Design of Low Noise Detectors

Marvin Johnson

Abstract

This paper describes methods for minimizing common mode noise in electronic detector systems. It discusses grounding issues, proper design of the signal path and experiment wide methods for low noise design. These principles are illustrated by several examples.

I. INTRODUCTION

Detectors for high energy physics experiments have changed significantly in the last few years. The twin goals of higher resolution and lower cost have moved the readout electronics from circuit boards located in a counting house to dedicated chips mounted directly on the detector. This effort has resulted in much better detectors but at the expense of ever decreasing signal levels. Thus, control of electrical noise is becoming an increasingly important feature of detector design.

Since there are many good text books describing methods for minimizing common mode noise, I will concentrate on methods of applying these methods to detector design. I will illustrate these ideas with several examples that I have been involved in. The first section discusses grounds and noise currents. The next section covers some features of detector design while the last section discusses more general aspects of experiment design.

II. GROUNDS AND NOISE CURRENTS

The term "electrical ground" means different things to different designers. A designer of a radio tower wants an electrical ground that can safely absorb several thousand amps from a lightning bolt. A building designer wants a ground that can keep the parts of a building and surrounding area at roughly the same potential as the center tap of the local power transformer. Detector designers have little need for either of these features. A good detector ground has a large capacitance so that noise currents flowing onto the ground do not change the voltage of the ground. It should also have a large surface area so that the current flow is not concentrated into a small area. This minimizes any magnetic field effects. From the detector point of view such a ground makes the noise current "disappear". The vacuum shell for the large CMS magnet is an example of a good detector ground.

An important feature of noise signals is that they are almost never a voltage source. That is, the noise source has some internal resistance so shunting even a small part of the current to a ground may significantly reduce the amplitude of the noise

signal. One should always ground detectors even if the connections are not ideal.

Most detectors operate at high frequency so low frequency noise is usually not important. However, the high frequencies mean that inductance almost always dominates over resistance in determining impedance to ground. For example, a 20 cm long wire 500 μm in diameter has nearly 10 ohms of inductive impedance at 40 MHz.

III. FRONT END DESIGN

Many contemporary detector designs have average signal levels of only a few thousand electrons. To put this in perspective, if a detector is sensitive to a constant current of 56 nA for 10 nS, it will accumulate 3500 electrons in a charge sensitive amplifier. This amount of noise current can be generated from magnetic coupling between 1 cm of wire (such as a silicon strip) located 1 cm away from a conductor (such as a cooling pipe) carrying about 100 μA of 10 MHz noise current. This example assumes an amplifier with 100 ohm input impedance. The obvious solution to this noise problem is to ground the pipe. If the pipe is 5 mm in diameter and 1 meter long, its impedance at 10 MHz from self inductance is over 7 ohms so grounding the pipe at one end may not eliminate the noise signal.

This example illustrates that electrical properties of mechanical components are often important to the overall detector design. Many of the noise problems that I have worked on are the result of "unintended consequences" of other systems interacting with the readout electronics. It is also true that electrical design might solve mechanical problems. For example, one might use some of the mechanical support structure as a ground return so that overall detector mass is reduced. I think that it is important that there be one design team for the detector - not separate teams for mechanical, cooling and electronics. It may seem wasteful to have electronic engineers sit through a discussion of cooling but if the cooling pipes are conductors, how these are routed and grounded could well be crucial to the success of the detector.

Most designers do a good job on the basic input circuit for a detector. This is not the case for the return part of the circuit. This is best illustrated by an example. Fig. 1 shows a schematic of a simple liquid argon readout cell for a calorimeter. One side of the cell is at high voltage and the other side is the readout plate. Charged particles passing through the cell ionize the argon atoms. The electrons drift to

the anode and are collected by the pre amp.

Fig. 1. Circuit diagram for a simple detector circuit.

Charge flowing into the preamp must be balanced by charge flowing out of the ground of the preamp and back to the cathode of the detector cell. Otherwise, the charge on the cathode would continue to increase. An electrical circuit must be a complete path back to the starting point. Thus, when charge is collected from the argon cell, a similar amount of charge is sent out the amplifier ground which must return to the high voltage side of the cell. Think of a simple common emitter circuit shown in fig. 2.
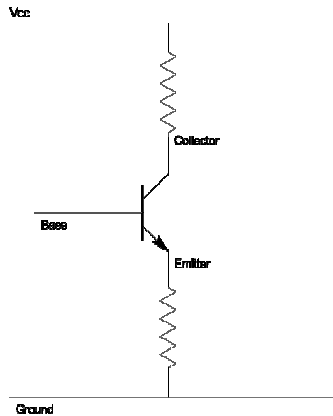

Figure 2: Common emitter amplifier. Any current injected into the base flows out the emitter and then back to its source.

Charge flows into the base and out the grounded emitter. For the circuit to be complete, this charge must flow back to the high voltage side of the argon cell. If the return path encloses any varying magnetic fields, noise signal will be induced into the circuit by Ampere's law. In particular if the return path is though a remote high voltage supply (as shown in fig. 1), the detector is likely to be quite noisy. The best design is to install a capacitor between the high voltage line to the cell and the amplifier ground (fig. 3). This capacitor should be as close to the amplifier ground as possible. Additionally, adding a resistor in the ground return of the high voltage supply will force all the return current through the capacitor as well as breaking any ground loops involving the high voltage system.


Figure 3: This is the same as fig. 1 but with the addition of a capacitor to provide local signal return to the cathode plate.

This is a straight forward design but it can have subtle problems. A muon system employing both anode and cathode readout had the following problem. The noise level was satisfactory when the chambers were installed but over the next few months the noise increased roughly linearly with time Fig. 4 shows a simplified schematic of this chamber. The designers have installed capacitors for proper return of the ground currents to the HV system. This looks fine on paper until one looks in more detail at the detector. This problem was traced to a poor ground connection between the anode and cathode boards. Both boards needed to be removed easily so the ground connection was made with a screw. Over time the surface of the screw oxidized thereby increasing the resistance between the two grounds. The actual schematic looked like the one shown in fig. 5 where R represents the resistance of the screw. As R increases, more of the return current is forced onto different paths. If these paths enclose fluctuating noise currents, some of this noise will appear in the signal. The simple solution of adding an explicit ground connection between the two circuit boards eliminated the noise problem.


Fig.4. Wire chamber with both anode and cathode readout. Note that there is only one signal return capacitor.

Fig. 5. This is identical to fig, 4 but with a resister shown in the return path between the anode and cathode amplifiers. The resistor represents the added resistance of the oxidized mounting screw.

Another example is a precision drift chamber with both anode and cathode readout. It worked well in test beams and in test setups outside the experiment. But when it was installed in the experiment and all the amplifiers installed, it would break out into stable oscillation after a few minutes. The time it took for the oscillations to start was variable. This behavior was the result of a poor design of the high voltage system itself. The drift chamber used a graded voltage system so that the drift velocities were roughly uniform throughout the detector. A schematic of the voltage distribution is shown in fig. 6.
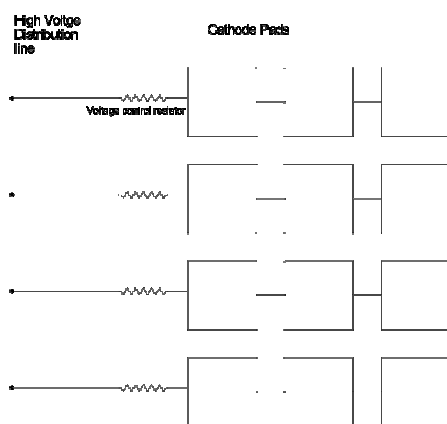

Fig. 6. Schematic of the high voltage distribution for a precision drift chamber. The high voltage distribution line was 32 times the length of the chamber.

The cathode pads were fed from a common line through resistors which set the pad voltage. This common feed wire ran back and forth across the chamber 32 times. The far end of the wire was open and the near end was terminated in a large resistor. The entire circuit was etched on a polyimide sheet and installed with the cathode pads mounted directly over the preamp inputs. The source of the oscillation was the high voltage line which functioned as a cable resonator. That is, when some of the preamp output was coupled back into this line (through accidental coupling), it excited the natural resonance frequency of the line. The most likely feed back path was through the feed back capacitor via a poorly grounded ground plane. Of coarse, the feedback from the preamps was random but the line selected out its natural frequency. When there were enough preamps feeding energy into the line, the signal exceeded the preamp threshold and the entire chamber started to oscillate at the resonance frequency of the high voltage line. The oscillations

started on noise signals so the start time just depended on achieving enough noise signal at one time to start the oscillation. A very simple fix for this problem would have been to have one line across the end of the chamber and 32 branch lines going to the preamps. The line would still have resonated but the frequency would have been above the bandwidth of the amplifier so no oscillations would have occurred.

There are many other structures in detectors such as cooling lines, cables and so on that could form resonant systems. All that is needed for oscillations to occur is a resonance in the bandwidth of the preamp, electrical coupling to the preamp input and some coupling of the preamp output to the structure. The key to preventing this type of problem is to make sure conducting mechanical structures are well grounded and electrical structures are short enough so that any resonance is above the bandwidth of the amplifier.

A third example is a wire chamber that is read out from both ends. This example reads out the cathode on one end and the anode on the other but it could also read out both ends of a wire in order to get the coordinate along the wire. A simplified schematic is shown in fig. 7.


Fig 7. Schematic of a wire chamber that is read out from both ends. All the anode channels are read from one end and all the cathodes from the other end. The gap in the ground shows that there was a very poor ground connection between the two ends.

The return circuit is only at one end and there is a break in the return ground plane. The gap in the return circuit is equivalent to an infinite value of R in the muon chamber example so one might expect that this chamber did not work at all and that was the case. The signal return path through the external electronics was so long that the phase of the returned signals was shifted to give positive feedback so that the preamps oscillated. The symptom was that oscillations would occur depending on output cable position. What was happening was that the propagation velocity of the return signal depended on the capacitance of the ground line to the surrounding world. That is, the formula for the velocity of signal propagation on a cable is

$$v = \frac{1}{\sqrt{LC}}$$

where L and C are the inductance and capacitance per unit length. When the cable position was changed, the capacitance changed which then changed the signal delay time. The overall delay was close to that needed for positive feed back so one position of the output cable would cause oscillation and another would not.

The fix for this problem was identical to the previous example: connect the grounds between the two ends. This eliminated the oscillation problem but the detector was still noisy. The circuit with the grounds connected is shown in fig. 8.


Fig. 8. This is identical to fig. 7 but the gap in the ground plane replaced by a noise generator. The gap was shorted together but the difference in potentials between the two grounds causes current to flow through the ground plane. The resistance in this connection causes a noise voltage.

I have included a noise generator in the circuit. The ground potential at the two ends of the detector is not the same so some ground current flows through the new connections. The connections have resistance so this results in a noise voltage that is directly in the return path for the one set of preamps which means that the noise is in the readout. This is a difficult problem to solve. Making the return path have very low impedance will minimize the noise. The noise can only be eliminated by isolating the grounds of one or both sets of preamps so that no external ground current can flow. The next example describes the use of ground isolation to eliminate this ground loop.

Note that adding a capacitor to provide a local return for the high voltage is likely to make the noise problem worse. Now the ground current is flowing through the high voltage plane so both sets of preamps will see the noise. Also, the high voltage plane is likely to have more impedance than a well constructed ground connection so the noise signal will be larger.

What do you do if not everything is close together? This could be a large liquid argon calorimeter where the high voltage port is separated from the signal port or a silicon detector where the preamps are connected to the sensors by a flex cable. This is just an extension of the previous example so we know the answer; either isolate the grounds of the preamps or make a very good ground connection.

Since most detectors involve high frequency signals, inductance is usually much more important than resistance. The formula for the self inductance of a rectangular conductor is

$$L = .002l\left(Log\left(\frac{2l}{H+W}\right) + \frac{1}{2} - Log(e)\right)$$

where l is the length of the conductor and H and W are the height and width of the conductor. The skin depth of copper at 1 MHz is 66 μm so H is small for most detectors. Thus, the most efficient way to distribute material for a low inductance connection is to make a wide thin sheet.

The formula for a wire (or cylinder since they are the same) is

$$L = .002l\left(Log\left(\frac{2l}{R}\right) - \frac{1}{2}\right)$$

where R is the radius of the wire. Again, we see that a large radius is important for a wire to have a low value of inductance.

Since the dependence in both cases is logarithmic, one rapidly reaches a point of diminishing returns. Also, these formulas break down as the width or radius approaches the length. But they do give us a guideline on how to proceed.

A good example of both a low impedance ground plane and an isolated ground preamp is the layer 0 silicon detector for D0. This device has a radius of only 18 mm so that the chips could not be mounted directly on the sensors. Rather, we used a roughly 300 mm long polyimide cable to attach the sensors to the chips. In order to minimize intrinsic noise, the cable capacitance must be made as small as possible. Thus, the cable was made without a ground plane. There is only one small trace to provide a return path for the bias voltage which has a resistance of 4 ohms. The impedance at 10 MHz is more that twice this value. This is far too high an impedance for a low noise design so we must use some other connection. We must also keep the overall mass as low as possible. The best option is to use some of the mechanical structures as electrical elements. The cooling lines are plastic so they will not work. However, the body of the device is a 12 sided carbon fiber polygon with a diameter of 35 mm. A 35 mm cylinder 300 mm long has an inductive impedance of less than an ohm at 1 MHz. If we can make the support structure conductive, our problem is solved.

High modulus carbon fiber is quite conductive if one can make good electrical contact with the carbon fibers[1]. We have developed a method of taking 50 micron thick polyimide film coated with a 5 μm thick layer of copper, etching a mesh pattern on it and then co curing the polyimide with the carbon fiber. That is, we etched a mesh ground plane on a piece of 50 μM thick polyimide that was coated with a 5 μM thick layer of copper. Standard printed

circuit vias were used to bring contacts to the reverse side of the material. This material was laid up with the copper layer facing the carbon fiber and the assembly was cured as a unit (fig. 9). This process results in a very low resistance device that is within a factor of 2 of an all copper structure at high frequencies The high voltage coupling capacitor was mounted on the sensor so that the support structure remained at ground. The trace length between the capacitor and the sensor bias plane was kept as small as possible.


Fig. 9. Copper mesh co cured onto the carbon fiber mechanical support.

The conductivity of a detector's mechanical structure can be very useful for some aspects of detector design but it can also create ground loops through the detector. Any detector with multiple independent readout sections is subject to possible ground loops. All one needs is to have different sections grounded to different locations and a conducting path through the detector. This was described in the third example above. The usual solution to this problem is to provide a dielectric break in the mechanical design that isolates the different readout sections. Sometimes design constraints prevent this. This was the case for the layer 0 detector. The small diameter and long length of this detector required a continuous carbon fiber structure. The only solution that eliminates this loop is to isolate the local electronics ground from the outside world grounds so this is what we did.

We can break up the isolation problem into three main parts: 1)the download and readout system, 2) the power supply, and 3) the circuit board layout. The readout system for layer 0 is LVDS so we chose to use the differential drivers themselves to isolate the readout. Other methods such as optical or magnetic coupling were studied but none were satisfactory for this environment. Power supply isolation was achieved by using a separate power supply for the isolated system. We selected a supply that had good AC isolation at high frequency. The circuit board was designed with minimum overlap between the two grounds. With all components

installed, the resulting board had 33 ohms isolation between the grounds at 7 MHz with little frequency dependence. There are six boards in parallel in this system so the overall impedance is around 6 ohms. This is a lower limit since the cables connecting the detector to the outside world have some inductive impedance.

IV. GENERAL TECHNIQUES

A. Ungrounded Conductors

One very common problem with detectors is isolated conducting components. By this I mean pieces of metal that are isolated from the rest of the detector. This isolation can be caused by oxidized aluminum or by glueing parts together with non conducting adhesives. Either one of these results in a conductor which can be at an arbitrary voltage. Signals induced on these components spread over the entire surface of the component by Gauss's law (fig. 10).


Fig. 10. Schematic of a noise source coupling to a front end circuit through an ungrounded piece of metal.

If one part of the component is close to a sensitive part of the circuit, the noise may simply be channeled directly to the sensitive component. Ungrounded components are one of the most common problem areas in detectors.

Bare aluminum oxidizes immediately. In order to ground aluminum one must establish a connection through this oxide to the base metal. This can be done by either a mechanical connection or by plating the metal. I will cover plating first. There are 2 common plating methods: Alodining and tin plating. Both methods work well but they have somewhat different applications. The Alodine process coats the Al with a coating that is only a few molecules thick. Thus, there is no change in the dimensions of the parts but the surface is easily scratched. It is most suitable for parts that have critical mechanical dimensions and will not be disassembled often.

Tin plating typically coats the material with 250 µM of tin so it has good mechanical robustness but the parts have grown in size. This method is good for cable trays and other parts that may need to be disassembled or be exposed to rough handling.

A mechanical connection can be made by use of a star washer or similar mechanical device. Star washers are lock washers with many sharp points. If they are properly tightened, they will cut through the

aluminum oxide and form a good connection to the aluminum underneath. The main problem with this method is maintaining enough pressure on the washer to maintain a gas-tight connection. Otherwise, the aluminum will reoxidize under the washer. With careful application, these connections can last for several years.

Sometimes the aluminum oxide problem is only recognized after the detector is completed. There are some things that can be done after assembly. One is to use star washers. However, if there are few mechanical connections, this will not work. A second solution is to use a product called an alodine pen which allows alodining small sections of an aluminum part. One can then make reasonably good connections with only mechanical pressure such as with a clamp.

B. Power Distribution

Transformers come with 0, 1 or 2 shields. A single shield is typically made as conducting screen between primary and secondary coils. A doubly shielded one usually has screens wrapped around the primary and secondary coils. A single shield reduces noise from capacitive coupling between primary and secondary by a factor of about 100. Two shields give an additional factor of 10. A single shield is connected directly to ground. Both shields of a doubly shielded transformer can be connected to a local ground but a better way is to isolate the ground of the secondary. That is, the secondary is attached to a ground isolated system. Then the shield of the secondary is connected to this ground. This arrangement would be suitable for a very sensitive experiment such as a dark matter search.

There is a serious safety issue with an isolated secondary ground. If the transformer fails, the ground of the secondary could raise to the voltage of the secondary. Since the grounds are isolated there would be nothing to trip the primary circuit breaker. This problem can be eliminated by attaching a saturable inductor between the two grounds. At very low currents, the inductance is high and there is a break between the two grounds. If a larger current flows (few hundred milliamps), the core saturates, the relative permitivity drops to 1 and the coil will present very little resistance to current flow.

C. Cables

Covered cable trays grounded every few meters to a good instrument ground provide the best protection from noise pick up in signal or power cables. If a covered tray is not possible, lining the bottom of the tray with a thin copper foil will usually give some benefit. it provides magnetic shielding for fields from below and it forms a ground plane for the cables passing over it. Of course, this works best for cables that lie directly on the copper ground plane. Also, the copper must have periodic ground connections.

D. Cable Shields

Grounding cable shields is often controversial. It is usually best to ground only one end of the cable. Otherwise, you risk forming a ground loop. Sometimes a capacitor coupling is used at one end to break the low frequency ground loops. It is important to choose the best ground for the cable shield which could be either at the source or destination end. If the grounds are equal, I usually choose the rack end rather than the detector end because I want to route energy picked up by the cable away from the detector.

E. Racks and Other Infrastructure

Racks and other support structure should be welded together and connected with a low impedance connection to a good ground. This is especially true if the rack is being used as a cable ground. Connections between structures that are painted and bolted together are rarely adequate for for a good ground.

V. SUMMARY

Successful design of any low signal level device require great attention to detail by skilled designers. I find it very useful to draw a simplified schematic of the entire detector including all the mechanical components that are potential conductors. Coupling strengths can be estimated by using simple formulas or by using various field calculation programs. I have been quite successful using finite element codes to calculate the capacitance between circuit elements. Once this is done, one can eliminate components that have negligible coupling to the electronics and make sure the others are adequately grounded.

VI. REFERENCES

1. W. Cooper et al. Nucl. Inst.and Meth. A 550 (2005) 127