

Kelvin Probe Charging Research

**A look at test mass charging problems
in the field of gravitational-wave interferometry**

**Noah Hurst
University of Glasgow, Institute for Gravitational Research
Summer 2009**

INTRODUCTION

Gravitational wave detection has proved to be one of the most difficult physics experiments ever attempted, for the simple reason that wave amplitudes are extremely small. Although very carefully engineered, interferometric wave detectors have as of yet been unable to pick up a meaningful signal due to noise in the system. In order to find these elusive waves, the interferometer must be very sensitive, but as a result it often picks up other signals which cover up the gravitational-wave signal. Most of the research being conducted in gravitational-wave astronomy at the moment is devoted to eliminating noise from the interferometer system – it is widely believed that a few more years of work could reduce noise to below the level of typical wave amplitudes passing through earth. Progress such as this would allow scientists to finally announce for certain the existence and nature of the ripples in spacetime that were originally predicted by Albert Einstein at the beginning of the 20th century.

Noise comes in a wide variety in gravitational-wave interferometers, due to their incredible sensitivity. Seismic noise occurs when the detector picks up vibrations from the environment - an earthquake, wind, or even a large truck driving past can be an issue. Thermal noise comes from thermal lattice vibrations in either the mirrors or their suspension systems. Shot noise arises from quantum uncertainty in the behavior of the laser light as it passes through the interferometer. Noise can also be introduced by imperfections in the mirror structure, or from residual gas in the imperfect vacuum system. The focus of this paper, however, is charging noise.

For the ultra-sensitive interferometer to work properly, the mirrors in the interferometer must be held completely still – any motion at all will be detected and will mask the presence of passing gravitational waves. Even the slightest force felt by the mirror can displace it enough to cause a problem. If the test mass acquires a net charge (whether positive or negative), it experiences Coulomb forcing due to attraction towards the grounded suspension system, which results in a slight displacement of the test mass and therefore introduces mechanical noise into the signal. Additionally, charge present on the mirror surface can attract tiny dust particles in the vacuum chamber and lead to absorption or scattering of laser light at the mirror surface, which shows up both as quantum shot noise and as thermal noise.

Originally it was thought that the 1st generation of interferometric gravitational-wave detectors such as LIGO, GEO, and VIRGO would be sensitive enough to pick up wave signals, but these instruments have been disappointingly unsuccessful. Disappointment turned quickly to encouragement as scientists immediately began to think up ways to improve 1st generation detectors. Within a few years, our current facilities will be upgraded to what will be known as the 2nd generation of detectors (for example, Advanced LIGO). Hopes are high that the new technology introduced in the 2nd generation will be sufficient for successful wave detection, however, this can only happen if the charging issue is effectively resolved. In addition, plans are being made for a gravitational wave detector in space, following the earth in its orbit around the sun. This device, dubbed LISA (Laser Interferometer Space Antenna), will be sensitive to waves in the lower frequency range, and will also experience charging problems.

I spent my summer of 2009 studying the charging problem with the Institute for Gravitational Research at the University of Glasgow, Scotland. I worked mainly with first year grad student Paul

Campsie, and we periodically corresponded with Professor Giles Hammond when we needed help. This paper is a summary of my activities and progress throughout the summer. I have tried to organize this paper in a more or less chronological fashion – I begin with research into past work regarding the charging problem and Kelvin probe instrumentation, then I give a brief overview of the experimental apparatus, and then proceed to cover problems I faced throughout the summer.

TEST MASS CHARGING

The charging problem in gravitational wave detection has to do with net charge accumulating on the surface of the test mass. This net charge generates an electric field and therefore a Coulomb force between the fused silica optic and the metallic suspension frame. Similarly, ambient magnetic fields can create Lorentz forces on the charged mirrors [1]. The forcing displaces the test mass in a way such that noise is created in the signal, covering up the actual gravitational wave signals that we are interested in. Another effect of charge buildup is the attraction of charged dust particles present in the vacuum system. Dust accumulates on the optic surface, where it can scatter or absorb incident laser light. Absorption results in heating in the center of the test mass, which leads to thermal noise as well as lensing. Scattered light can introduce shot noise into the system, specifically phase noise and laser power loss. Difficulties with charging also have to do with dynamic charge patterns on the test mass. For example, a sudden change in charge magnitude creates a forcing discontinuity that results in swinging of the optic. Mobile charge is associated with a fluctuating electric field, giving optic displacement noise at lower frequencies (less than 1 kHz) [2]. In order to address charging issues, the source of the charge must be clearly understood. Possible sources of charge include dust contact during pump-down, earthquake stops, First Contact protective coating, and possibly cosmic rays [2, 3]. Although earthquake stops and pump-down are not relevant for space-based detectors, LISA is expected to experience charging problems related to cosmic rays as well as interactions between housing and proof mass [1].

Earthquake stops are used in systems such as LIGO to prevent equipment damage in the event of a violent seismic event. The stops were originally made with viton (a rubber-like elastomer) tips that come into contact with the test mass during a high-magnitude earthquake. Although the viton stops protect against physical damage to the mirrors, they can deposit charge on the fused silica if contact occurs, at which point charging noise comes into play. In May 2006 a strong earthquake shook the Livingston LIGO site and resulted in a contact between one of the test masses and the viton-tipped earthquake stop. This necessitated breaking the vacuum in order to free the test mass, but once the vacuum was reinstated and the instrument became operational, it was noticed that a low-frequency noise contribution had disappeared [4]. This has been attributed to charging effects, and has spurred further research concerning the charging problem. One of the steps taken during the recent Enhanced LIGO upgrade was to replace the viton tips with fused silica tips. Because the stop and the test mass are now made of the same material, it is expected that charge transfer between the two will be minimized. However, other sources of charge still need to be addressed.

It has been suggested that cosmic ray interactions can transfer electrons to the fused silica optics from the surrounding metallic structures (vacuum equipment, test mass suspension system) because the critical energy of fused silica is higher than that of the metals [2]. Moscow State University has observed a steady rise in negative charge of 10^5 electrons per square centimeter per month on a fused silica optic sitting in a vacuum [2, 5]. This could be attributed to either the cosmic ray explanation or

perhaps geological radioactive processes [5].

In order to protect the surface of the test mass before installation in the interferometer system, it has been common practice to use a product called First Contact which is designed to serve exactly this purpose. First Contact is a polymer solution that is applied to the test mass surface as a viscous liquid which then dries and forms a protective coating over the optic and may be peeled off later when the mirror is safely installed at the interferometer site. Although this solution is very helpful for protecting the optics, it is believed that the coating leaves charge on the test mass when it is peeled off. If indeed this is true, then the optics will be significantly charged even before operation begins. First Contact solution is currently being studied by a few different research groups in an effort to understand its effect on test mass charging.

Each of these issues needs to be addressed in order to move on to the 2nd generation of detectors. Even though extreme precautions are taken to keep the vacuum system clean, there will always be some amount of dust in the system that could lead to charging during pump-down. It is expected that the fused silica earthquake stops should not be able to charge the test masses during a major earthquake event, but it may be hard to make any conclusions until such an event occurs again in the Enhanced LIGO system. Several techniques for reducing charge left by First Contact are being investigated – these include mechanical peeling of the coating in a vacuum, as well as the use of conductive First Contact coating. After all these issues have been resolved, there still exists a steady climb in charge as observed by Moscow State. Whether the culprit is cosmic rays, radioactivity, or something else, this requires a steady discharging technique. The main candidate at the moment is the use of UV illumination for photoelectric discharging, however, many people are concerned that long-term UV exposure could damage the reflective coating on the optics [2].

UV CHARGE MITIGATION

With the understanding that residual charge on the interferometer optics is inevitable, along with the idea that even a slight charge can produce significant noise, it is agreed that a method must be developed for successful mitigation of charging effects. Presently, one of the most promising charge mitigation technique is the use of UV illumination. The idea is that UV photons can liberate unwanted electrons from their dielectric trap on the fused silica optic surface through the photoelectric effect. In the case of a positively charged test mass, a control cathode could be placed in close proximity to the test mass. UV illumination would then free electrons from the control cathode, and these electrons would be electrostatically attracted to the positively charged test masses. However, it is well known that UV radiation can damage certain materials due to relatively high photon energy.

In order to successfully liberate excess electrons from the optic, UV light must be used – longer wavelength photons have insufficient energy to accomplish this task. Unfortunately, this means that the incident UV photons are also strong enough to alter chemical bonds in the silica structure itself. This is why researchers are worried that charge mitigation through UV illumination would eventually end up damaging the optic surfaces. That being said, both discharge and damage process due to UV light have associated time scales, and a thorough understanding of these rates will show whether or not UV can be the solution to the charging problem.

All other charging effects aside, charge naturally decays off the optic surface into the

surrounding environment. Theory suggests that adsorbed hydrogen ions from humidity in the air create conductivity on the surface of the optic, by which charge may be transferred across the optic surface towards ground [5]. The Moscow State research group has experimentally shown that charge decay is related to the humidity of the environment as well as the surface preparation of the optic [5]. In the lab, I measured the charge relaxation time constant to be around 2 hours, using a 2.5 cm diameter fused silica disk cleaned with isopropanol. Naturally, the time constant is much higher in a vacuum due to a lower concentration of humidity. Moscow measured the time constant to be ~2400 hours in vacuum without prior cleaning, and upwards of 12000 hours (about a year and a half) when prepared with an ultrasonic bath [5]. What now needs to be determined are the time constant for charge mitigation and corresponding optic damage rate. For example, what is the optimum intensity and wavelength of UV illumination that will give highest discharge rate and lowest damage rate? Once the optimum has been found, is the discharge rate high enough to keep charging noise out of the interferometer signal? Is the damage rate low enough to avoid problems?

Both Moscow State and Trinity University have experimented with tunable xenon lamps for UV illumination, while Stanford has tried new UV LED technology for discharging. Through experimental procedure and the use of equipment to measure charge on the surface of a fused silica sample, it has been shown that discharge rate is at a maximum around 215 nm, and that there is an approximately linear dependence between UV intensity and discharge rate [3]. Increased photon absorption by the test masses has been identified and linked to UV damage, the result being a rise in absorption of laser light by the optics. Despite uncertainty as to the damage rate, it would be advantageous to keep UV exposure at a minimum when possible, such that shot noise due to absorption is avoided. Therefore, despite the steady rise in surface charge, efforts must be made to minimize charge on the optic during preparation and installation such that potentially hazardous solutions such as UV illumination are used sparingly.

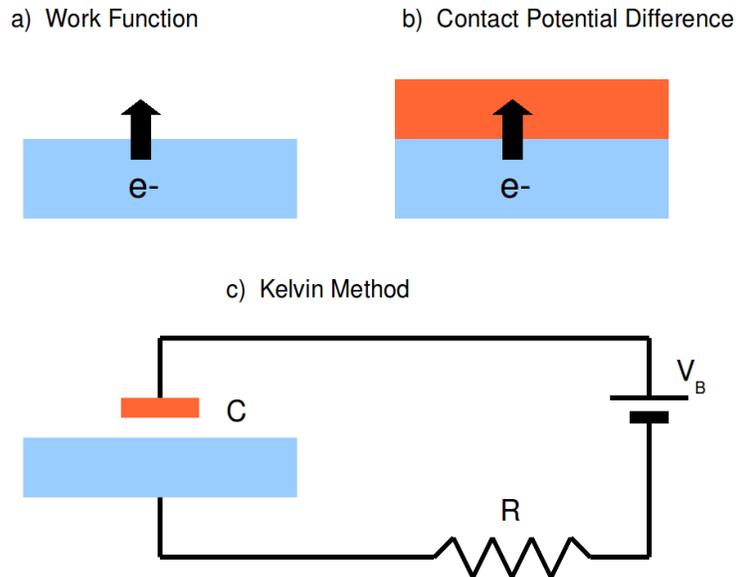
KELVIN PROBE

In order to study characteristics of charge on the fused silica surface, it helps to know not only the magnitude of charge on the surface, but also spatial variations in charge density across the surface. The best way to accomplish this is through the use of a device called the Kelvin probe. One of the concepts studied by Lord Kelvin was the contact potential difference between two solids. Essentially, this is the difference between the work functions of the two materials. When contact is made between the two materials, the substance with the higher work function will acquire electrons from the other substance, resulting in a potential difference. Therefore, if the work function of one of the materials is known, and the potential difference is measured, the work function of the unknown material may be deduced. The Kelvin method of measuring work functions is extremely versatile – it is a fairly simple concept that can be applied to a variety of different materials under a variety of different conditions. Additionally, it is very non-invasive – the sample being measured is almost entirely unaffected by the measurement process. Fig. 1 depicts specifically how measurement of the contact potential difference is accomplished.

Figure 1: Simple Schematic of the Kelvin method

a) The work function is the energy required to remove an electron from the bulk of a material. It is dependent on material characteristics as well as surface preparation. **b)** The difference in the work functions of two contacting solids is known as the contact potential difference (CPD). When connected, electrons will flow towards the solid with the higher work function, resulting in a voltage V_{CPD} across the junction.

c) The Kelvin apparatus consists of a probe surface (orange) and a sample surface acting together as a parallel-plate capacitor. Contact potential voltage V_{CPD} is generated across the circuit, and can be continuously measured. The work function of the probe is known, so the work function of the sample can be determined. Backing voltage V_B can be adjusted such that voltage across the circuit disappears, at which point V_B is equal to V_{CPD} . The probe is generally vibrated up and down, and scanned across the sample surface.



Since Lord Kelvin came up with this basic method for measuring work functions, further research has idealized the process to the point where very precise measurements are possible. In order to eliminate noise in the capacitor system, it was proposed to vibrate the Kelvin probe in a direction normal to the surface being measured such that the probe produces an AC voltage at the same frequency as the vibration. Therefore any portion of the Kelvin signal at a different frequency is noise and may be subtracted from the measurement. The AC voltage is then rectified and continuously balanced with an applied DC backing voltage V_B which can then be read off electronically for continuous monitoring of the contact potential difference. The vibrating probe technique also results in a much more stable signal than otherwise.

Error in the Kelvin probe signal is due largely to both unwanted electromagnetic pickup and stray capacitance [6]. Both sources of noise are interactions between the Kelvin probe and other equipment nearby, so the obvious way to increase signal-to-noise ratio (SNR) is to make sure the probe is very close to the specimen and far away from everything else. In fact, it has been theorized that SNR is maximized for high values of a/d , where a is the vibration amplitude and d is the distance from the probe to the sample (obviously a/d must be less than unity) [6].

Electromagnetic pickup occurs when the Kelvin probe senses noise from electric and magnetic fields in the vicinity. This could originate from other current-carrying wires or instruments nearby. Additionally, ambient magnetic fields must be kept very small, or the Kelvin probe will generate current by inductance. The best way to deal with this problem is to screen all other components from direct interaction with the Kelvin probe. Phase-sensitive detection also helps to ensure that EM pickup

is not included in the reading. However, stray capacitance is a huge source of noise that happens to possess the same frequency as the actual Kelvin signal.

The Kelvin probe is a vibrating capacitor where the parallel plates are the probe and the sample, and the insulating medium in between is a vacuum. Although the Kelvin probe is generally placed very close to the sample and farther away from other structures, it can still generate small stray capacitances. The trouble with this is that the stray capacitances between the vibrating Kelvin probe and other stationary structures will produce a signal at the same frequency as the real Kelvin probe signal coming from the sample. In addition, efforts to increase the actual signal strength, such as increasing vibration amplitude, can result in an increase in the stray signal as well. Stray capacitance effects have been shown to increase linearly with distance between probe and sample, with vibration amplitude, and with contact potential difference between probe and source of stray capacitance [6]. Therefore, it helps to keep vibration amplitude small, to keep the probe very close to the sample, and to make sure work functions of probe and other screens and/or structures in the vicinity are similar if at all possible.

The accuracy of the work function measurement of the sample is dependent on the certainty in the work function of the probe material. The work function of a material surface is often highly dependent on surface quality, cleanliness, and temperature. The general consensus is that a probe made of tungsten, gold, or tin oxide will work reasonably well at room temperature in a vacuum, providing the material is kept clean and smooth [6]. Despite all the research put into reducing noise in the Kelvin probe signal, a reasonably accurate signal is attainable without too much trouble. As long as basic precautions are taken, the Kelvin method is a very accurate way to measure work functions, especially for large, flat surfaces such as the interferometer test masses. In order to map spatial charge distributions on the optic surface, a smaller probe may be scanned across the surface of a larger sample – this aids in understanding not only the charge magnitude and spatial distribution on the optic, but also characteristics of charge mobility in the case of successive scans at different time steps.

EXPERIMENTAL SETUP

The goal of Kelvin probe research at the University of Glasgow is to study test mass charging by looking at the behavior of charge on smaller fused silica samples. The Kelvin probe is used to study charge behavior on the samples with regard to spatial charge distribution as well as temporal charge behavior. In order to deposit charge on the samples, a few different methods are used. A mock earthquake stop was designed in order to test charging effects from both viton and fused silica earthquake stops. First Contact solution is used to coat the samples, and then charge behavior is studied after the coating is peeled off. An electron gun is currently being built and installed, which will be used to deposit negative charge on the samples. Subsequently, a UV LED will be installed with which to study UV discharging of the samples. Eventually, plans are to also install a torsion balance, which would be able to measure the Coulomb force generated by charge on the sample. Once everything is in place, the Glasgow charging lab will be equipped to study almost all aspects of the charging problem currently under consideration. However, because the experiment is relatively new at Glasgow, much of the equipment is not set up yet. As a result, most of my time at Glasgow was spent working on lab equipment, building and fixing different parts of the experiment. The down side of experimental physics, as I found out this summer, is that everything has to work perfectly before scientific research can take place. I was a bit disappointed that I didn't have the opportunity to extensively study charging, but I realize that lab setup is a very important part of the experiment, and I think I learned quite a few

useful concepts related to experimental science (including patience).

For our purposes, we made an attempt to keep the Kelvin probe as sensitive as possible, because even minimal charge can be important regarding gravitational wave interferometers. In our lab setup, we used thin copper foil sheets to shield the Kelvin probe from other wires and electronic junctions in the vacuum chamber. We noticed that the Kelvin signal would jump considerably during operation of the magnetron ion gauge we use to measure lower pressures inside the vacuum chamber. A baffle plate made of aluminum and copper was installed to block EM interaction between the magnetron and the probe, and we also make sure not to operate the magnetron while taking probe readings. A reasonably large vacuum chamber was used so that the Kelvin probe parallel-plate capacitor could be placed in the center of the chamber, far away from the chamber walls and other sources of stray capacitance. Our probe was normally operated at a distance of about 1 mm and a vibration amplitude of ~ 0.5 mm. The probe consists of a small disk of clean gold mesh a few millimeters in diameter, and the samples we measured were mostly 2.5 cm diameter fused silica disks. The disks were positioned on a three dimensional translation stage consisting of three Thorlabs micro-positioning stepper motors. The step resolution of the motors was much finer than used for any of our experiments. Both the Kelvin probe and the translation stage were attached to a rigid cube-shaped frame that fit neatly inside the vacuum chamber.

When I arrived at the beginning of the summer, flanges on the vacuum chamber consisted of two viewing windows on the top and the side of the chamber, a pumping flange with a turbo pump and a roughing pump, two transducer pressure gauges, a magnetron ion gauge for measuring low pressures ($< 10^{-2}$ torr), an electrical feedthrough for the positioning stages, an electrical feedthrough for the Kelvin probe, and a sample loading chamber to minimize pump-down time. The sample loader was a small chamber closed off from the main vacuum system with a gate valve, with a manual fork for feeding samples into the chamber. During normal operation, the main chamber was held at around 10^{-5} – 10^{-4} torr. When a sample was loaded, the fork was retracted into the sample loading chamber and the gate valve was closed. Then the vacuum was broken in the loading chamber, and a new sample was placed on the fork. Then the loading chamber was sealed again, and pumped down with a roughing pump. Once the loading chamber was reasonably evacuated (so as not to significantly alter the vacuum in the main chamber), the gate valve was opened and the fork was manually pushed into the main chamber, where the new sample could be lifted off the fork with the vertical translation stage. Therefore the vacuum in the main chamber had to be broken only when a problem occurred with normal operation, for example, if the new sample was not loaded properly, or if something in the main chamber needed maintenance.

The Kelvin probe system was manufactured by a German company called Besocke Delta Phi. It consists of an electronics box which supplies a signal to the modulator in order to vibrate the Kelvin probe, and also receives and processes the probe signal. We fed the probe signal into an analog oscilloscope so as to visually analyze the signal, and into a lock-in amplifier by Stanford Research Systems for reading the amplitude and phase of the sine wave generated by the Kelvin probe. A LabVIEW program was used to perform probe scans across the sample surface and read the amplitude and phase from the lock-in amplifier at each point on the scan. From there it was a simple matter of generating a surface plot of the electric potential present on the silica sample in order to analyze charge behavior. Successive scans could also be put together into a short movie of charge evolution over a certain time period. The entire lab setup is shown in Fig. 2.

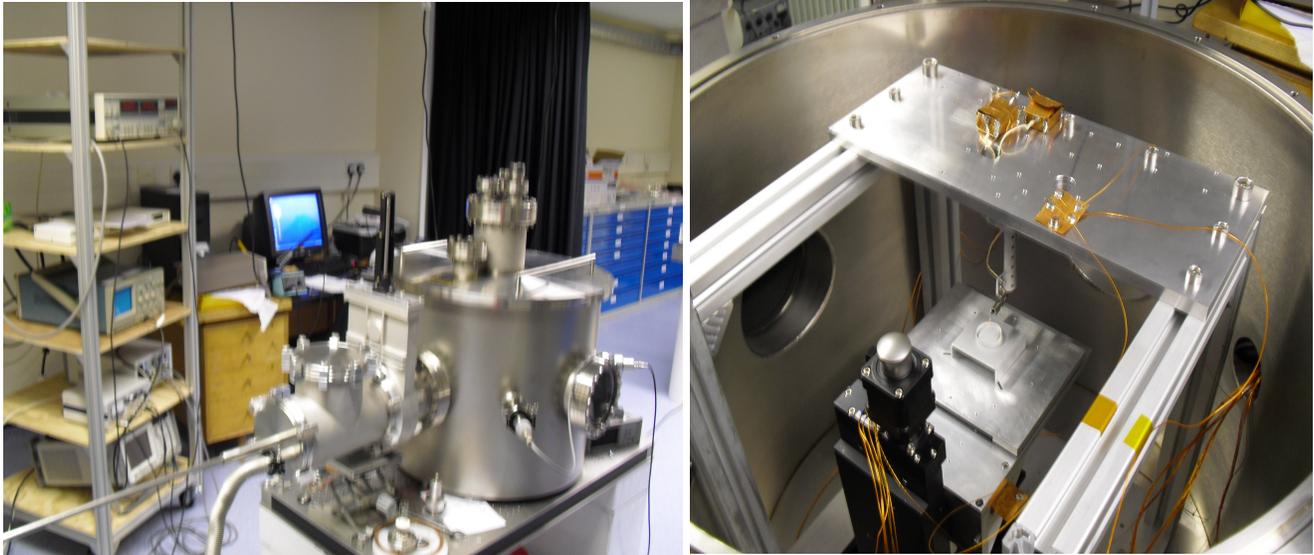


Figure 2: Lab Setup – The left frame shows the apparatus in its entirety. The stack of instruments on the far left, from top to bottom, are the lock-in amplifier, a few unused LabVIEW instrument connectors, the analog oscilloscope, the Kelvin probe electronics box and the stepper motor driver box, and finally an unused signal analyzer. The computer in the background controls the whole setup. In the foreground, the nearest flange on the vacuum chamber is the sample loading chamber with the extended loading fork and gate valve, a few pressure gauges on the right side, and the electron gun feedthrough on top. On the unseen side of the vacuum chamber are the rough and turbo pumps, as well as a viewing window. The right frame shows the contents of the vacuum chamber. The Kelvin probe is shown in close proximity to a small circular fused silica sample in the middle of the chamber. The sample is sitting on a 3-way translation stage – the z-direction stage is visible, while the x and y motors used for 2 dimensional scanning are located at the bottom of the tank. A heavy, rigid cube-like frame supports everything.

After being introduced to the lab setup, I began working simultaneously on building the electron gun and studying First Contact charging characteristics. In most cases, we performed ten subsequent ten by ten millimeter square scans, which takes about an hour and a half, so we had time to work on the electron gun while the scans were taking place. However, soon after I began work, one of the stepper motors failed. This put a temporary halt to our Kelvin probe usage, in which time we had to determine the reason for the motor failure and redesign the scanning process such that the replacement motor would not also fail. We spent quite a while on this, because Thorlabs took their time sending us our replacement motor. During this time we also continued work on the electron gun. Once the motor was returned and installed, we were able to continue First Contact research for a short time before I left, as well as continue work on the electron gun. For the sake of organization, I will cover the electron gun, the broken motor, and the First Contact separately.

ELECTRON GUN

Commercially available electron guns tend to be very expensive. Our creative solution to this problem was to use the cathode ray tube (CRT) from an old analog oscilloscope. We had more than enough old oscilloscopes sitting around the department, so we chose a Scopex 4D10 analog

oscilloscope from 1970 that we found in the lab. After taking off the casing and removing the CRT, we began studying how the electronics were laid out. The data sheet for the Mullard D10-160GH CRT was very helpful – it explained what each wire was used for, as well as what voltage and current each wire required. Connections from the oscilloscope to the CRT included 4 wires for 2 – dimensional electrostatic deflection of the electron beam, 3 wires for the cathode heater, a wire for the cathode, a wire for the accelerator electrodes, and a wire for the focusing electrode.

We started by measuring voltages carried by each wire (being very careful with the high-voltage wires) in order to understand how to control the gun. We were able to coarsely control the deflection of the electron beam with an external two-channel 0-30 V power supply. The idea was that we would be able to spatially control the electron beam in order to deposit different patterns of charge on the fused silica samples. Upon further thought, it became obvious that it would be easier to keep the electron beam stationary and move the samples with the translation stages, which have sub-micrometer precision as opposed to the inaccuracy of the deflector electrodes. We opted to snip the four deflector wires, which, as expected, resulted in a stationary beam right down the center of the gun.

Regarding interferometer noise, concerns are associated with very small amounts of charge, while the electron gun passes a fairly sizable current, so it was clear that we would also have to design some type of shutter for the electron beam in order to limit charge deposited on the silica sample. Among our ideas for an electron shutter were a mechanical solenoid-operated shutter, an electrostatic electron catcher, or a high-voltage switch in the accelerator line. The solenoid shutter would be a grounded metal plate attached to a solenoid armature, which could be controlled externally through a vacuum electrical feedthrough. We initially had problems getting the solenoid to behave precisely and reliably. It was difficult to find a good place to mount the solenoid in the vacuum chamber – we wanted to build a spring-loaded system such that the solenoid would only be operated briefly (for a fraction of a second) to deposit charge, after which the spring would return the grounded plate into the beam path. Our troubles consisted of both limited space in the vacuum chamber and the fact that the solenoid was a bit weak and the armature displacement was short. We turned to other solutions in the meantime. The high-voltage switch was ruled out because although we could cut the accelerator voltage rather efficiently, the hot cathode would continue to emit electrons thermally until the heater had a chance to cool down, and we were looking for a sharp cutoff in beam current.

I began to consider the usage of an additional electrode with which to collect electrons on their path towards the sample. I figured that a positively charged conducting ring closely encircling the beam would be able to catch all electrons and stop current flow depending on geometry, electrode potential, and electron speed. I performed calculations in order to find out whether this solution could potentially work. I wrote a MATLAB program to solve Laplace's equation and determine the electrostatic potential due to a positively charged ring conductor. Using spherical coordinates with azimuthal symmetry, I got a nice-looking surface map of potential surrounding the ring. My plan was to use Lagrangian mechanics to calculate electron dynamics, so that I could find a suitable geometry to ensure that the electrons would be collected. Around this time, it occurred to me that it might be easier to use the spatial deflector electrodes already in place as a more efficient electron catcher. Although my calculations did not turn out to be useful, I thought it was a good practice problem in electrostatics, and could turn out to be useful experience at some point throughout my career. At that point in time, other issues came up, and as of yet we have not gotten back to working on the shutter. This was partially because we had other problems with the electron gun itself, and we figured a shutter was

pointless if the actual gun didn't work.

After we obtained a good understanding of how the oscilloscope electronics were set up, we determined which CRT connection wires were necessary, and we soldered these onto the vacuum feedthrough. We used two feedthroughs at first – one had twenty pins, but a voltage rating of around 500 V. Because the accelerator electrodes use 1500 V, we were forced to also use an additional coaxial feedthrough in order to handle the high voltage wire. The next step was to take apart the glass enclosure for the gun. Assistance was provided by the glassblower in the chemistry department. He used a hot wire to score a circular cut in the tube. Despite concerns that the huge pressure difference between the inside and outside of the CRT would result in a shattered tube, the cut worked out very nicely, and we now had the electron gun and the phosphorous-coated oscilloscope screen separated. We placed the phosphorous screen inside the vacuum chamber on one of the translation stages so that we would be able to see the electron beam if the gun worked. Unfortunately, the gun did not work.

Initially we had problems with the coaxial feedthrough. It had been welded together by the machine shop at Glasgow, and had a small leak that we were able to locate by dripping isopropanol around the joints. With the faulty feedthrough, the chamber wouldn't pump down past about 5 torr. After smearing vacuum grease around the leak, we were able to pump the chamber down to 0.5 torr. We knew that the electron gun had to be operated in a substantial vacuum – generally the oscilloscope CRT's are operated at around 10^{-5} torr, because any gas atoms in the way of the beam will be ionized and will return to the cathode and cause damage. We figured that for a simple test fire, we could use the 0.5 torr vacuum without much trouble. When the oscilloscope was turned on, we immediately noticed arcing inside the oscilloscope box, and shortly thereafter, the fuse blew. After replacing the fuse with a slightly stronger one, the same arcing pattern persisted. It seemed that current was jumping over two small diodes in the center of the box, near the wires leading to the gun.

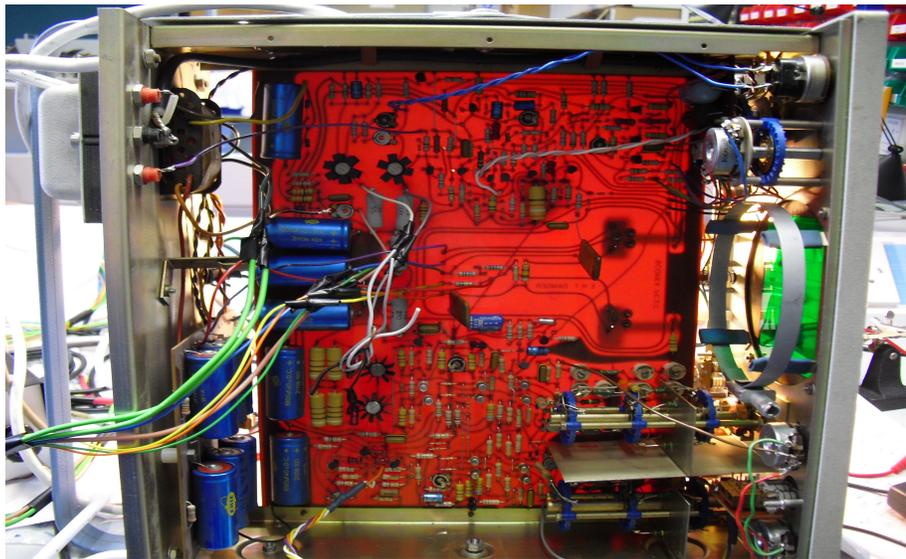


Figure 3: Scopes 4D10 Oscilloscope Box – This picture was taken while we were tracking the circuitry and trying to figure out why the diodes had been fried. The circuit board is lit up from behind in order to show the connections on the back of the board. The wires leading out of the box to the left are the connections for the CRT. The diodes were removed before this picture was taken, but they were located right at the center of the board directly left from the screen.

Our first thought was that the diodes were trying to prevent current from flowing the wrong way, which could mean that a couple wires had been hooked up backwards. However, we carefully checked all the wires, and they were attached exactly as they should be. Our next thought was that we had altered the circuit by cutting off the end of the tube, which had a conductive coating attached to the 1500 V accelerator line. Then we measured across the questionable diodes and found that they exhibited infinite resistance in both directions, meaning that they had been fried. While in the process of trying to replace the diodes, and after a significant amount of time spent trying to understand the circuitry in the forty year old oscilloscope, we opted to give up and supply our own power to the electron gun. We found two high-voltage power supplies sitting around the department, which we could use for the accelerator voltage and the focus electrode. We had a smaller 0 – 30 V power supply which we decided to use for the cathode heater. The cathode was connected to ground with the understanding that a large positive accelerator voltage in combination with cathode heating would be sufficient to draw electrons from ground into the beam. We found a brand new 4-way coaxial feedthrough that we could use for the electron gun – this would allow us to achieve a reasonable vacuum, and we had managed to reduce the number of wires that needed to be fed through to the gun to three. This setup was tested, and we did not observe an electron beam of any sort. We found no clues as to why this could be, but further work will hopefully result in a working electron gun.

BROKEN STEPPER MOTOR

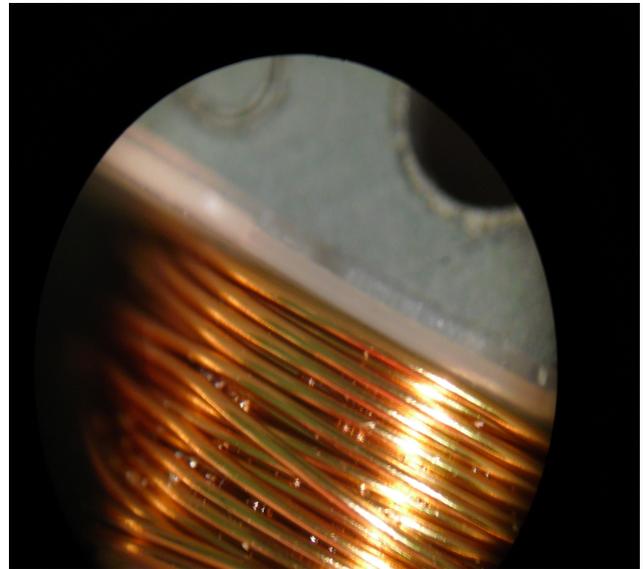
After I had just arrived in Glasgow, and I was being trained to use the Kelvin probe apparatus, one of the things that was mentioned was that the x-axis translation stepper motor had been making strange noises recently. After a few weeks of work, the motor died altogether. The motor began to make strange noises during a scan, and subsequently stopped moving. At first, it would turn slightly and then fall back into position as if too weak to turn the armature, but after a short amount of this behavior, it quit working completely.

Initially, I figured the problem could be either some type of mechanical blockage preventing the motor from turning, or something could be wrong with the motor itself. We broke the vacuum and examined – nothing seemed to be blocking the armature from moving, and the armature could be manually turned without any resistance. The next step was to examine the motor. We took the entire translation stage assembly apart and took the faulty stage out of the vacuum chamber. We were able to remove the stepper motor from the stage assembly without too much trouble. At first glance, nothing seemed to be wrong. The coils all seemed to be intact, and none of the connecting wires were loose. The stepper motor consists of 8 coils, half of which are connected to the phase A current line, the other half of which are phase B. The coils are arranged circularly, and staggered (A, B, A, B, etc). Out of curiosity, we measured resistance across each phase circuit, and found the resistances to be around 3 ohms, whereas the other two working motors showed a resistance of over 5 ohms – good evidence that some type of shorting had occurred. Again, out of curiosity, we had a look at the coils under a microscope, which revealed another interesting fact – we noticed small bubbles all over the wires, shown in Fig. 4. The bubbles were just small enough to evade regular vision, but the microscope made it quite obvious that something had happened. I popped one of the bubbles with a small screwdriver, and noticed that the material was weak and flaky.

Magnetic windings often use a coating material on the wires called Formvar. Formvar is an

enamel coating with a high dielectric breakdown strength such that wires in magnetic windings may be packed in very close to one another without shorting. After learning this, we were immediately suspicious that the bubbles were actually melted Formvar. The flaky nature of the bubbles was consistent with characteristics of enamel-like materials. This would explain why the resistance we measured was lower – if the enamel coating had melted, the winding wires could be in contact with one another, which would short the circuit. Additionally, the motors were not rated as vacuum compatible. Vacuum compatible motors are much more expensive, and are not always necessary. The idea is that a motor in vacuum cannot easily give heat off to its environment, due to the lack of surrounding air. Vacuum compatible motors are properly heat-sinked such that the motor does not overheat and damage itself. Regular motors are often successfully operated in vacuum, so long as the loading or usage is not too great, however, we were convinced that our x-direction motor had been overworked. At this point, we had to perform further research in an attempt to prove our theory, and we also had to determine the source of the problem and correct the procedure such that the new motor (which was pretty expensive) would not also fail.

Figure 4: Formvar Bubbles – This is a slightly magnified view through the microscope of one of the 8 coils in the broken stepper motor. Note the small bubbles all over the surface of the wires. We are fairly certain that these bubbles are made of Formvar that melted when the motor overheated. This resulted in an electrical short that rendered the motor useless.



The LabVIEW program used for Kelvin probe scans was written by a summer student the summer before I was in Glasgow. The scan was designed to perform a rectangular scan of user-specified dimensions. The program would step the x-motor along a number of points in the x direction, then return the x motor back to its starting point, step the y motor down once, and repeat the process for however many y steps were required. At first, nobody had any objections to this process, but now that we knew the x motor had been overworked, it seemed like a poorly thought out procedure. Although the program had been relatively simple to write, it relied far more on the x motor than on the y motor. For a ten-by-ten step square scan, the x motor traveled 200 steps – ten steps down and ten back, repeated ten times – while the y motor traveled only 10 steps total. No wonder the x motor had failed!

The first idea was to modify the original LabVIEW program such that the motors would perform an S – shaped scan. I did this without too much trouble, but I realized that the x motor was still working ten times as hard as the y motor, and I saw a better solution to the problem. I proposed to write a new LabVIEW program in which the scan would be performed in a spiral pattern, such that the x and y motors would perform an equal number of steps, and each would have an opportunity to cool down while the other was stepping. It was not an easy program to write – especially because I had very little experience with LabVIEW. I began by writing a MATLAB program to plot a square spiral pattern with a given number of steps. Once I had successfully found a way to accomplish this programming task in a text-based manner using MATLAB, I took my code to the lab and began learning how to use

LabVIEW. It was relatively easy to learn, and without too much trouble, I was able to create a working spiral scanning procedure that could read the Kelvin probe and export data with which I could then use MATLAB to create a surface plot. The task now was to prove that the old scanning procedure had overheated the x motor, and that my new spiral program would be able to perform the scan without the motors reaching dangerous temperatures.

We did not want to take apart either of the working motors for fear that we would break something, and we were pretty sure that although the x motor had an electrical short, we could use it to measure temperature for each scanning procedure. We placed the broken motor in the vacuum chamber, and fed the four phase wires through to the motor driver. We tucked a platinum resistance thermometer between the coils, and smeared the area with heat-sink putty. The leads from the resistance thermometer were fed through to a Keithley multimeter, where we could measure resistance and deduce temperature. We pumped out the chamber and began to run tests. Temperature of the broken x motor was monitored while each scan was performed – the original scan, the S – shaped scan, and the square spiral scan. Additionally, we measured the heat dissipation time constant so that we could predict how fast the motor would cool off when given a specified amount of time to relax. Results of this experiment are given in Figs. 5 and 6. I was able to show that the S – shaped scan was better for the motor than the original scan, and that the spiral scan was better yet. Thorlabs recommends that the motors be operated between 5 and 40 degrees Celsius for maximum lifetime. When sitting still, the motors apply a holding current to keep the armature from moving – this current holds the motor temperature around 38 degrees. The original scan heated the motor to over 55 degrees on a regular basis, and undoubtedly higher for longer scans. Using data from the heating measurements as well as the heat dissipation time constant from Fig. 6, I calculated that my spiral scan, with a 60 second break between scans, could keep the motor around 44 degrees. We agreed that this was a reasonable temperature to ensure the safety of the motors.

Figure 5: Motor Temperature Experiment - The original scan, the S-shaped scan, and the spiral scan were all performed 10 times in succession while temperature was recorded. Each scan was an 11 by 11 square with 1 mm steps every 3.2 seconds. The original and S scans incorporated a ten second delay between scans, which can be seen as a sharp drop in temperature on the graph. The spiral scan shown here does not have delays between scans. Since this data was taken, I have programmed a scan delay into the spiral scan procedure.

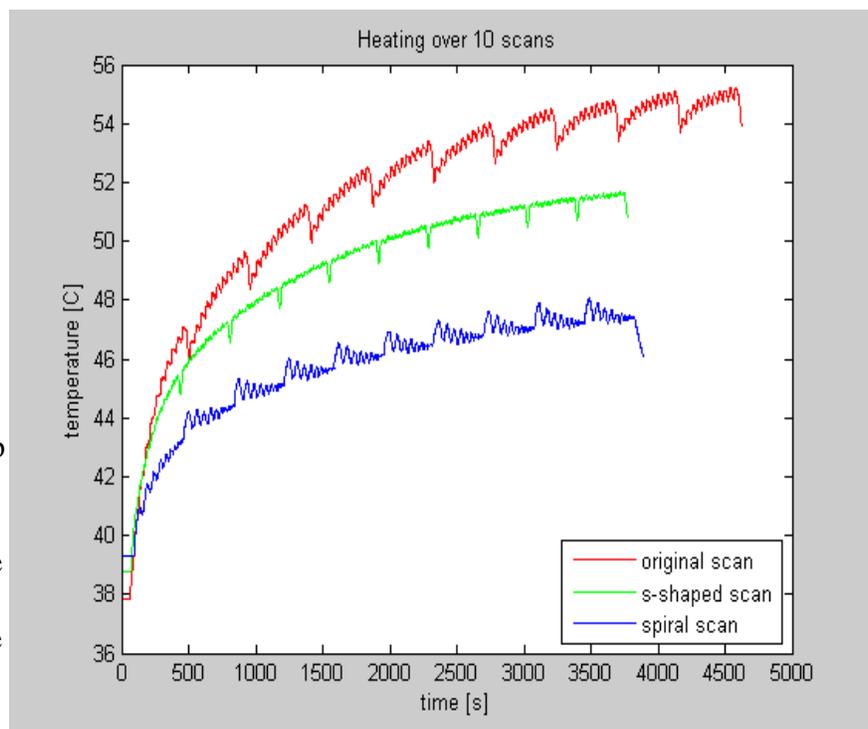
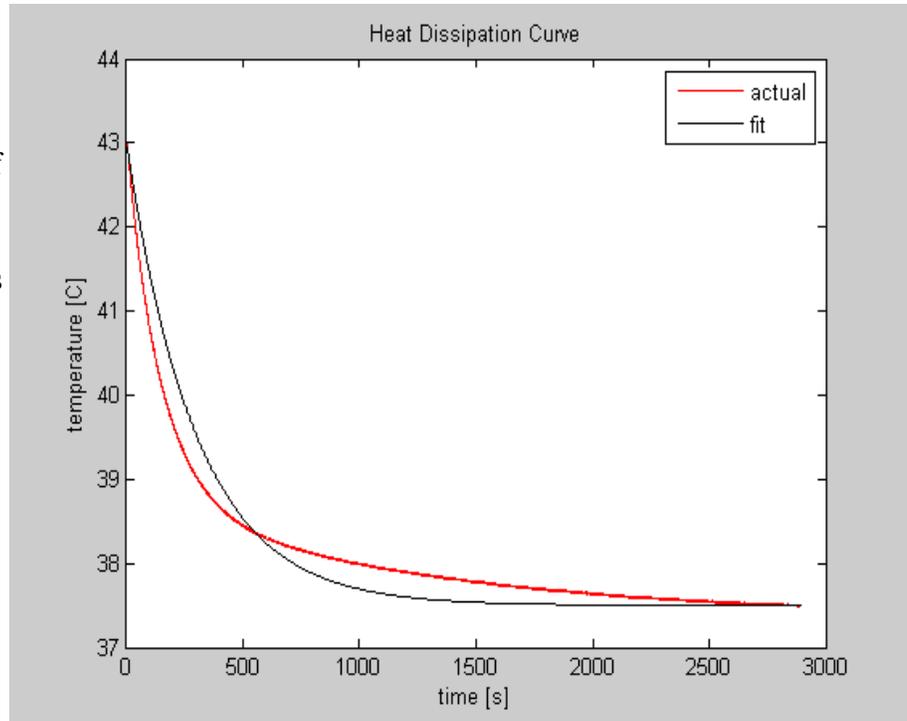


Figure 6: Heat Dissipation Rate of Thorlabs Stepper Motor in a Vacuum

Starting at 43 degrees Celsius, I monitored the temperature of the stepper motor as it cooled down in a $\sim 10^{-4}$ torr vacuum. The heat dissipation curve was not exactly exponential, but could be closely approximated as such. The time constant of the fit shown here is 300 seconds, however, the decay rate is a bit faster early on ($\tau \sim 250$ seconds), and then slower later on ($\tau \sim 350$ seconds).



FIRST CONTACT RESEARCH

As I discussed earlier, my time spent on actual charging research was very limited as a result of problems with equipment. The short amount of time I had to use the Kelvin probe was spent studying charge left on the fused silica samples as the First Contact coating was peeled off. Currently in the LIGO systems, the test masses are installed with First Contact coating in place, and peeled manually before the chamber is evacuated. We were interested in determining whether it would be advantageous to peel the coating off mechanically in a vacuum, and whether the charging was dependent on the speed at which the coating was peeled. Additionally, the company that makes First Contact sells a few different varieties, and we were curious as to the differences in charging between the different types.

The loading fork provided us with a reasonable method of peeling First Contact. Before pumping down the sample loading chamber, we taped one corner of the coating to the loading fork. When the sample was loaded onto the scanning stage, we held the fork still and retracted the sample with the x-direction stepper motor. As the stage moved away from the loading fork, the tape became taut and peeled off the First Contact fairly reliably. Plus, we could choose the speed at which the stepper motor moved the stage away, and therefore the speed of the peel. Peeling speeds between 0.5 mm/s and 4 mm/s did not seem to be correlated with charging (see Fig. 7), however, extremely slow peeling speeds (like 0.05 mm/s) and faster peeling speeds (performed manually with the fork because the stepper motor's max speed is 4 mm/s) both seemed to deposit a significantly higher amount of charge. It should be noted that we did not take a sufficient amount of data to come to any definite conclusions – the research presented here is very speculative, and only based on a few measurements.

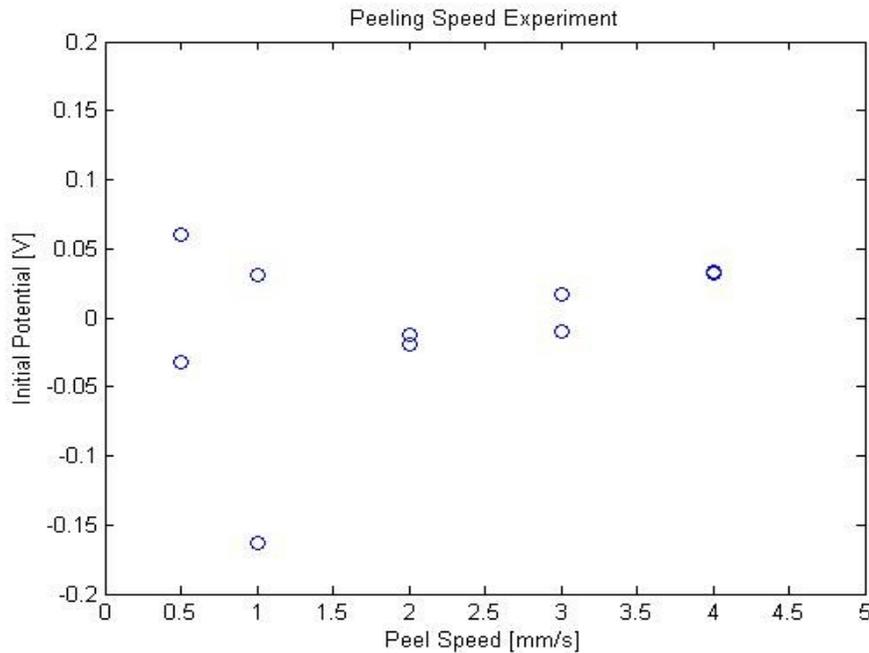


Figure 7: Peeling Speed – A few data points suggest no significant correlation between First Contact peeling speed and silica charging. However, more data is necessary in order to make a definitive conclusion.

First Contact comes in a few different varieties. Initially, First Contact was a clear polymer solution, but eventually, red dye was added in order to make the coating more visible. Surprisingly, it was noticed that the red coating seemed to leave less charge than the clear stuff, although it is completely unclear why, and this is also just an observation and not supported by extensive research. It is suspected that a conductive coating could potentially leave less charge on the optics, and for this reason, we ordered a vial of First Contact mixed with carbon nanotubes. The company obligingly sent us the requested solution, but did not specify the density of carbon tubes in the solution. We did not immediately see any correlation between charging and nanotubes, but in order to be sure, we have put in an order for a few different vials consisting of varying densities of carbon nanotubes, which should help relate polymer conductivity to charging effects.

CONCLUSION

Unfortunately, the small amount of scientific research I was able to perform this summer ended with very unclear, inconclusive results. I was able to help the charging experiment run a bit smoother, and I made progress on some of the lab equipment, but I was not able to make any strong conclusions regarding First Contact and charging. I will be interested to see what happens with this research eventually, and I have asked Paul and Giles to keep me posted regarding research results. I will also be interested to see what insights can be gained by using the electron gun, the UV LED, and eventually the torsion balance. With any luck, the Glasgow group will be able to assist in making Advanced LIGO a reality in the next few years, and with even more luck, the 2nd generation of wave detectors will be able to unlock the holy grail of gravitational physics by finally picking up a wave signal.

REFERENCES

- [1] Ke-Xun Sun et al, "LED deep UV source for charge management of gravitational reference sensors" *Classical and Quantum Gravity* Vol. 23 2006 S141-S150
- [2] D. Ugolini et al, "Charging Issues in LIGO" *30th International Cosmic Ray Conference* 2007
- [3] D. Ugolini et al, "Three years of charging research at Trinity University" MIT LIGO Seminar 2008
- [4] R. Amin, "Brief of the May 2006 LIGO Charging Event" LSC/VIRGO Charging Workshop, MIT 2007
- [5] V. P. Mitrofanov, L. G. Prokhorov, "LIGO charging research at Moscow University" 2007 Workshop on Charging Issues in Experimental Gravity
- [6] N. A. Surplice, R. J. D'Arcy, "A critique of the Kelvin method of measuring work functions" *Journal of Physics E: Scientific Instruments* Vol. 3 1970 477-482

Other helpful documents, not cited above:

R. J. D'Arcy, N. A. Surplice, "The effects of stray capacitance on the Kelvin method for measuring contact potential difference" *Journal of Physics D: Applied Physics* Vol. 3 1970, 482-488

V. Mitrofanov, "Overview of Charging Research" LSC/VIRGO Meeting, Caltech 2008

J. S. W. de Boer, et al, "Analysis and Improvement of the Kelvin Method for Measuring Differences in Work Function" *Review of Scientific Instruments* Vol. 44 No. 8 August 1973 1003-1008

I. D. Baikie, et al, "Analysis of stray capacitance in the Kelvin method" *Review of Scientific Instruments* Vol. 62 No. 3 March 1991 725-735