Floating Nanoammeter

Adam West

March 29, 2017

1 Introduction

While performing tests with the electrostatic lens, we found that X-rays were being emitted due to field emission from the electrodes. This phenomenon was accompanied by a significant leakage current due to the electron emission, up to several μA in size. We also found that the current reported by the power supplies is not particularly precise or reliable; tests with nothing connected to these supplies showed reported currents at the μA level which were also correlated with the size of the voltage applied.

Since leakage current is an indicator of field emission, we would like to monitor it precisely to ensure it is minimised. To do this we are using a floating ammeter which closely follows a design developed by the Imperial group: link.

2 Design

The basic design is shown in Fig. 1 The current flowing through a sense resistor ($R_{\text{sense}} = 10\ \text{M}\Omega$) gives a voltage

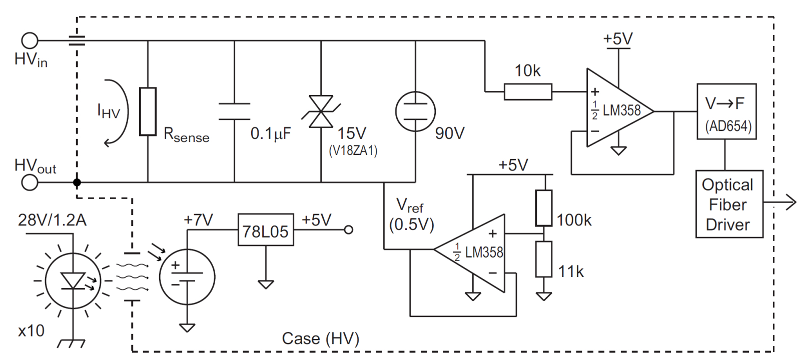

Figure 1: Simple schematic of the ammeter circuit design from the Imperial paper. See main text for a description.

drop that is fed through a follower to a voltage-to-frequency converter which outputs optical pulses. The voltage, V_{sense} is related to the frequency, f by V_{sense} A voltage divider and follower are used to give a 0.5 V offset which allows for bipolar current sensing. The whole circuit is powered from a photovoltaic converter. Note that we use a laser rather than LED to provide the optical power, and we do not float the enclosure.

The more detailed schematic for the circuit we use is shown in Fig. 2. Jim was responsible for the design and construction of the circuit. Some small modifications were made to improve upon the Imperial design, most notably reducing the power consumption by significantly reducing the pulse width output by the voltage to frequency converter.

There are a number of elements in the circuit designed to protect it from high voltages (that is, large potential differences at the input), including a GDT, capacitor and varistor.

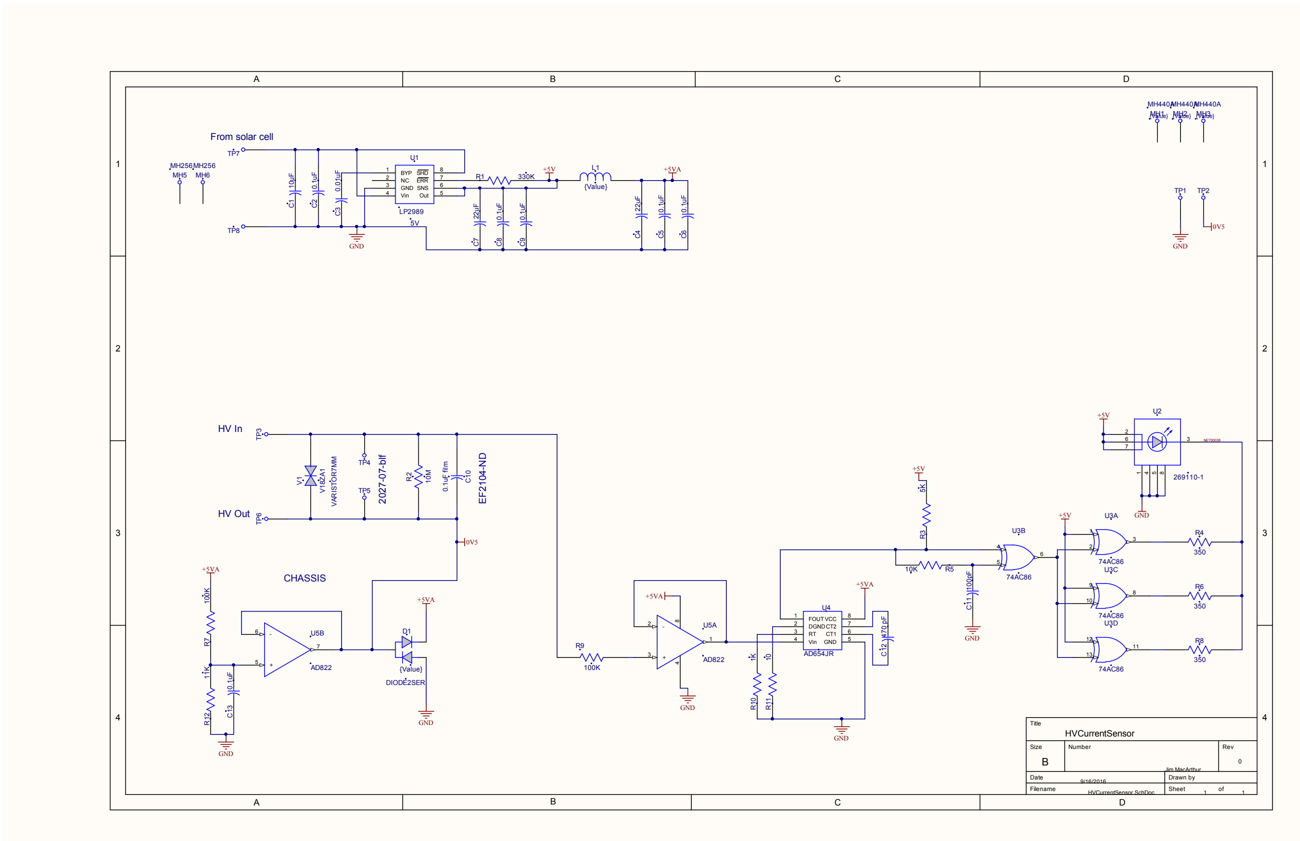

Figure 2: Circuit diagram for the floating ammeter. Some small changes to the output signal pulse generation are not reflected here.

3 Testing

3.1 Linearity

3.2 Precision

3.3 Stability

3.4 Durability

Since the circuit we are using will be inaccessible after it has been potted for high voltage, we want to make sure that it will not break when exposed to possible damaging events, such as arcs in the apparatus. The first test I performed was to ensure that the device was happy operating at a fairly high voltage. For this I floated it up to 5 kV, without potting. When doing this, I found that there were a couple of problems.

Firstly, I observed arcing — an audible spark followed by the power supply tripping. This was found to be due to the anti-static tote in which the circuit was placed; since the tote material has some small but non-negligible conductivity we concluded that something was arcing to it.

Secondly, the current reported by the device was sensitive to other objects/people moving in its vicinity. After discussing the effect with Jim he concluded that RF noise was being picked up and aliased within the circuit. In particular, noise from components ‘upstream’ of the ammeter were likely causing part of the problem. Swapping HV In and HV Out (so as to be opposite to that in Fig. 2) helped mitigate the effect — because the connection labelled HV Out acts as a reference, offset by 0.5 V to allow bipolar operation, it was less sensitive to AC noise than HV In. This, combined with some non-linearity in the circuit, was our best guess for what was causing the effect. Swapping the connections was seen to reduce the effect somewhat.

Following the simple HV test, we performed a more destructive test whereby an arc was deliberately created, so as to check if the circuit would survive such an event, which is likely to occur in actual operation. For this, I designed and machined a spark gap, shown in Fig. 3. This gap was placed in parallel with the ammeter and the

Figure 3: Spark gap design. Spherized aluminium electrodes have an external tap so that they screw into aluminium cylinders. These cylinders attach via set screw to nylon cylinders. On top of the aluminium cylinders are 1/4-20 threads to make connections via spade connectors. The bottom of the nylon cylinders are threaded such that you can screw in a pedastal base so that they may be clamped to a breadboard. The separation of the electrodes can be finely adjusted by screwing the electrodes in or out.

voltage applied gradually increased until a spark was produced. At this point, the power supply would trip. This was then repeated with increasing gap until the spark voltage was $\approx$5 kV (the limit of the power supply). During all of these tests, no adverse effects were observed on the ammeter. You can see the spark gap in operation in the video at this link.

3.5 High Voltage Testing

3.5.1 Run 1

In this section I describe the tests that were performed at voltages in excess of 5 kV.

The first set of data taken was without the floating ammeter in place, with the high voltage passing from the supply into a bullet which was potted inside a box. The data are shown in Fig. 4. As the voltage is stepped (via a 0–10 V control voltage), we see the voltage reported by the supply step accordingly. We also see spikes in the current reported by the supply, and we also see the current rise from zero to around 0.7 µA — this is why we need the floating ammeter! The ammeter, which is not connected, is quite stable with mostly just the odd fluctuation of a few pA.

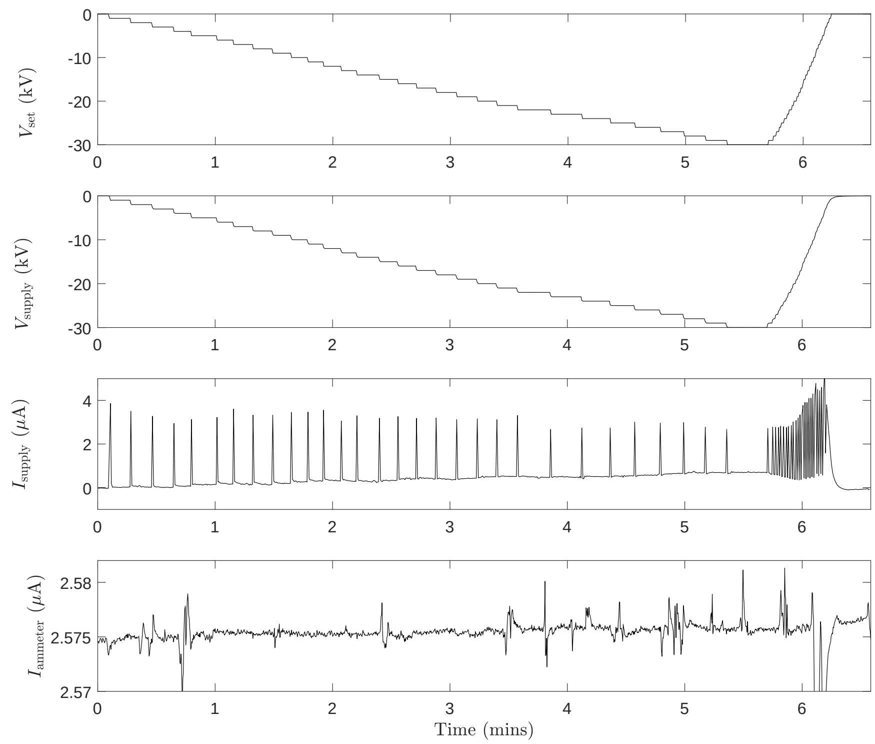

Figure 4: Data taken with the high voltage supply connected directly into a potted 'bullet'. The nanoammeter is not in the high voltage circuit.

Next, the ammeter was included, between the supply and the bullet. Some weird behaviour was seen. First, it was noted that there were some oscillations in both the voltage and current reported by the HV supply. This is illustrated in Fig. 5, which shows some data taken at a nominally constant voltage of -1.9 kV. The oscillations are quite slow, with a period of ∼1 minute. We also note that the voltage and current oscillations seem to be approximately $\pi/2$ out of phase, which one would expect for something charging and discharging[1]. The voltage oscillations are quite small — around 1 V out of 2 kV, i.e. ∼0.1 %, however the current oscillations are quite significant, at around 40 nA amplitude. There is no sign of such oscillations in the nanoammeter readout.

For now, I don't have a model for what causes these oscillations. It is notable, however, that they also occur when there is zero voltage applied, as shown in Fig. 6 If we look back at the data without the ammeter connected, we do not see these kinds of oscillations, which is quite peculiar. This would suggest something happened to make these oscillations start happening, or that they are caused just by the ammeter circuit being connected, independent

[1]Note that there may be a sign error on the current, which I should check.

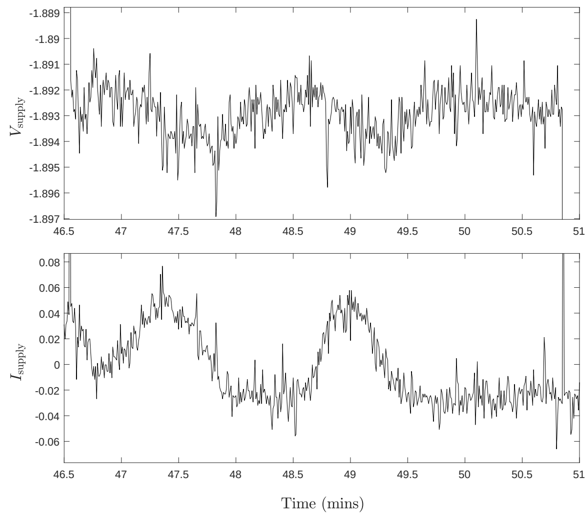

Figure 5: Oscillations in the voltage and current reported by the power supply with the nanoammeter connected at high voltage. Again the circuit is terminated by a potted 'bullet'.

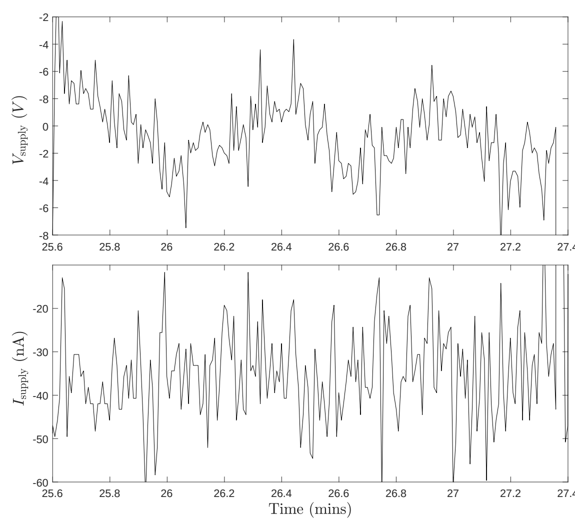

Figure 6: Oscillations in the voltage and current reported by the supply with a nominally zero voltage.

of any high voltage effect. It would perhaps be interesting to go back, remove the ammeter, and verify that the oscillations do indeed go away.

Whilst these oscillations are quite peculiar, they do not necessarily prevent the apparatus from operating as required, since the oscillations do not show up in the nanoammeter readout. However, there was other, more concerning behaviour.

If we look at the current reported by the ammeter as the voltage is increased, we see that there is a significant change, as shown in Fig. 7. Each time the voltage is changed, the current reported by the ammeter settles, over a

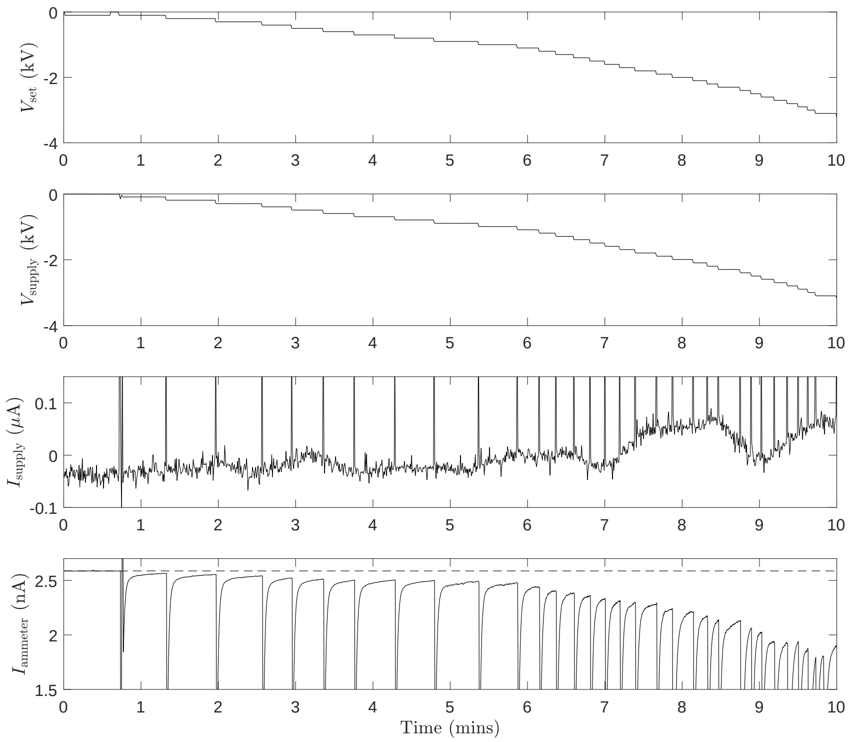

Figure 7: Change in the current reported by the nanoammeter as the voltage is increased.

fairly long time period (increasing with voltage), however it is still quite clear that the steady state value changes significantly as the voltage is changed. In Fig. 7, the voltage is increased to around 3 kV and the nanoammeter's reported current changes by around 50 pA. This is perhaps not that worrying in itself, since 50 pA is small, however, at these low voltages we would really not expect such a change, particularly as we were able to demonstrate greater stability before the ammeter was potted up — between 0 and 5 kV we observed a change of ≈10 pA.

A second unusual effect that we observed was an instability in the nanoammeter reading. At higher voltages, the readout would spike in a quite regular manner. This is illustrated in Fig. 8. The nominal voltage was set at

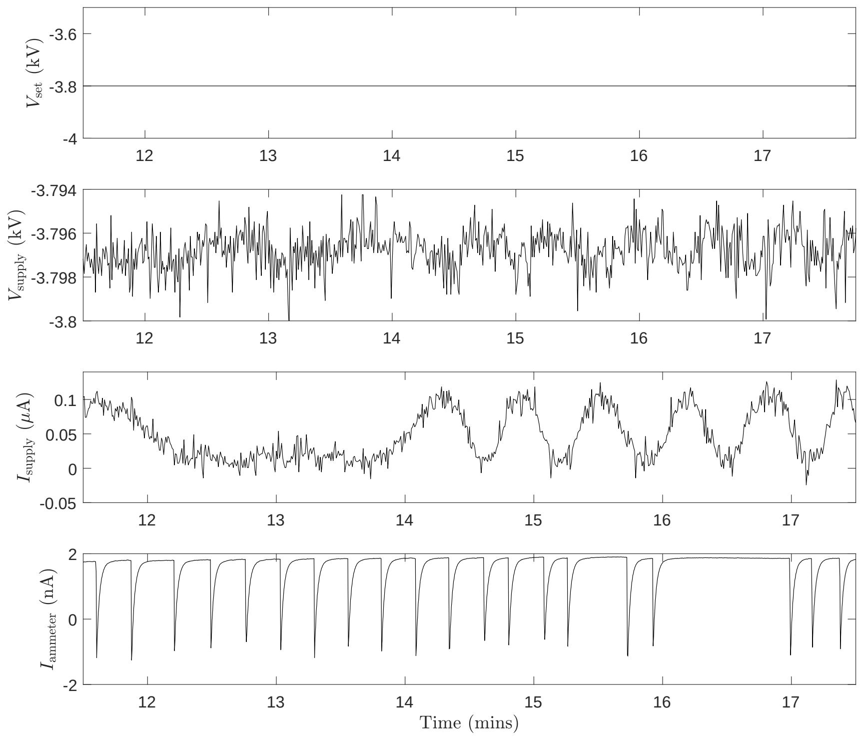

Figure 8: Current spikes reported by the ammeter with a constant voltage applied (-3.8 kV).

-3.8 kV and unchanged. Note that the HV supply reports it to be accurate to this value within a few volts, and that both the voltage and current it reports show the aforementioned oscillatory behaviour again. However the most striking behaviour is from the ammeter, which shows regular spikes in the reported current of around 3 nA in size, occuring every 10 s or so. We saw that this spiking at constant voltage began to occur when the voltage was raised to around 3 kV, and that the frequency of such spikes increased as the voltage increased.

Again, I am not sure what is causing this behaviour, however I believe that the current is real, since it is clearly voltage dependent. If it is indeed real, the question is of course where is it. It seems reasonable to assume that it is coming from the nanoammeter circuit since it represents a big array of emission points. We cannot see the same spikes in the current reported by the supply, however we would not expect to since the magnitude of the spikes is less than the noise. If we change the polarity of the applied voltage we see very similar behaviour, but with the sign of the current spikes reversed.

The data in Fig. 9 show the behaviour as the voltage is increased further. The aim was to increase in steps of 1 kV, unfortunately the voltage was mistakenly stepped from 10 kV to 20 kV at around the 3 minute mark. From 0

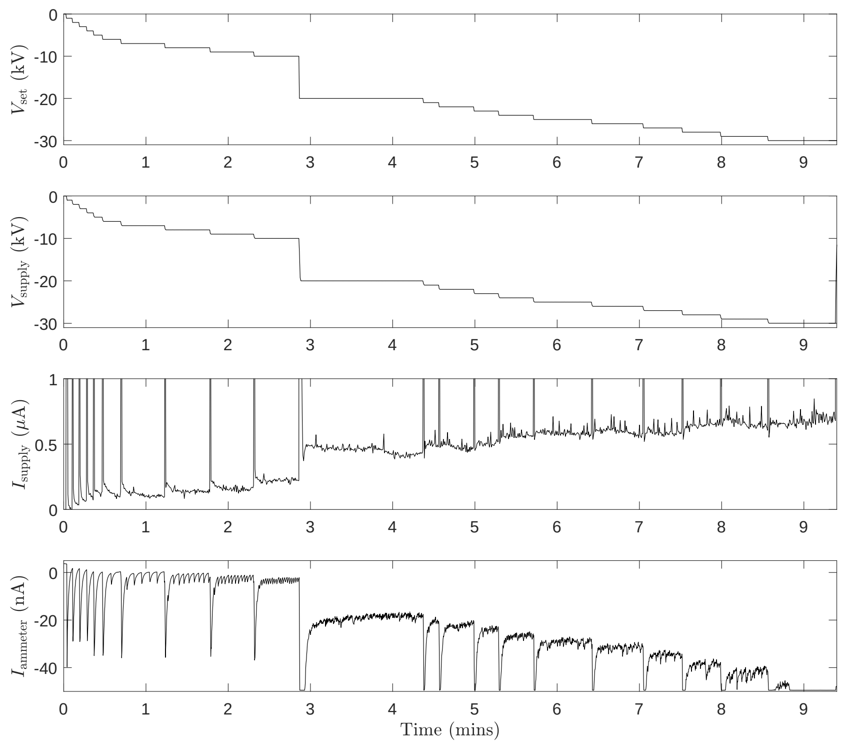

Figure 9: Data taken whilst increasing the voltage up to its maximum value of 30 kV. Note the large rise in the current reported by the nanoammeter.

to 10 kV one can clearly see the spikes in the nanoammeter readout, inbetween voltage steps. You can also see that their frequency increases with voltage. Unfortunately, as the voltage is increased further, the nanoammeter offset continues to increase, quite significantly. Indeed, at the maximum voltage of 30 kV, the current reported by the ammeter reaches its maximum value of 50 nA[2]. Additionally, the signal becomes much noisier, with fluctuations of several nA.

3.5.2 Run 2

After observing the behaviour described above, I returned to do some further testing, to try and discern whether the observed leakage current was just in the ammeter box or elsewhere as well. While repeating tests in the previous configuration, I saw sparking when the voltage reached around 28 kV. Fig. 10 is perhaps the best picture of it. It is occurring between the two optical fibers, i.e. the fiber supplying power to the ammeter circuit, and the fiber carrying the data out of the circuit. This behaviour was very repeatable, as shown in the following montage: link

As the voltage neared the level where these sparks happen, I could hear what sounds like microdischarges occurring - it was hard to locate where.

The distance between the two fibers is approximately 1.5 cm. The dielectric strength of air is around 30 kV/cm. This suggests that the voltage difference between the fibers must have been close to the total applied voltage in order for this breakdown to occur, i.e. one fiber was at approximately ground, and one was at approximately 30 kV.

The inside of the ammeter box is shown in Fig. 11. On the right we see the two fibers. The left-hand one is used to provide the light to the photovoltaic converter which powers the device. It has a metal connector/ferrule.

[2]The maximum value is a little larger for current of the opposite sign

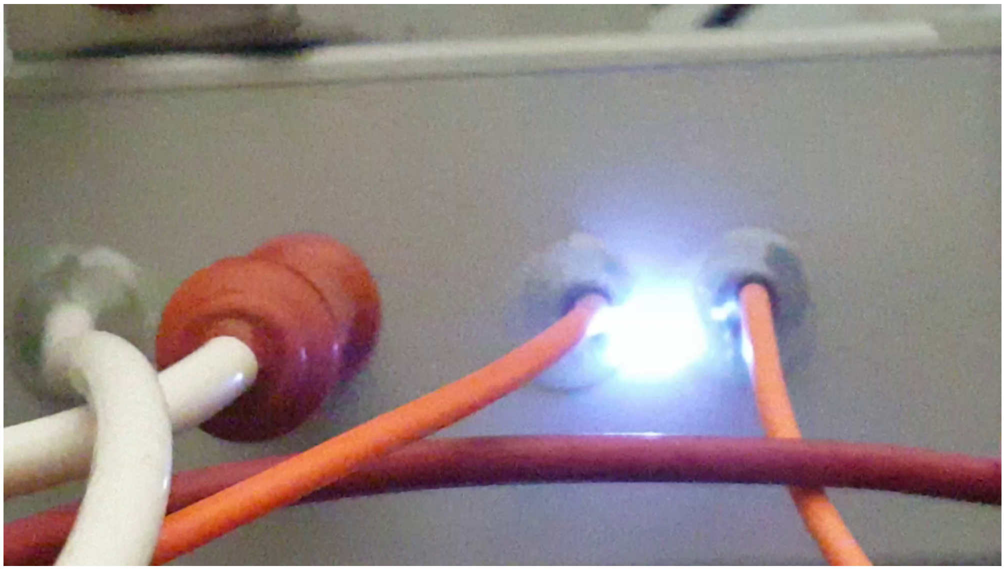

Figure 10: Spark observed between the two optical fibers connecting to the ammeter.

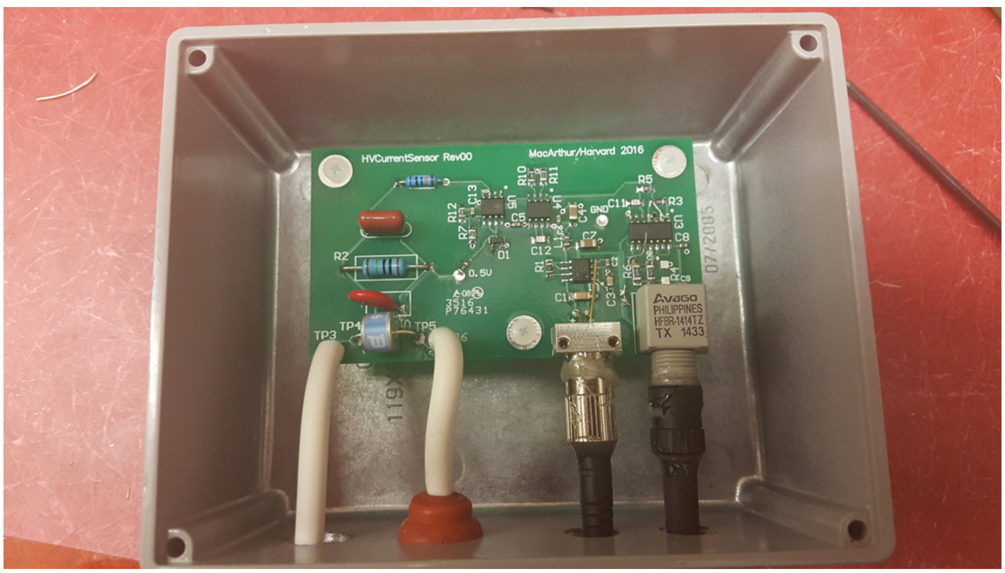

Figure 11: The inside of the nanoammeter box before being potted.

The right-hand fiber is used to extract the signal from the device. It has a plastic connector. The fibers are glued in place with 5 minute epoxy in order to prevent any potting compound entering the connector. On the left we see the connections for the high voltage. The fibers are also epoxied in place at the box, to prevent leakage of potting compound.

Whilst the fibers are nominally insulating, they will have a non-infinite resistance. I tried to measure this resistance using an electrometer but was not able to — even without the leads connected to anything, the electrometer read around 20 GΩ initially, gradually rising up to its limit of 200 GΩ. This may be worth looking into further.

In future we should improve the setup by having a larger section of optical fiber potted inside the box and grounding the optical fibers to the box as they pass through — for this we can try using conductive epoxy.

I was able to extract the circuitboard from the potted box; cutting through the box and potting compound to close to the circuitboard allowed the silicone to then be torn off. Fig. 12 shows a couple of pieces removed. We see that the details of the circuitboard are clearly impressed on the silicone, suggesting that the efficiency of filling and

Figure 12: Pieces of silicone that were used to pot the ammeter.

conforming to the surface was high. In particular, the bottom picture shows some very small features which were filled. In addition, whilst removing the silicone there was no evidence of air gaps or bubbles at all.

After extracting the circuitboard, I gave it to Jim to perform a post-mortem. He found that there was damage to an op-amp through a diode, and he believes that the source of the damage is a discharge emanating from the photovoltaic converter that passed to a nearby capacitor. A picture of the capacitor, taken through a microscope, is shown in Fig. 13. We see that the solder joint on one side of the capacitor seems to be damaged — it looks like there is a small crater taken out of the solder on one side. There is no evidence of damage from any nearby components. To try and prevent this in future we could try moving the photovoltaic converter further away from the rest of the board components. Also, the case of the photodiode (and hence presumably the case of the photovoltaic unit as a whole) is currently not connected to anything, so it may be better if we simply connect it to either the positive or negative terminal of the photodiode.

3.5.3 Run 2

From hereon, when reporting the current read by the floating ammeter I have removed the zero-voltage offset.

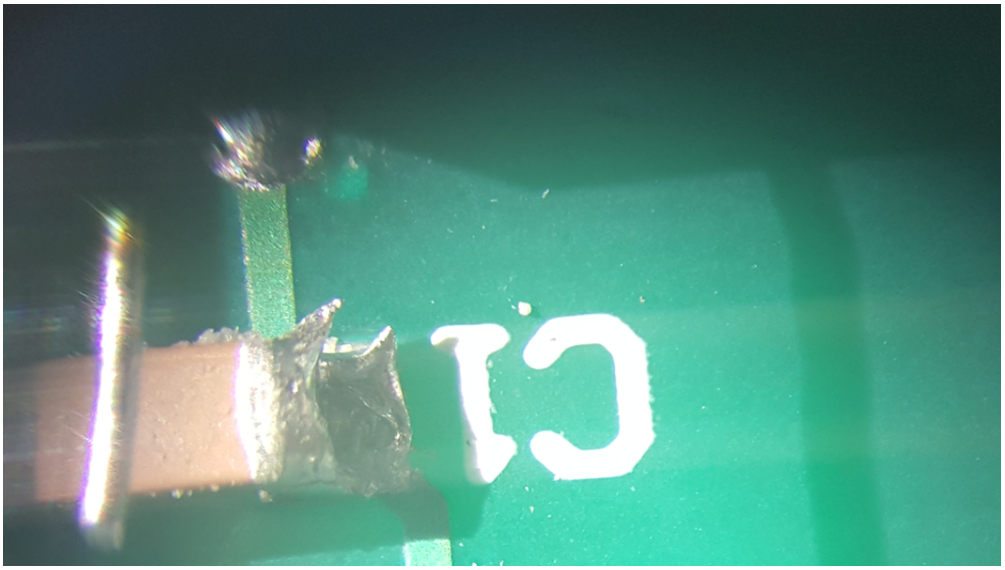

Figure 13: Picture of a surface mount capacitor which looks as if it may have been damaged by an electrostatic discharge.

For the next iteration of testing we made a few changes. To avoid the problem of arcing between the optical fibers we used silver solder to seal around the fibers where they passed through the metal box. This should explicitly ground them and prevent high voltage passing along their surfaces to outside the box where they can arc. We also explicitly tied the case of the photovoltaic converter to the positive lead of the photodiode to prevent the case floating to an unspecified voltage, which may have encouraged discharging to occur inside the box. Additionally, the box in which the ammeter was potted was made significantly bigger so that the high voltage wires and the fiber optic cables were inside the silicone for a greater distance before exiting the box. Fig. 14 shows pictures of this new box.

After these changes, we tested the system again to see if the problem of arcing had been mitigated, and if the leakage current which we had previously observed had been reduced. The resulting data are shown in Fig. 15.

Firstly, we did not see or hear any evidence of arcing happening between the optical fibers — it seems that this problem was solved. However, the current levels reported by the floating ammeter were still quite high; we see that the current has increased by 1 nA when the voltage is at -6 kV, 5 nA by -20 kV, and reaches a 10 nA increase for -30 kV. Similar behaviour is observed with a positive applied voltage. We also see that there are large (greater than 10 nA) spikes with every 1 kV step in voltage, and at higher voltages the current becomes quite unstable, fluctuating by several nA for a nominally fixed voltage.

If we plot the reported current vs the applied voltage, as shown in Fig. 16, we can see the general correlation a little more clearly. The large spread in current values is due to the spikes associated with stepping the voltage.

After observing that the problem of leakage current had not gone away, we decided to try insulating the floating ammeter with transformer oil instead. The oil that we are using is Shell Diala S4 ZX-I and was purchased from here. The dielectric strength of the oil is very comparable to the silicone, but there should be no ambiguity as to whether it is conforming to the shape of the PCB etc. since it is liquid. It also makes it easier to remove the PCB after testing, make changes, and retest it. Images of the bucket which was modified to house the ammeter are shown in Fig. 17. The high voltage connections to the ammeter are made in a very similar manner — silicone cables are epoxied in place on the lid. The optical fiber connetions for the power and signal are made via bulkhead connectors which are screwed into the lid. A separate connection is provided for explicit grounding of the bucket. The fiber bulkhead connectors have good connection with the ground connection. The PCB is suspended from the lid by 4 inch long nylon standoffs which are screwed into the lid. The fiber connection on the PCB were again sealed with 5-minute epoxy however we note that it is likely that during assembly/modification of this setup that oil got into the fiber connections on the lid, and did not seem to prevent successful operation.

One small drawback in the current design is that the fibers are much longer than they need to be and are thus coiled up inside the bucket. They were tied down to a nylon disk in the bottom of the bucket with mushroom tape

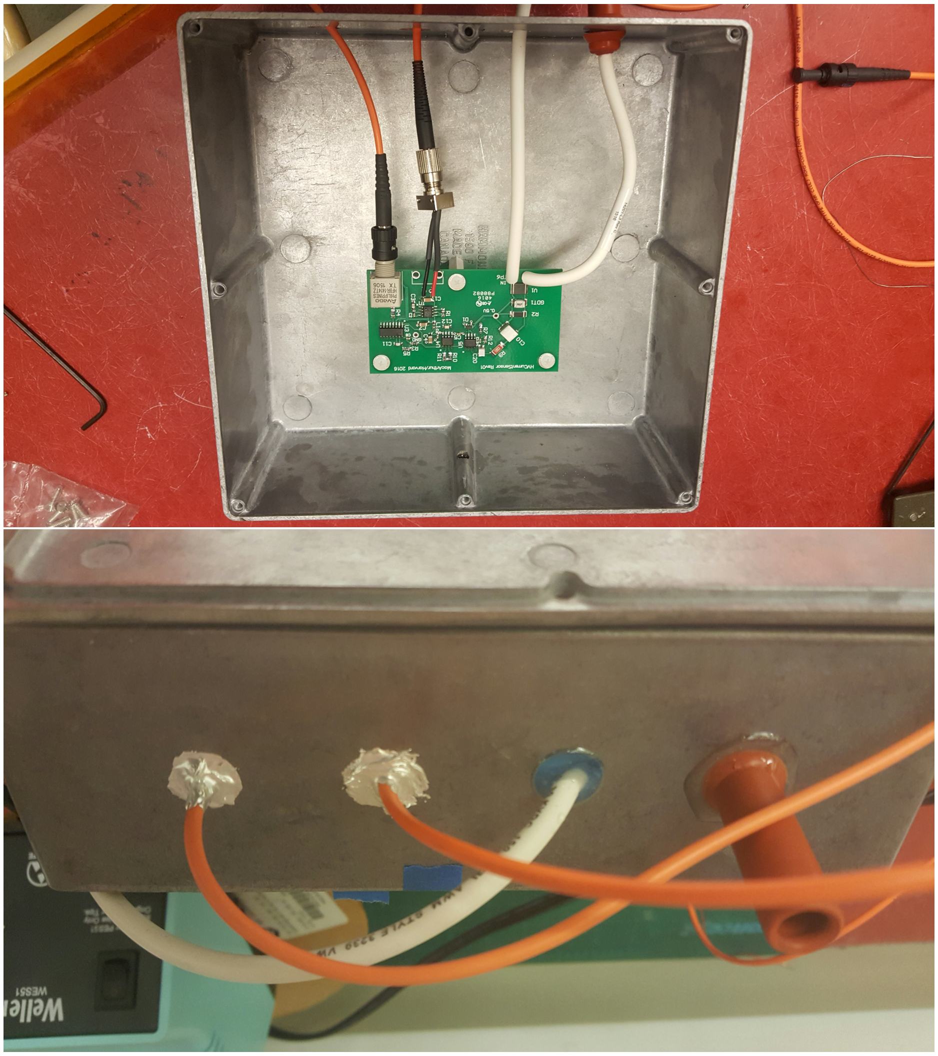

Figure 14: A floating ammeter was potted with silicone inside a new, larger box. Silver epoxy was used to seal around the optical fibers so that they were better grounded.

so that they would not touch the metal bucket. However, this made removal/reinstallation tricky as it requires one to reach into the oil to remove/arrange the fibers!

All of our cable assemblies are hipot tested to ground. If the assembly does not have a shield it is laid across a ground plane or other means and tested. As a rule of thumb the assemblies generally draw less than 1A. Allowable limit is less than 3A. Unfortunately our hipot test set does not measure below the A level. In theory if you went to the larger 14 series and the diameter between the bulkhead and the conductor was larger the leakage current should be reduced. Adding a semiconductive layer between the bulkhead and the insulation would also help smooth out the ground.

I would recommend abrading and priming the bullet connection to eliminate current caused the air ionizing around the bullet. Air can ionized at voltages above 500v.

I would also recommend that any insulation that will be encapsulated be abraded prior to encapsulation as well as adding a taper to the end of insulation. This will ensure there is no trapped air and you achieve a good bond to the cable insulation. May also want to look at the curing system of the potting and make sure it is compatible with the cured cable. Ie Platinum cure versus peroxide cured systems. Most potting materials tech data sheet should say if they are not compatible with certain cure systems.

Not sure what 5 minute epoxy you are utilizing to secure the leads prior to breakout but silicone acts as a sponge and will absorb some chemicals which can affect its dielectric properties. You could try baking the box at a sufficient

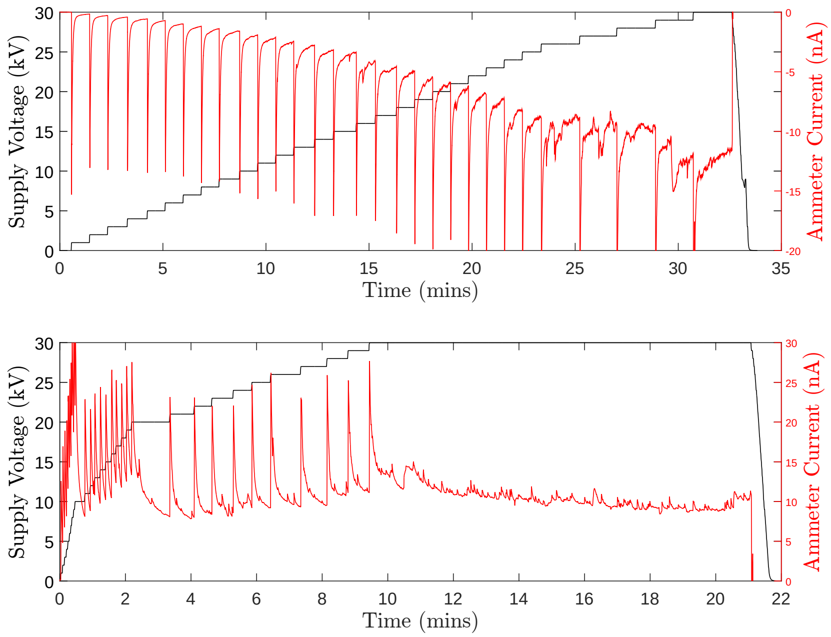

Figure 15: Current reported by the floating nanoammeter and voltage reported by the power supply as a function of time. Ammeter is in silicone in a large box. The top (bottom) plot shows data with a negative (positive) applied voltage.

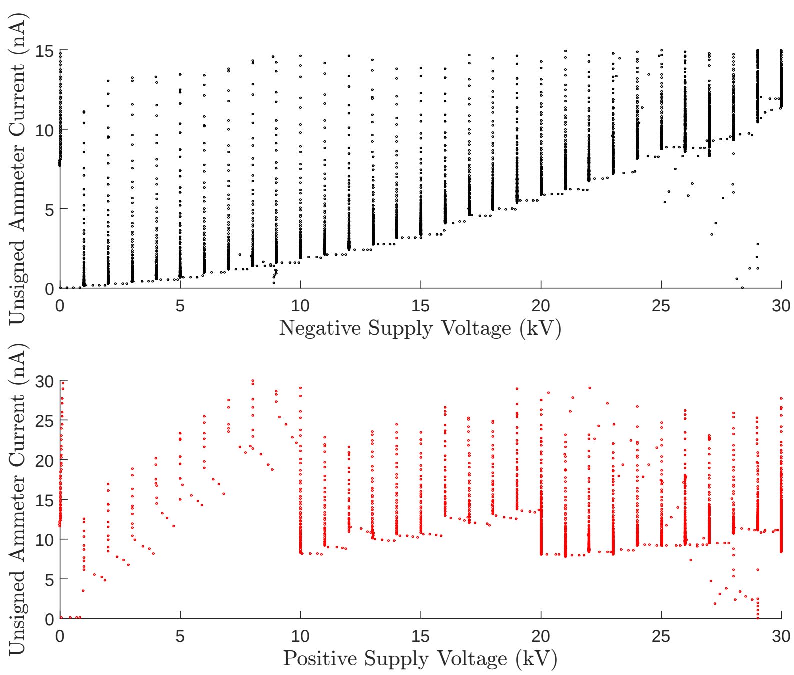

Figure 16: Current reported by the floating nanoammeter vs voltage reported by the power supply. Ammeter is in silicone in a large box.

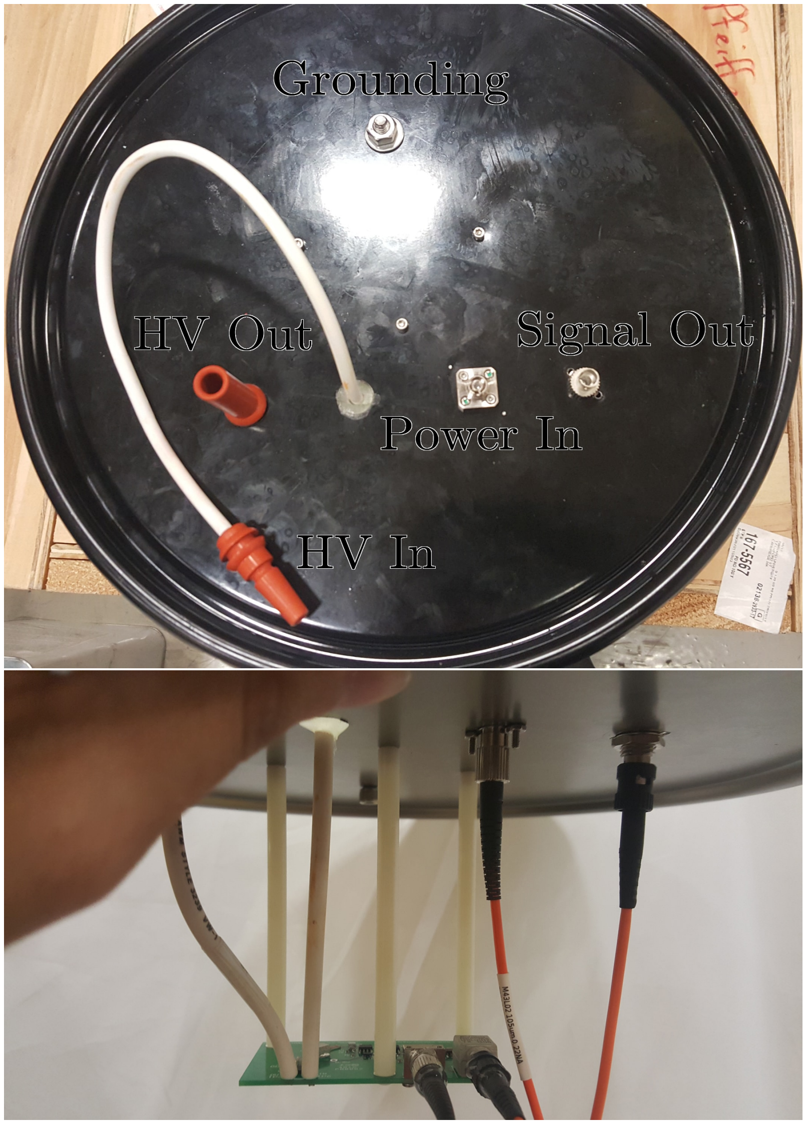

Figure 17: Nanoammeter mounted inside a bucket which will contain transformer oil. The top image shows the connections in/out. The bottom image shows the PCB mounted to the underside of the lid.

temperature to drive off anything that may have been absorbed by the silicone without damaging anything. Prior to doing so I would ensure the potting is fully cured. Most cure systems are 90% cured within 24 hours but do not fully outgas and completely cure for about a week. I ran into this issue on an application that I build recently. To eliminate the silicone from absorbing the chemicals in the epoxy I wrapped the leads with Teflon pipe thread approx 10 wraps prior to encapsulating them and this seemed to elevate the issue.

3.5.4 Run 4

We suspected that there may have been some surface effects occuring which were producing the observed leakage currents, in particular along the standoffs we were using. The standoffs were made of nylon (and were purchased from McMaster). While nylon is a good insulator, Teflon is supposedly better, both in terms of bulk resistivity and surface resistivity (although these properties are seemingly not very well known). To try and see if this was the problem, I bought some Teflon stock and machined some new standoffs (also 0.5 in longer). After replacing them, we performed the same tests as before — slowly ramping up the voltage and observing the measured current. The results were very similar, perhaps even slightly worse, i.e. several nA of current.

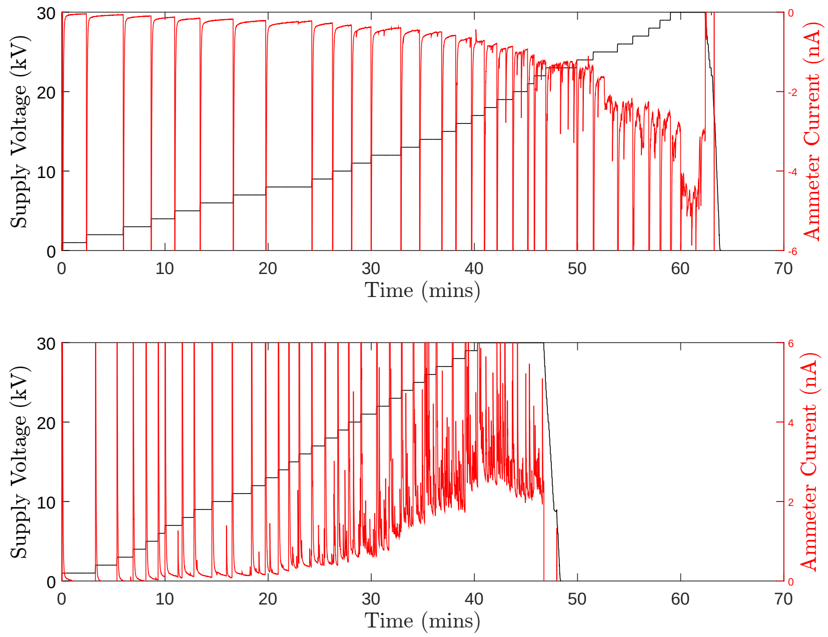

Figure 18: Current reported by the floating nanoammeter and voltage reported by the power supply as a function of time. Ammeter is in bucket of oil. The top (bottom) plot shows data with a negative (positive) applied voltage.

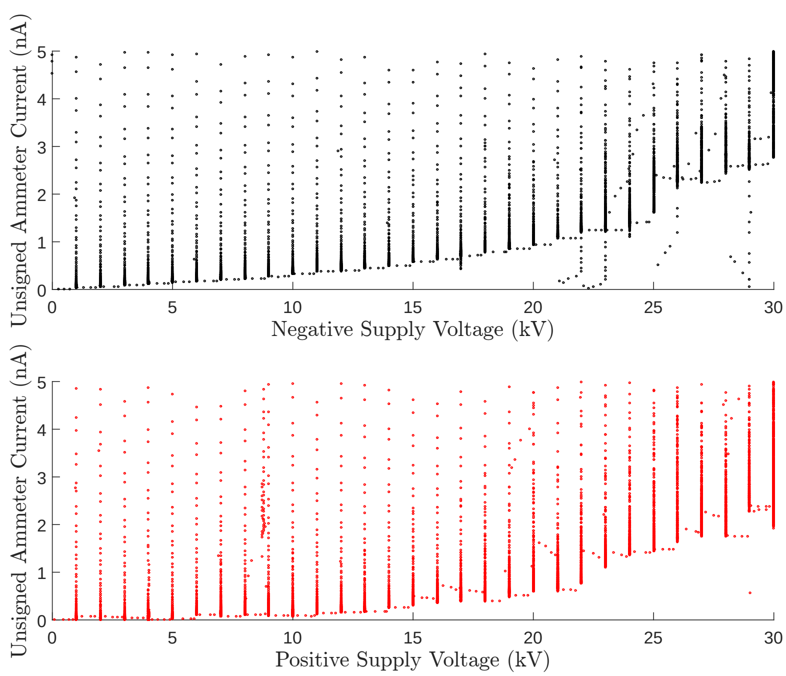

Figure 19: Current reported by the floating nanoammeter vs voltage reported by the power supply. Ammeter is in bucket of oil.

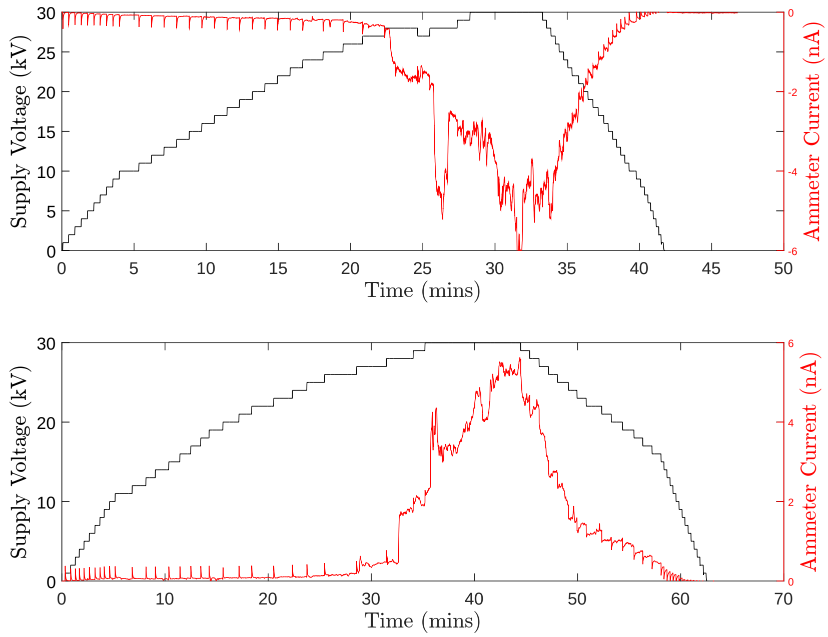

Figure 20: Current reported by the floating nanoammeter and voltage reported by the power supply as a function of time. Ammeter is in bucket of oil with output disconnected. The top (bottom) plot shows data with a negative (positive) applied voltage.

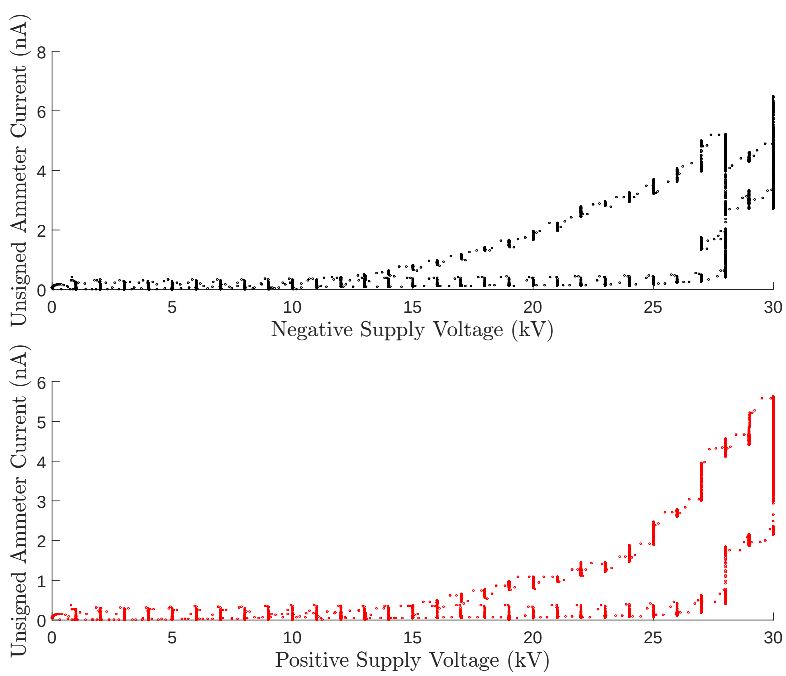

Figure 21: Current reported by the floating nanoammeter vs voltage reported by the power supply. Ammeter is in bucket of oil with output disconnected.

3.5.5 Run 5

To try and diagnose where the leakage current was coming from in the circuit we decided to bring it up to voltage in air. This way we know that we will see some kind of corona/discharge, which will highlight the weak points in the circuit. To do this, I made another pail with the appropriate feedthroughs etc. and also with an acrylic window to allow for observation with a camera during operation.

Initially, no signs of corona etc. were observed, even when charging up the circuit to 30 kV. While perhaps reassuring, this was not useful. In order to try and induce corona we moved the circuit board closer to the lid of the pail by shortening the standoffs which attach it to the lid. This was done incrementally, first shortening them to 3 in, then 2 in and finally 1 in. It was not until the standoffs were only 1 in long that we observed any corona. The corona/arc is shown in Fig. 22. The top left image shows the ammeter through the window in the pail. The top

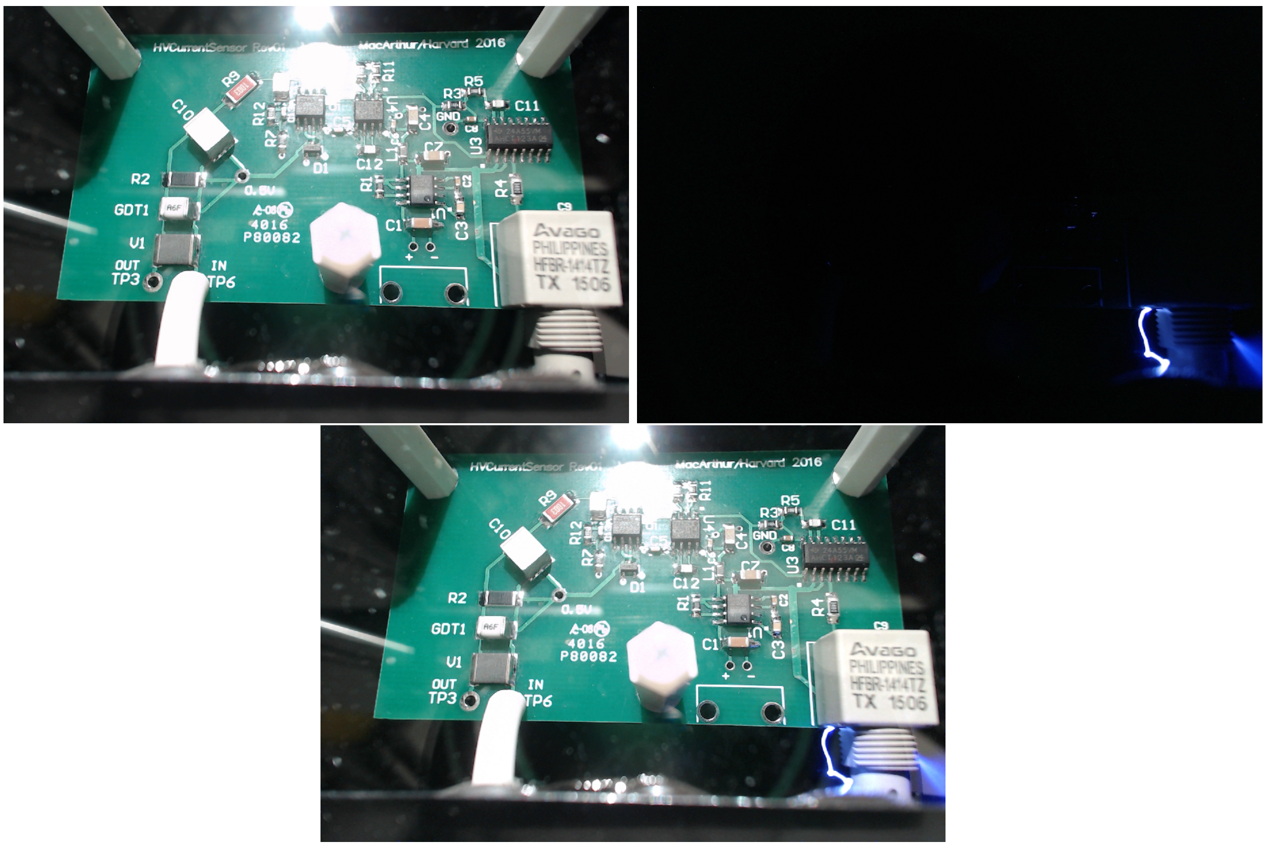

Figure 22: Pictures of the ammeter circuit showing the position of corona/arc events.

right image is a combination of two photos which contained an arc (to the left) and some corona (to the right). The bottom image is a combination of the upper images, showing the position of the discharges. The corona seemed to start occuring when the voltage reached 25 kV, and arcing occurred at around 28 kV. There was some observed hysteresis — the corona was maintained until the voltage was reduced to 20 kV. Looking at the images we see that the discharges seem to be concentrated around the optical pulse generator which is used to transmit the current reading through an optical fiber.

After observing this I closely examined the optical pulse generator to try and identify the source of the problem. The images below show what I found while dismantling this component. The first image (starting from top left) shows the component from the rear side, giving the clearest view of what we suspect the problem to be: the red circles highlight very small holes in the plastic housing (present on both the front and back). Set slightly back from the holes, there is metal. Upon disassembly it seemed that this metal performed no function, other than mechanical. As far as I could tell, there seemed to be a small piece of metal encapsulated in the housing to provide rigidity. At the holes highlighted there seemed to be a sharp edge/corner of that metal. The component has eight pins — the central four (gold coloured) connect to the LED inside (see image 2), and the outer four are there solely

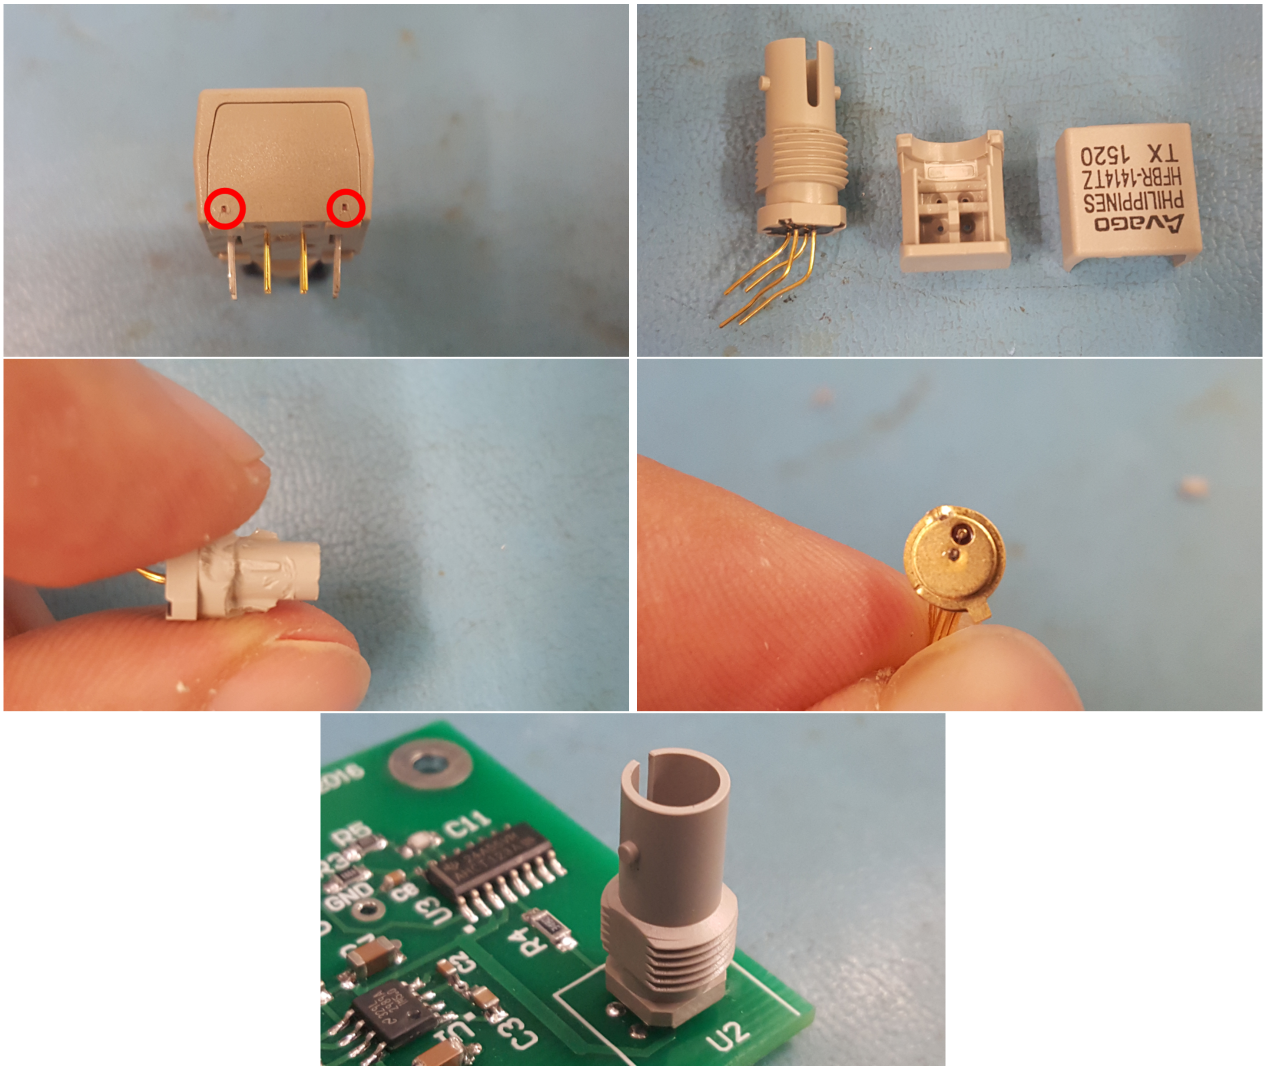

Figure 23: Pictures of the optical pulse generator being dismantled. See main text for description.

for mechanical purposes and seem to connect to the aforementioned encapsulated metal. Image 2 shows that it is thankfully relatively easy to remove part of the casing, separating the LED and ST fiber connector from the rest of the package. Image 3 is after having cut away the ST connector part of the package — the LED is encapsulated in plastic still. Image 4 shows the rest of the plastic having been cut away, leaving just the LED itself. Image 5 shows the LED and ST connector (left-hand piece in image 2) soldered onto the circuit board in place of the original package. Note the now-unoccupied pins around the component which were previous used for mechanical support (the component is still plenty rigid).

Hopefully this will have reduced or solved the problem of residual leakage current which we have been observing. We can either test again in air to see if there is any corona, or test in oil to measure the leakage current. One slight problem is that we are running out of the photovoltaic converters. Of the 4 originally purchased, one is on the circuitboard currently in oil. One is in a silicone-potted circuit board, one is on a previous iteration of the circuit, but not usable since it had the grounding pin cut off. One had a leg break when being handled (presumed faulty). Alternatively, we have 'panel mount' style photovoltaic converters, shown in Fig. 24. It is somewhat bulkier and might not fit on the board as nicely. It would also require modification due to the mis-sizing of the FC-connector.

We should bear in mind that this may not be the only problem on the ammeter circuit. In particular it is worthwhile noting that the optical pulse generator is probably the component which is closest to ground (the lid of the pail), due to the presence of the (acrylic) window used to observe the circuit.

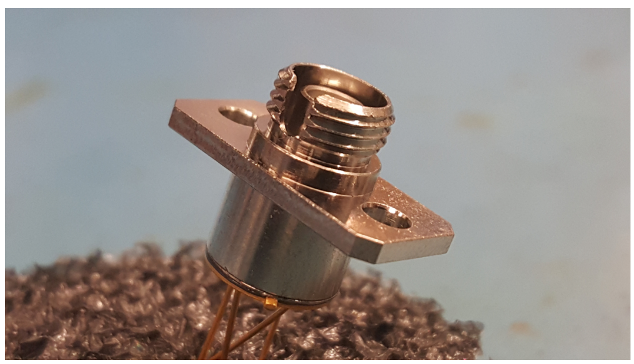

Figure 24: ‘Panel mount’ style photovoltaic converter.