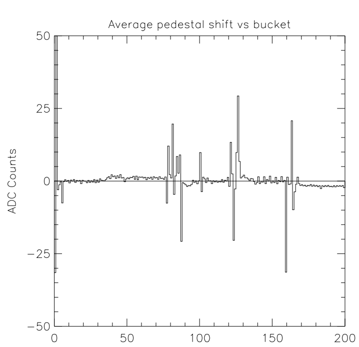

The first measure on data quality is *pedestal shift*. To find the pedestal shift for a bucket, the pedestal level average taken over all good channels in that bucket is calculated. From that number is subtracted the average pedestal level for all good channels in *all* buckets. The difference is the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

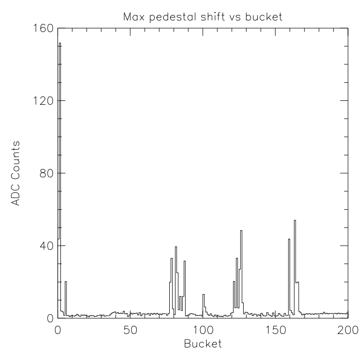

The next quality measure tells about deviation from the pure pedestal shift. *Maximum pedestal deviation* is defined as the greatest deviation from the pure pedestal shift for any one channel. To find maximum pedestal deviation for a bucket, the pedestal shift for each channel in the bucket, relative 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

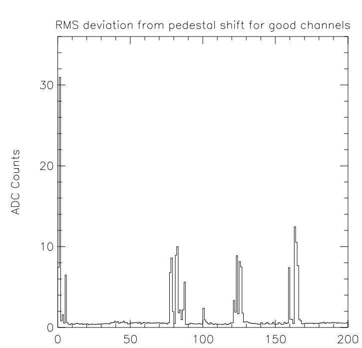

The *RMS pedestal deviation* is calculated in a way similar to the maximum pedestal deviation. Instead of taking the largest absolute shift, one calculates the RMS over the individual shifts of the channels in a bucket. Together with the maximum pedestal deviation the RMS pedestal deviation gives a statistical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

Maximum noise

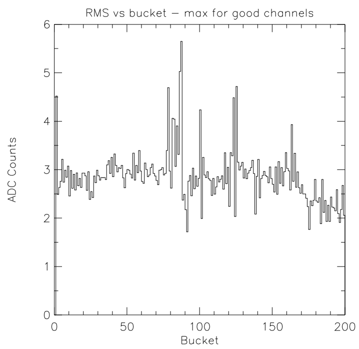

In contrast to the above measures which all deal with averages for a bucket, the noise level describes a *variation* in the pedestal level for a certain channel in a certain bucket. This is simply RMS noise in the same sense as it was used for the non-deadtimeless measurements. *Maximum noise* is the noise level of the noisiest channel in a bucket.

Average noise

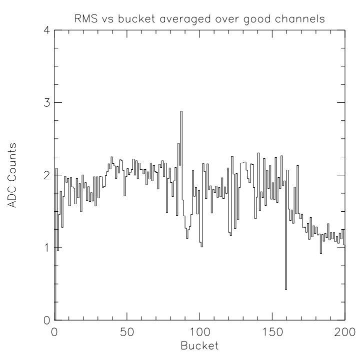

Average noise is the average RMS noise in the pedestal level for good channels. Together with the RMS pedestal deviation it gives an overall measure on noise; unwanted signal that cannot be subtracted away.

5.3.2 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 bucketwise 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 5.3–5.8 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 5.9 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 5.9 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.

5.3.3 Identification of noise sources

Figure 5.2 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 5.3.1. The dashed curve is the deviation from pedestal shift, as defined in section 5.3.1. In this panel time is measured in buckets, as explained in section 2.3.3. 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 5.2 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 me 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.

5.3.4 Biasing

Total bias voltage

Figures 5.3 and 5.4 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 5.2. 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 5.3.2. 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.

Bias splitting

The splitting of bias voltage has not been found to have any effect on the performance of the readout module. Figure 5.5 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.

5.4 Grounding and Shielding

5.4.1 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 4.2. 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.

5.4.2 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 chapter 4, 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 5.7. Figure 5.6 shows some combinations of grounding, shielding and bending of the data cable. In contrast to figure 4.3, 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)

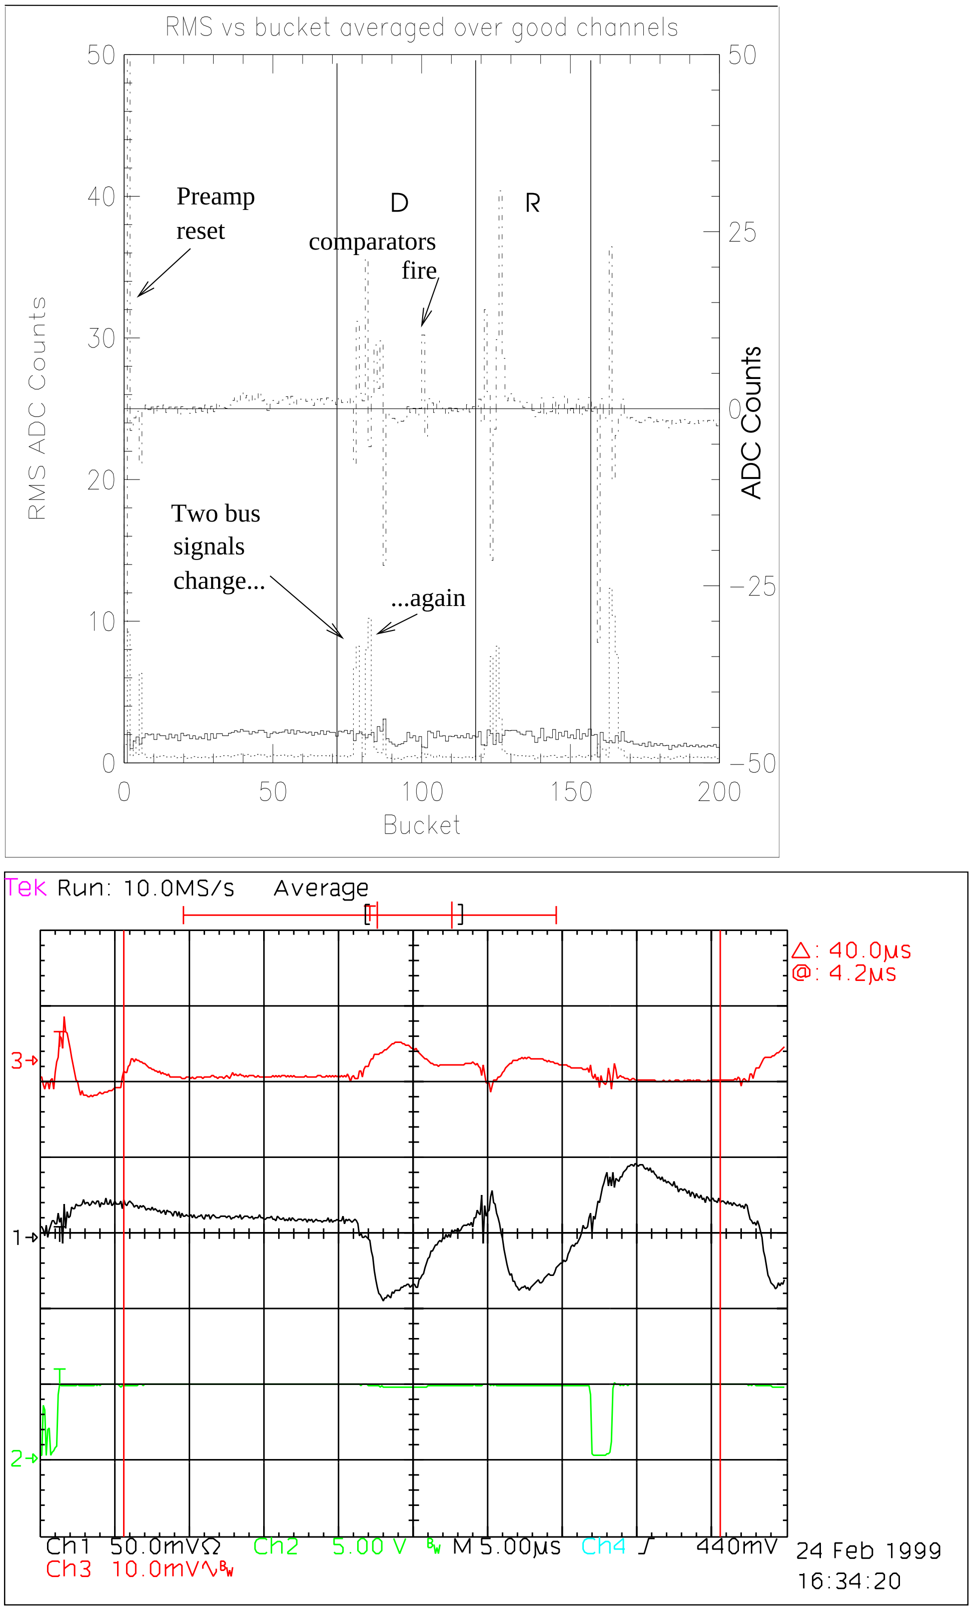

Figure 5.2: Overall performance in different modes. Top: Pedestal shift, pedestal deviation and noise. Bottom: Oscilloscope measurements of voltage supplies. The two panels are synchronized. The sensor is biased from the bias supply, using $48 + 48$ volts. A plastic chassis is used. For an explanation of what happens in this figure, see section 5.3.3.

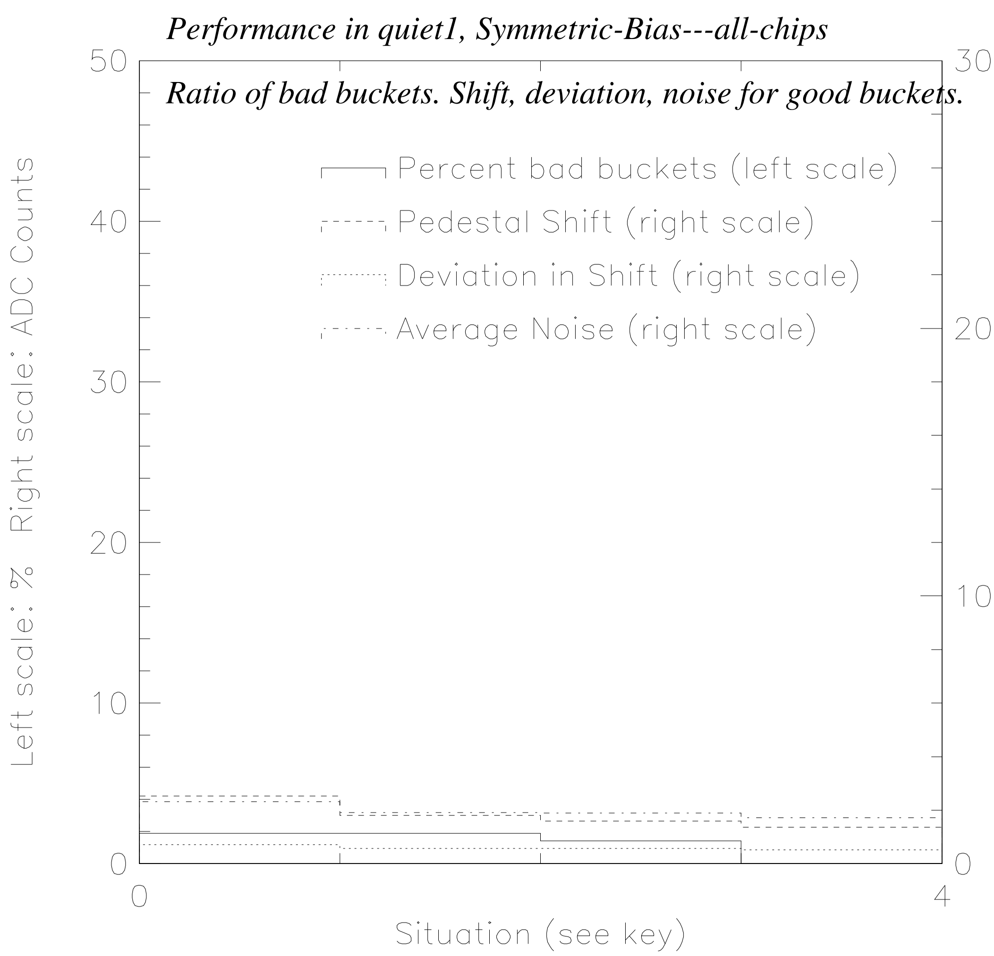

Figure 5.3: Performance of the test readout module in four bias situations. From left to right: Bias voltage $0 + 0$, $16 + 16$, $32 + 32$ and $48 + 48$ volts. The solid line is the percentage of rejected buckets. The other curves are the aggregate performance indicators. Note that the solid line refers 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 5.2. The first two buckets (preamplifier reset) are not included.

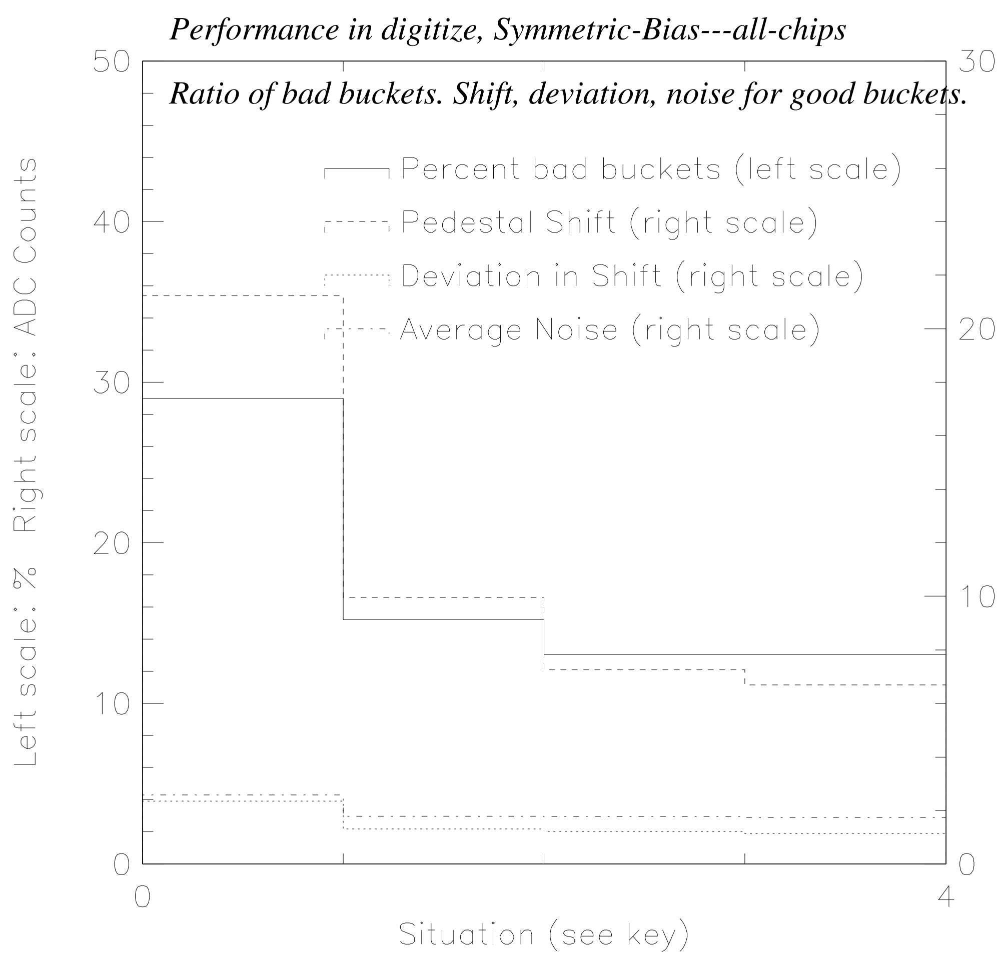

Figure 5.4: Same as figure 5.3, but for data taken while the chip was in digitize mode. Refer again to figure 5.2 for the definition of the modes. Note that performance is much worse compared to data taken with a quiet chip.

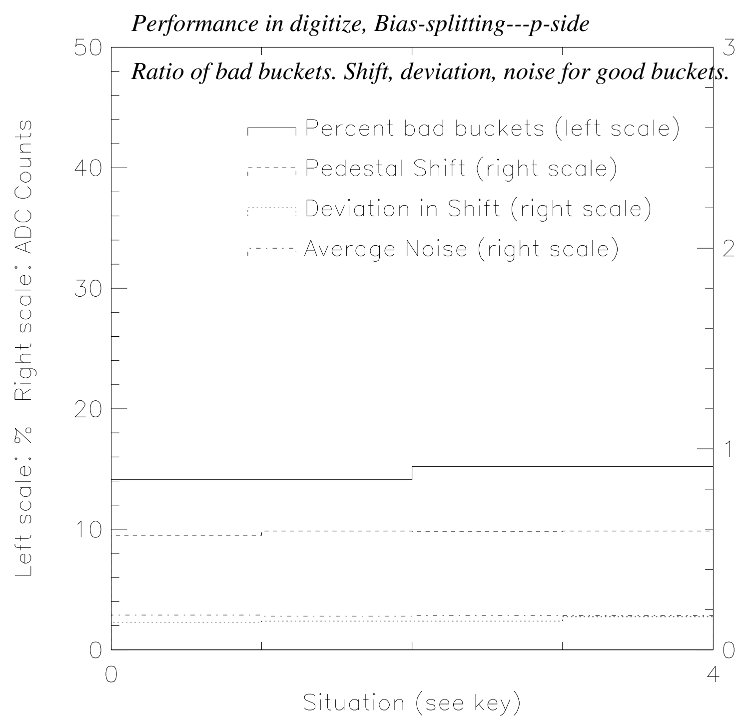

Figure 5.5: Performance in digitize mode. This time we are only looking on the ϕ-side. Unlike what was the case in the previous figures, total bias voltage is held constant at 80 volts. Instead it is the splitting of those 80 volts between the two sides of the sensor that is varied. For the four situations 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.

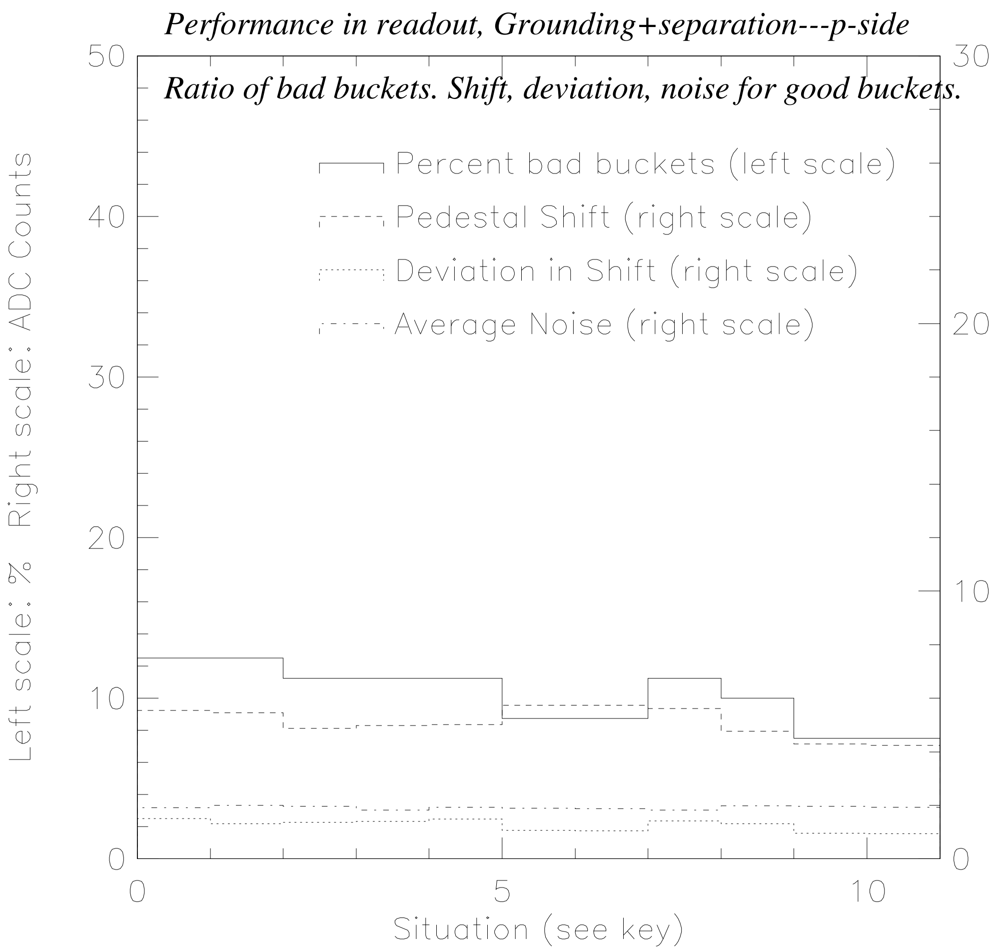

Figure 5.6: Chip performance for different ground configurations. Key to situations on the x axis:

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?	–	$\surd$	$\surd$	$\surd$	$\surd$	$\surd$	$\surd$	–	$\surd$	$\surd$	$\surd$
Shield grounded?	–	–	$\surd$	$\surd$	$\surd$	$\surd$	$\surd$	–	$\surd$	$\surd$	$\surd$
Bias ground to shield:	–	–	–	near	far	–	near	–	–	–	near
Cable close to silicon?	$\surd$	$\surd$	$\surd$	$\surd$	$\surd$	–	–	$\surd$	$\surd$	–	–

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.

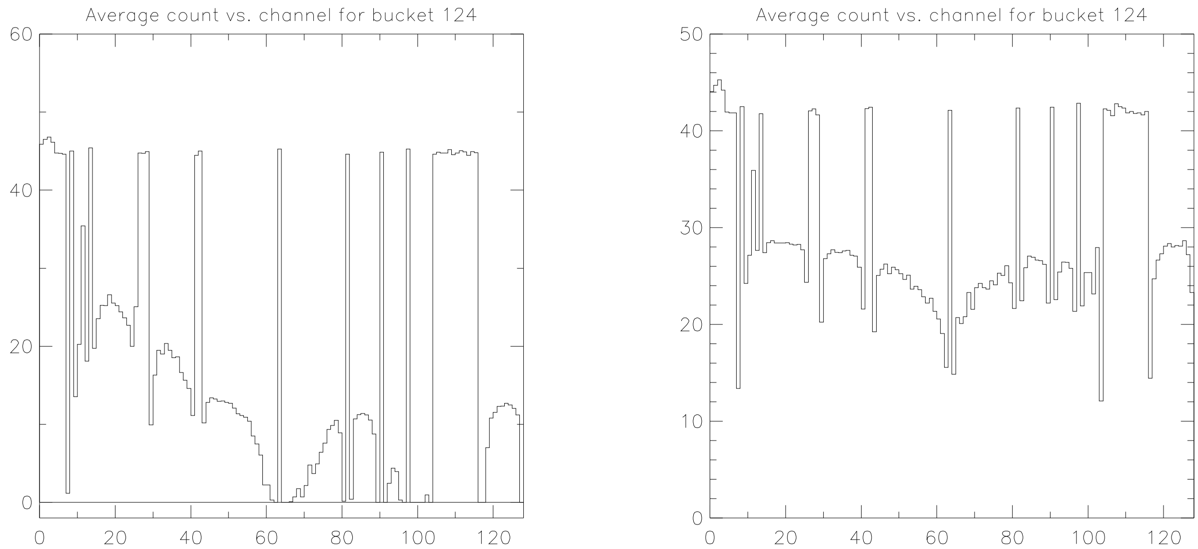

Wire on top of sensor

Wire away from sensor

Figure 5.7: 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.

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.

5.5 Capacitive coupling

In chapter 4 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 5.8 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.

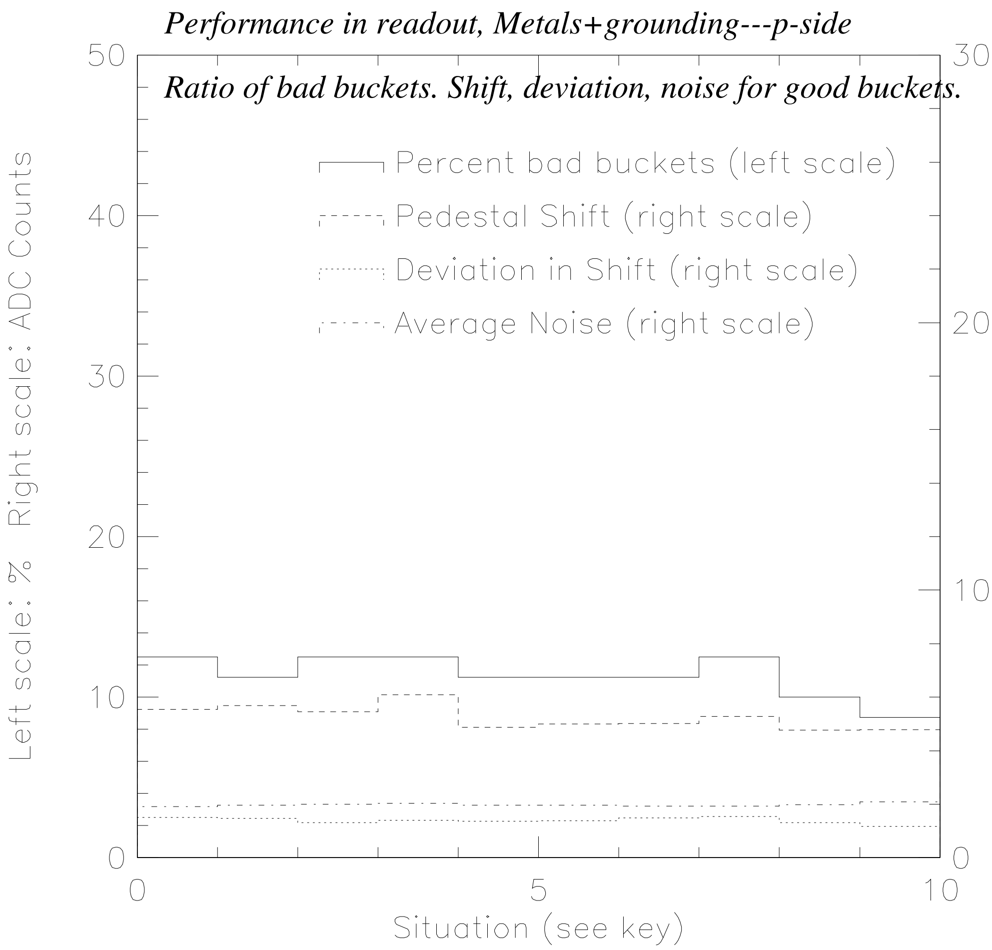

Figure 5.8: Capacitive pickup. Comparison of performance in digitize mode with and without metal pieces.

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 5.1: 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 5.2: Lost buckets with DPS capable of removing a shift of 30 ADC counts.

5.6 Dynamic pedestal subtraction

The dynamic pedestal subtraction mechanism 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.

In figure 5.9 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 5.1 and 5.2. 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 included; they correspond to preamplifier reset and mini digitize-readout operations.[1] 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.

5.7 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.

[1] 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.

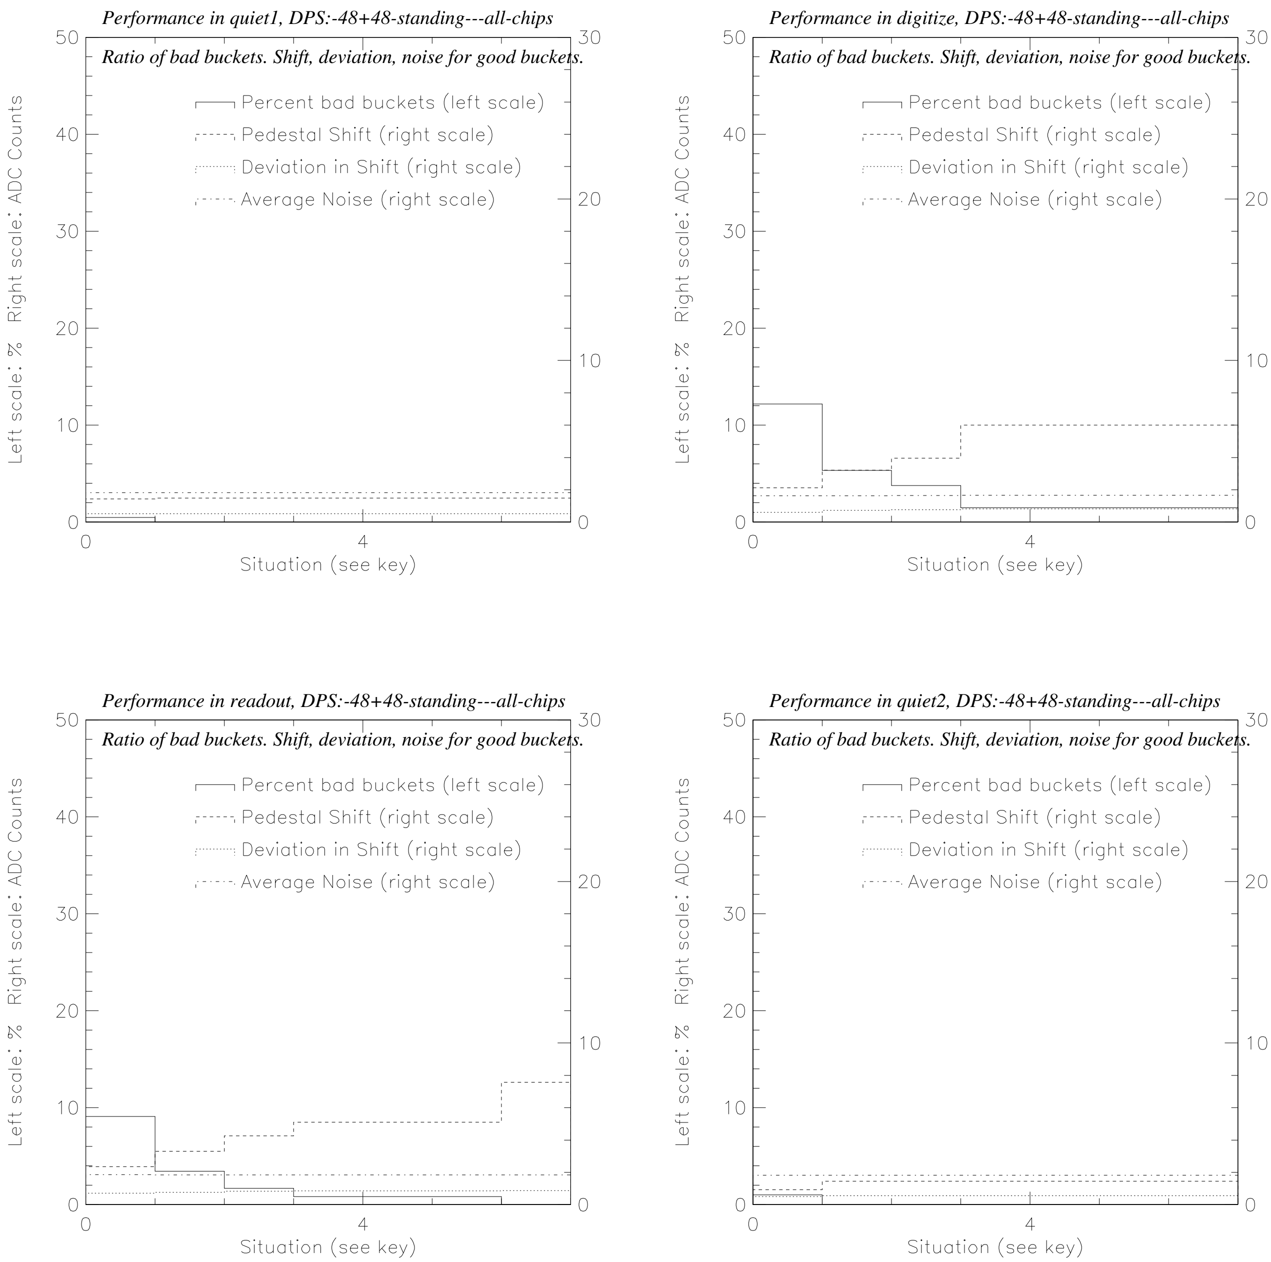

Figure 5.9: Here is overall performance in all four chip modes. The test module is fully biased, and standing on its edge, in order to move the data cable away from the silicon. The DPS is supposed to be capable of removing a pedestal shift of 0, 10, 20, 30, 40, 50 and 60 ADC counts in any direction 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.

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.

Appendix A

Definitions of noise

A.1 Noise

In this report, “noise” by itself is always understood to mean “RMS noise”. RMS noise is the deviation in a series of measurements in the following sense. Suppose a channel in the detector is read out several times, under the same conditions. They will not always be the same, but will vary by some amount. Suppose that we have read out the following: 10, 13, 11, 11, 9, 13, 12. The average value can be calculated to be 11.29. The *deviations* from this average is -1.29, -0.71, -0.71, -2.29, 2.71, 1.71. Next we calculate the *squares* of those deviations: 1.66, 0.5, 0.5, 5.24, 7.34, 2.92. Finally we take the *mean* of those, 3.03, and the square *root* of that – 1.74. Then, by the italized words, we arrive at the Root of the Mean Square Deviation, or RMS deviation to be 1.74. For those who prefer statistical notation the definition is $\sqrt{\left\langle (X - \langle X \rangle)^2 \right\rangle}$, $\langle \cdot \rangle$ denoting expectation value.

I will frequently calculate the n-side noise or the p-side noise. Those are the average noise of the n-side and p-side channels on the sensor, respectively. But because of the large number of bad channels, those had to be excluded to yield a meaningful result. Among the excluded channels are all the pulled (not connected) channels and their closest neighbors. This has been taken into account in all cases where noise figures are quoted.

A.2 Differential noise

Sometimes it is useful to distinguish between noise that affects many channels at ones, and noise in an individual channel. A measure of this can be obtained by looking at the difference between two channels. The RMS noise of this difference can be shown to be $\sqrt{2}$ times larger than the RMS noise of individual channels, *in the case when the channels are uncorrelated*. The differential noise is therefore defined as the RMS of the difference, divided by $\sqrt{2}$. The “d-noise” thus defined will be lower than the noise if the channels are correlated, higher if they are anti-correlated. The difference between noise and differential noise is therefore a hint on whether there is common mode noise in the system.

List of Figures

Bibliography

[1] Bortoletto, Daniela *et al*, *The CDF SVX II upgrade for the Tevatron run II* (Purdue University).

[2] Campbell, Peter, *Optimization of circuits and physical design in a high-speed bipolar RISC processor* (Rensselaer Polytechnic Institute 1997).

[3] Haber, Carl *et al*, *Electrical performance of the CDF silicon vertex detector* (Nuclear instruments & Methods in physics research A342 1994).

[4] Hoeferkamp, Martin *et al*, *SVX II Sensor/Preamplifier Noise Estimates* (New Mexico Center for Particle Physics 1996).

[5] Radeka, V, *Signal, Noise and Resolution in position-sensitive detectors* (Brookhaven National Laboratory).

[6] Sjögren, David, *The effects of Ionizing Radiation on the SVX3 Chip* (KTH 1998).

[7] Skarha, John, *Mechanical Design of the CDF SVX II Silicon Vertex Detector* (The Johns Hopkins University 1994).

[8] Tollnert, Christopher, *Testing of Readout electronics for the Silicon Vertex Tracker of the Collider Detector at Fermilab* (KTH 1998).

[9] Worm, S *et al*, *SVXII Silicon Strip Sensor Specifications and Conformance Testing* (Fermilab 1997).

[10] Zimermann, Tom, *Guideline to Measurements of SVX3 Chip Analog Performance* (Lawrence Berkeley National Laboratory).

[11] *The CDF II Detector Technical Design Report* (Fermilab, Pub-96/390-E)

[12] *Chronological Timeline of Fermilab Accomplishments* (http://www-stock.fnal.gov/directorate/public_affairs/history/chronology.html).

[13] *Encyclopaedia Britannica.*

[14] *Physics at Run II* (Web page collection, Fermilab 1999).

[15] *Run II Handbook* (Fermilab 1998).

[16] *SVX II web pages* (Lawrence Berkeley National Laboratory *and* Fermilab).