



University
of Glasgow

Campsie, Paul (2012) *Investigations of charging noise in future gravitational wave detectors*.

PhD thesis

<http://theses.gla.ac.uk/4055/>

Copyright and moral rights for this thesis are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Investigations of Charging Noise in Future Gravitational Wave Detectors



University
of Glasgow

Paul Campsie, B.Sc.

School of Physics and Astronomy
College of Science and Engineering
University of Glasgow

Presented as a Thesis for the Degree of

Doctor of Philosophy

Supervised by Dr. Giles Hammond and Prof. Sheila Rowan

2012

This thesis is my own composition except where indicated in the text. No part of this thesis has been submitted elsewhere for any other degree or qualification.

Copyright 2013, Paul Campsie

March 1, 2013

External examiner : Dr. Terry Quinn

Internal examiner : Dr. Bjoern Seitz

Convenor : Dr. Johannes Courtial

I would like to take this opportunity to thank my examiners, whose insightful comments improved the work contained within this thesis.

To the memory of my grandfather, David Carmichael.

Contents

Contents	iii
List of Figures	ix
List of Tables	xx
Acknowledgements	xxii
Preface	xxv
Summary	xxix
1 The Nature and Detection of Gravitational Waves	1
1.1 Introduction	1
1.2 The Nature of Gravitational Waves	2
1.3 Sources of Gravitational Waves	6
1.3.1 Burst Sources	6
1.3.2 Continuous Sources	7
1.3.3 Stochastic Sources	8
1.4 Gravitational Wave Detection	10
1.4.1 Laser Interferometry	10
1.4.2 Fabry-Pérot Cavities	14
1.4.3 Delay Line Interferometry	15
1.4.4 Power Recycling	16
1.4.5 Signal Recycling	17

1.4.6	Current Detectors	18
1.4.7	Second Generation Detectors	21
1.4.8	Third Generation and Future Detectors	23
1.5	Limiting Noise Sources of Ground Based Detectors	25
1.5.1	Gravity Gradient Noise	26
1.5.2	Seismic Noise	27
1.5.3	Thermal Noise	28
1.5.4	Photon Shot Noise	29
1.5.5	Radiation Pressure Noise	30
1.5.6	The Standard Quantum Limit	30
1.6	Conclusions	31
2	Charging Noise in Gravitational Wave Detectors	34
2.1	Introduction	34
2.2	Contact Electrification Between Different Materials	36
2.2.1	Contact Electrification Between Metals	36
2.2.2	Contact Electrification Between Metals and Insulators	38
2.2.3	Contact Electrification Between Insulators	40
2.2.4	Complications of Contact Electrification	42
2.2.4.1	Ion Transfer	42
2.2.4.2	Material Transfer	43
2.2.4.3	Other Contributing Factors	44
2.3	Charging Noise	45
2.4	Modeling Charging Noise	47
2.4.1	Defining Surfaces for the Model	47
2.4.2	Calculating Noise due to the Movement of Charge on the Optic Surface	48
2.4.3	Calculating Force Noise from a Moving, Charged, Earthquake Stop	50
2.4.4	Calculating the Total Charging Noise	52
2.4.5	An Example of Modeled Charging Noise	53

2.5	Sources of Charging in Gravitational Wave Detectors	55
2.5.1	Evacuation of the Vacuum Chamber	55
2.5.2	Contact with Earthquake Stops	56
2.5.3	Cleaning Procedure	58
2.5.4	Cosmic Rays	59
2.6	Possible Surface Charge Mitigation Techniques	61
2.6.1	UV Illumination	61
2.6.2	Conductive Coatings	61
2.6.3	Alternating Low-Energy Electron and Ion Beams	62
2.6.4	Ionised Gas	62
2.7	Conclusion	64
3	The Kelvin Probe	66
3.1	Introduction	66
3.2	Experimental Set Up	67
3.3	Working Principles of the Kelvin Probe	71
3.4	Method of Calibration	78
3.5	Noise Sources	83
3.5.1	Stray Capacitance	83
3.5.2	Leakage Current	85
3.6	Conclusion	86
4	Mitigating Surface Charge Using an AC Glow Discharge	87
4.1	Electric Glow Discharges	88
4.1.1	Introduction to Gas Discharges	88
4.1.2	The Glow Discharge Region	89
4.1.2.1	Aston Dark Space	90
4.1.2.2	Cathode Glow	91
4.1.2.3	Cathode Dark Space	91
4.1.2.4	Negative Glow	91
4.1.2.5	Faraday Dark Space	91
4.1.2.6	Positive Column	92

4.1.2.7	Anode Glow	92
4.1.2.8	Anode Dark Space	92
4.1.3	Paschen’s Law	93
4.1.4	AC Glow Discharges	94
4.2	Experimental Set Up	96
4.2.1	PID Servo Control System	100
4.2.1.1	The P, I and D Action	100
4.3	Results	102
4.3.1	Surface Charge Measurements	104
4.3.2	Transmission Measurements	110
4.3.2.1	Student’s t-test	113
4.3.2.2	Kolmogorov-Smirnov test	114
4.4	Discussion	116
4.5	Conclusion	119
5	Reducing Surface Charge During the Cleaning of Future Gravitational Wave Detector Test Masses	121
5.1	Introduction	121
5.2	SEM Images of Silica Discs After Being Cleaned with First Contact	123
5.3	Resistivity Measurements of Different Types of First Contact	125
5.3.1	The First Contact Samples	125
5.3.2	DC Resistivity Measurements	128
5.3.2.1	Experimental Set Up	129
5.3.2.2	Results	132
5.3.3	AC Resistivity Measurements	134
5.3.3.1	Experimental Set Up	134
5.3.3.2	Results	136
5.3.4	Surface Charge Density Measurements	138
5.4	Mitigating Surface Charge Using an AC Corona Discharge	140
5.4.1	A Corona Discharge	140
5.4.2	Experimental Set Up	141

5.4.3	Results	142
5.4.3.1	Surface Charge Measurements	142
5.4.3.2	Transmission Measurements	147
5.5	Discussion	150
5.6	Conclusion	151
6	Developing a Torsion Balance for Measuring Charging Noise	152
6.1	Introduction	152
6.2	The Torsion Balance	153
6.3	Experimental Set Up	155
6.3.1	The Vacuum Tank and Inner Structure	155
6.3.2	Torsion Balance Design	157
6.3.3	Capacitive Sensor Design	159
6.3.4	PID Servo Control System	161
6.4	Calibrating the Capacitive Sensor	163
6.5	Initial Results	166
6.5.1	Detecting Surface and Polarisation Charges on a Silica Disc	166
6.6	Improving the Sensitivity of the Torsion Balance and Locating Noise Sources	168
6.6.1	Decreasing the Fibre Diameter	169
6.6.2	Changing the Sensor/Acuator Configuration	169
6.6.3	Shielding the Bob and Sensor Plates	169
6.6.4	Contact Potentials	170
6.6.5	Natural Unwinding of the Fibre	173
6.6.6	Temperature Drift	174
6.6.7	Tilt	175
6.6.8	Magnetic Field Fluctuations	177
6.7	Using Multiple Regression Fitting to Remove Known Noise Sources	178
6.8	Sensitivity of the Torsion Balance After Improving the Set up	181
6.9	Torsion Balance - Mark II	182
6.9.1	The Torsion Bob	183

6.9.2	The Sensor Housing	184
6.10	Conclusion	185
7	Measuring Charging Noise with a Torsion Balance	187
7.1	Introduction	187
7.2	Experimental Set Up	188
7.3	Calibrating the Servo Control	191
7.4	Results	201
7.4.1	Estimating the Expected Charging Noise Level	204
7.4.1.1	Estimating the DC Torque Exerted on the Torsion Bob	205
7.4.1.2	Estimating the Charging Noise Correlation Time . .	205
7.4.1.3	Modeling Torque Noise Due to Charging	206
7.5	Discussion	209
7.6	Conclusion and Future Work	211
8	Conclusions	213
A	Derivation of the Charging Noise Equation for an Exponentially Decaying Surface Charge	217
B	Additional SEM Images	219
	Bibliography	224

List of Figures

1.1	An illustration of the effect of the two polarisation states of a gravitational wave on a ring of free particles.	5
1.2	This diagram highlights the orthogonality of the plus (left) and cross (right) polarisation states of a gravitational wave.	5
1.3	A plot of the change in the orbital phase in the binary pulsar system PSR B1913+16 over time.	9
1.4	A schematic of a basic interferometer set up.	11
1.5	An illustration of the effect of a gravitational wave with plus polarisation traveling through a gravitational wave interferometer.	12
1.6	A schematic of an interferometer with additional optics that form Fabry-Pérot cavities in the arms of the instrument.	15
1.7	Illustrations of two configuration types of delay line interferometer. .	16
1.8	A schematic of an interferometer with a power recycling mirror before the beam splitter.	17
1.9	A schematic of an interferometer with a signal recycling mirror before the photodetector.	18
1.10	Photographs of some of the first generation gravitational wave detector sites.	20
1.11	Artists' impressions of what ET (left) and LISA (right) could look like.	23
1.12	A plot showing the increase in strain sensitivity of ground-based detectors with each new generation of detectors.	24

1.13	A plot of the estimated advanced LIGO sensitivity curve based on calculations of the limiting noise sources.	26
1.14	A diagram of the uncertainty that arises in a normal coherent state compared with states where squeezing is applied to the amplitude quadrature and the phase quadrature.	32
2.1	A diagram showing how a contact potential arises between two metals in physical contact.	37
2.2	A diagram showing the localised states that arise in the band gap of an insulator.	39
2.3	A diagram showing the charge transfer between a metal and an insulator.	39
2.4	A diagram showing the charge transfer between two insulators. . . .	40
2.5	A diagram showing how ions are transferred in the contact electrification process.	43
2.6	A diagram of the surfaces of the optic and the earthquake stop defined in the model.	48
2.7	A plot of the convergence test data for the charging noise model. . .	53
2.8	A surface charge density map of a one inch silica disc after it was cleaned with acetone and a clean room cloth.	54
2.9	A plot of charging noise compared to the total aLIGO noise budget and the main limiting noise sources of the detector.	55
2.10	Kelvin probe data of the surface charge density across a silica disc before contact with Viton.	56
2.11	Kelvin probe data of the surface charge density across a silica disc after contact with Viton.	56
2.12	A plot of the expected charging noise after an aLIGO test mass has come into contact with a Viton tipped earthquake stop (blue) and a silica tipped earthquake stop (red).	57
2.13	Kelvin probe data of the charge on the surface of a silica disc before being cleaned with	59

2.14	Kelvin probe data of the charge on the surface of a silica disc after being cleaned with First Contact TM	59
2.15	Plot of charging noise after the removal of First Contact TM compared to the total aLIGO noise budget.	60
3.1	A photograph of a Kelvin probe similar to the one used for the investigations presented in this thesis.	68
3.2	A photograph of the vacuum tank where the Kelvin probe is situated.	69
3.3	A photograph of the Kelvin Probe and silica sample inside the vacuum tank.	69
3.4	A diagram showing the previous (a) and current (b) methods of scanning the surfaces of samples using the Kelvin probe.	70
3.5	A plot of the surface charge density on a silica sample changing over time.	70
3.6	A schematic diagram of the Kelvin probe set up.	71
3.7	A plot of the effective radius versus probe-sample separation distance.	74
3.8	A screen shot of a high speed Kelvin probe movie.	74
3.9	A screen shot of a high speed Kelvin probe movie that has each pixel more easily defined.	74
3.10	A zoomed in image of the Kelvin probe tip at maximum oscillation.	75
3.11	A zoomed in image of the Kelvin probe tip at minimum oscillation.	75
3.12	Plot of probe voltage versus probe current.	76
3.13	A plot of the probe output voltage versus the probe-sample distance.	77
3.14	A plot of the probe output voltage against oscillation frequency.	78
3.15	A schematic of the Kelvin probe set up represented in a circuit diagram.	79
3.16	Plots of V_{probe} against V_s for each of the four scenarios measured.	81
3.17	A Kelvin probe calibration curve that allows the surface charge density of the sample to be estimated from the Kelvin probe signal.	83
3.18	A plot of the surface charge on a silica sample changing over timescales that are much quicker than what would be expected from the correlation time of the movement of the surface charge.	86

4.1	A diagram of a classical gas discharge tube.	88
4.2	Voltage and discharge current characteristics of electrical gas discharges.	89
4.3	A diagram showing the structure of a glow discharge.	90
4.4	Paschen curves for air and argon gas.	94
4.5	Diagram (a) shows how the potential across the glow discharge may vary during the negative cycle of the square wave signal and diagram (b) shows how the potential across the discharge may change during the positive cycle of the square wave signal.	95
4.6	Glow discharge electrodes positioned inside the vacuum tank.	96
4.7	Glow discharge electrodes during operation.	96
4.8	A circuit diagram of the current-to-voltage converter used to convert the current measured by the Faraday cup into a voltage that could be measured with the Keithley multimeter.	98
4.9	A photograph of the sample stage set up for the glow discharge procedure.	99
4.10	A photograph of the Kelvin probe unshielded and in position to take a measurement of the surface charge density on the fused silica disc.	99
4.11	A schematic, aerial view, showing the glow discharge experimental set up.	99
4.12	A schematic diagram showing the PID feedback controller process.	102
4.13	An example plot of the voltage read by the Keithley multimeter from the current-to-voltage converter during the glow discharge process.	103
4.14	A photograph of a silica disc (left) and a silica sample with a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer coating (right).	104
4.15	A photograph of the set up required to show that the glow discharge procedure was capable of mitigating charge between two closely spaced optics.	104
4.16	Kelvin probe scan of the surface charge density across the coated silica sample before the glow discharge procedure.	106
4.17	Kelvin probe scan of the surface charge density across the coated silica sample after the glow discharge procedure.	106

4.18	Resulting charging noise after the glow discharge process compared to the aLIGO noise budget.	107
4.19	Kelvin probe scan of the surface charge density across the three inch silica sample before the glow discharge procedure.	108
4.20	Kelvin probe scan of the surface charge density across the three inch silica sample after the glow discharge procedure.	108
4.21	A rescaled plot of the Kelvin probe scan of the surface charge density across the three inch silica sample after the glow discharge procedure.	109
4.22	A photograph of the ellipsometer used to measure the transmission of the coated silica samples used in this investigation.	111
4.23	An example of the transmission data taken using the ellipsometer.	112
4.24	A histogram of all the transmission measurements of the glow discharge sample and the reference sample.	112
4.25	A plot of the empirical cumulative distribution functions of the transmission data obtained from transmission measurements of the reference sample and the glow discharge sample.	116
4.26	A photograph of a corona discharge gas ioniser designed by Liam Cunningham.	118
5.1	An SEM image of a silica disc before ordinary First Contact TM was applied.	124
5.2	An SEM image of a silica disc after ordinary First Contact TM has been removed from its surface.	124
5.3	An SEM image of a silica disc before First Contact TM with multi-walled CNT was applied.	125
5.4	An SEM image of a silica disc after FC with multi-walled CNT has been removed from its surface.	125
5.5	A photograph of one of the FC samples used in the resistivity measurements.	126
5.6	An example of a Talysurf scan of a Veroboard sample used in the resistivity measurements of FC.	127

5.7	A Talysurf scan of the surface of a PCB board with two tracks of copper etched off.	128
5.8	A schematic of the circuit used to measure the DC resistance of different types of FC.	129
5.9	A photograph of the complete PCB used to measure the DC resistance of different types of FC.	130
5.10	A plot of the measured resistance of two $5\text{ G}\Omega$ thin film resistors, connected in series, over time.	131
5.11	A plot of the measured temperature over the course of the test measurement against the DC resistance measured over the course of the measurement.	131
5.12	A plot of the measured resistance of the FC sample with a CNT mass concentration of $4.02 \times 10^{-2}\%$	133
5.13	A plot of temperature against resistance during the measurement of the FC sample with a CNT mass concentration of $4.02 \times 10^{-2}\%$. . .	133
5.14	A plot of the measured impedance of a $5\text{ G}\Omega$ resistor and a 22 pF capacitor connected in parallel, red, and the theoretical impedance, blue, over a range of frequencies.	135
5.15	A plot of the measured impedance of the FC sample with a CNT mass concentration of $4.02 \times 10^{-2}\%$	137
5.16	A plot of the Kelvin probe measurements of the charge deposited by FC with and without CNT.	139
5.17	The corona discharge experimental set up.	142
5.18	A close up of one of the pins that creates the corona discharge. . . .	142
5.19	Kelvin probe scan of the surface charge density across the coated silica sample before the corona discharge procedure.	143
5.20	Kelvin probe scan of the surface charge density across the coated silica sample after the corona discharge procedure.	143
5.21	Resulting charging noise after the corona discharge process compared to the aLIGO noise budget.	145

5.22	Kelvin probe scan of the surface charge density across the three inch silica sample before excess surface charge is reduced using the corona bar.	146
5.23	Kelvin probe scan of the surface charge density across the three inch silica sample after using the corona bar to mitigate excess charge. . .	146
5.24	A histogram of all the transmission measurements of the corona discharge sample and the reference sample.	148
5.25	A plot of the empirical cumulative distribution functions of the transmission data obtained from measurements of the reference sample and the sample used during the corona discharge investigation.	149
6.1	A schematic of a classic torsion balance set up.	154
6.2	A photograph of the vacuum tank and inner structure.	155
6.3	A photograph of the stepper motor set up.	156
6.4	A photograph of the torsion balance bob.	158
6.5	A photograph of the eddy current damper.	159
6.6	A cross sectional diagram of the eddy current damper.	159
6.7	A schematic of the capacitive sensor circuit.	161
6.8	A plot of capacitive sensor output voltage against rotation angle. . .	164
6.9	A plot of the decaying oscillations of the torsion balance after it had been excited.	165
6.10	Amplitude spectral density of the capacitive sensor output voltage showing a signal due to surface charge interaction at 1 mHz and a signal due to polarisation charge interaction at 2 mHz.	168
6.11	A diagram of the new sensor and actuator configuration of the torsion balance set up.	170
6.12	A plot showing the increase in sensitivity after the bob and sensor plates were shielded.	171
6.13	Plots showing the torque applied to the bob by the actuator plates against servo voltage.	173

6.14	Three plots showing the typical temperature variations measured by the three thermometers used in the experimental set up.	175
6.15	A plot showing the linear correlation between changes in temperature and changes in the capacitive sensor voltage.	176
6.16	A plot of the x and y tilt of the vacuum tank.	177
6.17	A plot of the magnetic field background recorded over a weekend in the room which contains the torsion balance.	178
6.18	The top plot shows the raw capacitive sensor data, in blue, before the least squares fit data, in red, is subtracted from it. The bottom plot shows the capacitive sensor data after the fit has been subtracted from it.	180
6.19	A plot showing the highest sensitivity reached by the torsion balance.	182
6.20	A photograph of the Mark II torsion bob.	183
6.21	A photograph of the sensor housing for the Mark II torsion balance.	185
6.22	The electrode in the roof of the housing is for applying the drive signal to the cube mass of the torsion balance.	185
6.23	A photograph showing an exploded view of the inside of the Mark II sensor housing.	186
7.1	A photograph of the torsion bob with new aluminium plates attached to the ends of each of its four arms.	189
7.2	A photograph of the silica sample positioned near the torsion bob to make a measurement of charging noise.	190
7.3	A schematic of the non-inverting op-amp circuit used to provide a servo control voltage to the actuator plates.	190
7.4	A schematic of the gains that have to be taken into account in the torsion balance set up.	193
7.5	A plot showing the change in servo effort, blue, as the torsion balance is moved in a sine wave motion, green.	195
7.6	A plot of the closed loop gain of the torsion balance PID system measured at different frequencies.	196

7.7	A plot of the torque sensitivity of the torsion balance while in free running mode, red, and while under servo control, black.	197
7.8	A plot of the amplitude of the servo voltage oscillation measured for different values of proportional gain.	199
7.9	A plot showing the change in servo voltage, blue, the fit to the servo voltage data, red, and the torsion balance sine wave motion, green, for the PID values used during the charging noise measurement. . .	200
7.10	A plot of the closed loop gain measured at different frequencies for the PID values used during the charging noise measurement.	200
7.11	Plots of the applied servo voltages during the charging noise measurement.	202
7.12	Plots of the typical servo voltages applied to the servo actuators when a charged sample is not present in the vacuum chamber.	203
7.13	A plot of the torque sensitivity of the torsion balance with a charged sample near the bob, red, and without a charged sample near the bob, black.	203
7.14	A plot of the Kelvin probe voltage against the servo voltage being applied to one of the actuator plates during the charging measurement.	204
7.15	Plots of the decaying Kelvin probe signal measured during the charging noise measurement (top) and the residuals of the fit to the data (bottom).	207
7.16	A plot of the torque sensitivity of the torsion balance with a charged silica disc present in the vacuum chamber (red) and estimates of the charging noise expected using the two different methods described in the text (blue and green).	208
7.17	A plot of the ratio of the torsion balance measurement and the theoretical estimate of charging noise against frequency.	209
7.18	A plot of the total aLIGO noise and charging noise.	210
B.1	An SEM image of a silica disc before FC with CNT is applied. This image was taken at the edge of the sample.	219

B.2	An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken at the edge of the sample. . .	219
B.3	An SEM image of a silica disc before FC with CNT is applied. This image was taken approximately 1 mm from the edge of the sample. .	220
B.4	An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken approximately 1 mm from the edge of the sample.	220
B.5	An SEM image of a silica disc before FC with CNT is applied. This image was taken approximately 2 mm from the edge of the sample. .	220
B.6	An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken approximately 2 mm from the edge of the sample.	220
B.7	An SEM image of a silica disc before FC with CNT is applied. This image was taken approximately 3 mm from the edge of the sample. .	221
B.8	An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken approximately 3 mm from the edge of the sample.	221
B.9	An SEM image of a silica disc before FC with CNT is applied. This image was taken approximately 4 mm from the edge of the sample. .	221
B.10	An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken approximately 4 mm from the edge of the sample.	221
B.11	An SEM image of a silica disc before FC with CNT is applied. This image was taken approximately 1 mm from the edge of the sample. .	222
B.12	An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken approximately 1 mm from the edge of the sample.	222
B.13	An SEM image of a silica disc before FC with CNT is applied. This image was taken approximately 2 mm from the edge of the sample. .	222

B.14 An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken approximately 2 mm from the edge of the sample.	222
B.15 An SEM image of a silica disc before FC with CNT is applied. This image was taken approximately 3 mm from the edge of the sample. .	223
B.16 An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken approximately 3 mm from the edge of the sample.	223
B.17 An SEM image of a silica disc before FC with CNT is applied. This image was taken approximately 4 mm from the edge of the sample. .	223
B.18 An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken approximately 4 mm from the edge of the sample.	223

List of Tables

2.1	An example of a triboelectric series.	41
3.1	Values of the capacitances and effective radii of the Kelvin probe and copper sheet for different probe-sample separation distances	73
3.2	Measured Kelvin probe oscillation amplitudes from the high frame rate movies	76
4.1	Results of the Kelvin probe measurements of the silica samples before and after the glow discharge procedure.	106
4.2	Results of the transmission measurements of the reference sam- ple and the glow discharge sample.	112
4.3	Results from the Student's t-test analysis carried out on the transmission data of the reference sample and the glow discharge sample.	114
4.4	Results from the Kolmogorov-Smirnov test analysis carried out on the transmission data of the reference sample and the glow discharge sample.	115
5.1	Measured gap distances between the copper pads of each First Contact TM sample and the length of each First Contact TM sample	128
5.2	Measured resistivities of First Contact TM with varying concen- trations of CNT using DC method	132
5.3	Measured resistivities of First Contact TM with varying concen- trations of CNT using AC method	136

5.4	Results from the corona discharge investigation using argon and nitrogen gas.	144
5.5	Results of the transmission measurements of the coated silica sample after the corona discharge procedure and a reference sample.	147
5.6	Results from the Student's t-test analysis carried out on the transmission data of the reference sample and the corona discharge sample.	148
5.7	Results from the Kolmogorov-Smirnov test analysis carried out on the transmission data of the reference sample and the corona discharge sample.	149

Acknowledgements

Jings, crivens and help ma boab! That last dash to the finish line was a lot more intense than I was expecting.

I have thoroughly enjoyed the last four years of pottering around in the lower switch room and, to a lesser extent, the pit. However, I feel the real enjoyment that you get out of something like this is from the many friends that you make along the way, and working with people with a real passion and enthusiasm for their work.

First of all I would like to thank my first supervisor, Giles Hammond. You have been an absolute pleasure to work with for the past four years and I feel like I have learned a lot under your guidance. Your enthusiasm for experimental physics has often been something that has kept me going through many frustrating moments. I would also like to thank Sheila Rowan and Jim Hough for giving me the opportunity to be a part of the Institute for Gravitational Research (IGR). You gave me the opportunity to present my work at various meetings around the world, something which I will always be grateful for, and always showed a keen interest in my research. I also owe a special thank you to Liam Cunningham who always had time to listen to my problems and complaints when there was no one else to go to. I thank the Science and Technology Funding Council (STFC) for the scholarship that allowed me to carry out the work presented in this thesis.

I would like to offer my thanks to many of the members of the IGR. A huge thanks should go to Iain Martin, Stuart Reid, Alan Cumming and Marielle Van Veggel for showing me how to use various instruments, helping me with prob-

lems that I just couldn't solve and for interesting and stimulating conversations at the biweekly 9:15 meetings. Jean Grieg deserves a big thank you for getting me where I needed to be on my travels and for giving me gentle reminders to make travel arrangements. I would like to say a special thank you to Matt Abernathy and Nicola Beveridge. We started our Ph.D.'s together in 461 and it has been a blast. I would also like to thank my other, past and present, office mates Keiran Craig, Chris Bell, Karen Haughian, Riccardo Bassiri, Keith Evans, Rahul Kumar and Peter Murray for thought-provoking physics discussions and also completely mental and hilarious discussions.

None of the experiments that I carried out for this thesis would have made it past scribbles in my note book had it not been for the excellent technical staff in the department. My hearty thanks go to Russell Jones and Mike Perruer-Lloyd for all their time, patience and expertise whenever I was needing help designing parts for an experiment or with SolidWorks. I also thank them for not getting too annoyed whenever I entered their office with a stack of technical drawings, all drawn in my own unconventional style. Thanks also to Colin Craig and Stephen Craig for making various bits and pieces on short notice. I would like to thank Steven O'Shea and Alastair Grant for all their help and advice whenever I was designing circuits for this work.

I would like to thank everyone that has joined me in a game of football over the past four years. I have had tonnes of fun playing football during my Ph.D. and it was always nice to forget about work for an hour on the pitch. I do apologise if you were on the receiving end of one of my shots during days when things in the lab weren't quite working out. I feel that an extra special thank you should go to Euan Bennet and Gail Penny for organising weekly games of football. If it wasn't for the both of you I would probably be a great deal heavier than I am now. Eric Yao also did a pretty good job of booking games when Gail or Euan weren't available.

A big thank you should go to my former band mates Andy, Adam and Dave.

Band practices and gigs were always a welcome distraction from difficult days in the lab but it wouldn't have been nearly as fun without your good humour and banter. I will make it along to a gig sometime now that I'm finished!

An extra special thank you should go to my girlfriend, Hannah, for her unconditional love and encouragement. If it wasn't for your support, especially in the final few months, I would have crumbled under the stress of it all.

Finally, I would like to thank my parents, my sister and my grandparents for their love and encouragement during my Ph.D.

Preface

This thesis presents an account of the research carried out at the Institute for Gravitational Research at the University of Glasgow between October 2008 and September 2012. This research specifically investigated charging noise in future gravitational wave detectors.

Chapter 1 discusses some of the important aspects of general relativity that differentiate it from classical Newtonian gravity. The idea of gravitational waves is introduced and the impact their observation would have on astronomy is highlighted. This chapter also discusses the status of current and future generation gravitational wave detectors and some of the advanced interferometric techniques that have been, or will be, used to allow these detectors to reach incredibly high sensitivities.

In Chapter 2 a brief overview of the current understanding of contact electrification is given. It is then shown that surface charges on insulators give rise to time varying forces which could create excess noise in future gravitational wave detectors. A detailed description of the method that the author used to model charging noise in a future gravitational wave detector, such as advanced LIGO, is discussed. A review of different scenarios that are believed to lead to excess detector noise caused by surface charges and possible charge mitigation strategies are given. Modeled estimates of the charging noise expected from some of these different scenarios are presented.

A description of the Kelvin probe experimental set up used for measuring surface charge densities in the investigations presented in this thesis is given in Chapter 3. The LabVIEW Kelvin probe scanning program was initially

written by the author with further improvements made by Mr. N. Hurst. It is shown that the Kelvin probe output signal can be accurately modeled. Models of the expected Kelvin probe signal are included for two scenarios: one where the probe-sample distance is increased and one where oscillation amplitude is increased. These models were created by the author at the suggestion of Dr. G. D. Hammond. The ANSYS models needed for determining the capacitance between the Kelvin probe tip and a reference surface were created by Dr. L. Cunningham and the analysis was carried out by the author for different geometries. The Kelvin probe had to be calibrated in order to obtain meaningful surface charge density measurements. A detailed description of the calibration process followed by the author is given.

In Chapter 4 an investigation into a charge mitigation process which uses a Proportional-Integral-Derivative (PID) controlled AC glow discharge is presented. This process was developed by the author at the suggestion of Dr. G. D. Hammond and Prof. S. Rowan. The experimental set up of the glow discharge charge mitigation process was created by the author through fruitful discussions with Dr. G. D. Hammond. The PID controller was programmed in LabVIEW by the author and the surface charge density measurements taken with the Kelvin probe were made by the author. The transmission of the samples used in this investigation was measured using an ellipsometer. These measurements were made by the author with assistance from Dr. L. Cunningham. The data from the transmission measurements was analysed using Student's t-test and the Kolmogorov-Smirnov test which was carried out by the author. The Kolmogorov-Smirnov test was carried out at the suggestion of Prof. M. Hendry. The chapter concludes with a discussion of suggested improvements to the glow discharge process. These improvements were determined through discussions with Dr. L. Cunningham.

In Chapter 5 an investigation into mitigating charge during the cleaning procedure of future gravitational wave detectors is presented. The work in

this chapter was carried out at the suggestion of Dr. G. D. Hammond and Prof. S. Rowan. The samples of First ContactTM with different concentrations of carbon nano-tubes mixed in them were provided by Prof. J. Hamilton of Photonic Cleaning Technologies. The resistivity measurements of the different types of First ContactTM were made by the author. The circuit used to measure the DC resistivity was designed by the author through stimulating discussions with Mr. S. O'Shea and Mr. A. Grant. The PCB circuit was fabricated at the department of electrical engineering at the University of Glasgow. Kelvin probe measurements of a silica sample cleaned with both ordinary First ContactTM and First ContactTM mixed with carbon nano-tubes were made by the author to ascertain whether the inclusion of carbon nano-tubes affected the amount of surface charge deposited on the silica sample. The procedure for mitigating surface charge using a corona discharge was created by the author through discussions with Mr. J. Patterson of Static Clean International and Dr. G. D. Hammond. The surface charge density measurements and the transmission measurements of the samples were made by the author. The analysis of the transmission data was carried out by the author and was the same analysis as carried out in Chapter 4.

Chapter 6 details the construction of a torsion balance apparatus that was used to directly measure charging noise. The work in this chapter was carried out at the suggestion of Dr. G. D. Hammond and Prof. S. Rowan. The experimental set up was created by the author with assistance from Dr. G. D. Hammond. The torsion bob and sensors/actuators were designed by the author. The author acknowledges the assistance of Mr. R. Jones and Mr. M. Perreux-Lloyd for stimulating discussions and suggestions in the design process of the torsion balance Mark I and II. The author also gratefully acknowledges the staff in the mechanical workshop in the physics department at the University of Glasgow for machining the parts of the torsion balance Mark I and II. The capacitive sensor used to sense the position of the torsion balance was

constructed by Mr. I. McCrindle and designed by Dr. G. D. Hammond. The original LabVIEW program used to control and read data from the torsion balance was created by the author with later amendments made by both the author and Dr. G. D. Hammond. The measurements presented in this chapter were made by both the author and Dr. G. D. Hammond, and the analysis of the data was carried out by the author.

A torsion balance measurement of charging noise, made with the Mark I instrument, is presented in Chapter 7. The method of calibration of the torsion balance was suggested by Dr. G. D. Hammond and carried out by the author. The charging noise measurement presented in this chapter was taken and analysed by the author. The author acknowledges the help of Dr. G. D. Hammond, Dr. B. W. Barr and Dr. I. S. Heng for useful discussions on the the calibration of the torsion balance under servo control and the analysis of the charging noise measurement. The theoretical charging noise estimates presented in this chapter were calculated by the author.

Summary

Investigations of Charging Noise in Future Gravitational Wave Detectors

Some of the biggest discoveries in astronomy have come from probing new regions of the electromagnetic spectrum. Penzias and Wilson famously detected faint whispers from the creation of the Universe using the Holmdel horn antenna at Bell telephone laboratories. The microwave signals that they detected strongly supported the big bang theory of the creation of the Universe. Radio wave observations have provided the first glimpses of rapidly rotating neutron stars, X-ray astronomy gave the first evidence for the existence of black holes and infrared images gave the first clear indication for the existence of brown dwarfs. If gravitational waves could be detected directly, and regularly, they would no doubt provide many spectacular discoveries. They would provide an entirely new window through which to observe the cosmos and further our understanding of the Universe.

Gravitational waves are a result of Einstein's theory of general relativity which postulates that these waves, created by asymmetrically accelerating masses, propagate through the Universe as ripples in the curvature of space-time at the speed of light. The only sources of gravitational waves that are considered detectable are astronomical ones. These sources include: asymmetrically rotating neutron stars, coalescing compact object binary systems and supernovae.

Detections of gravitational waves have been attempted for nearly 50 years

and there has still been no confirmed observation. This is because the force of gravity is extremely weak compared to the other fundamental forces of nature. This poses an extremely difficult challenge to experimental physicists for measuring strains in space-time caused by passing gravitational waves that are predicted to be of the order 10^{-21} and smaller. Kilometer scale interferometers are used to try and measure the strains in space-time created by gravitational waves. There is a world wide network of detectors that are currently being, or have already been, upgraded to an advanced detector status. In order to get these detectors to a sensitivity high enough to stand a chance of detecting gravitational waves, researchers around the world have been developing innovative and novel technologies, and investigating high quality, low mechanical loss materials, to reduce the limiting noise sources of these instruments. A brief overview of gravitational wave detection and the nature of gravitational waves will be given in Chapter 1.

One noise source that may become more conspicuous as gravitational wave detectors become more sensitive is noise due to the movement of electrical charges on the detector optics. Charge can be transferred to the detector optics through various processes such as; dust abrasion on the surface of the optics as the vacuum chamber is being pumped out, contact with nearby structures such as earthquake stops, cleaning of the optics and cosmic rays. In some instances the charge density transferred can be as high as 10^{-7} C/m² which would cause the sensitivity of advanced detectors to be significantly reduced at some frequencies. Once charge is deposited on a silica test mass, it slowly redistributes itself across the surface of the optic due to a very small amount of surface conductivity. This creates a fluctuating electric field, and therefore a fluctuating force, which acts on the test mass causing displacement noise in the detector. Chapter 2 will briefly discuss the physics of why electrical charge is transferred between materials during contact. Charging noise is then discussed, with a detailed description of how it is modeled throughout this

thesis, and some examples of different charging noise scenarios are given. It is shown that if the correct measures are not taken, charging noise could easily limit the sensitivity of future gravitational wave detectors.

For the investigations presented in this thesis the surface charge on fused silica samples had to be measured. A useful instrument for measuring surface charge on the surface of insulators is the Kelvin probe. Chapter 3 will explain how the Kelvin probe works and show that the Kelvin probe signal can be accurately modeled. There is also a detailed explanation of how the Kelvin probe signal was calibrated, as this was a very important process for making measurements in terms of surface charge density.

To minimise down time during science runs, future gravitational wave detectors should accommodate at least one charge mitigation procedure. There have been many different procedures that have been developed by institutions in the LIGO Scientific Collaboration, each of which is unique in its own way. The most promising method of mitigating surface charge on gravitational wave detector optics looks like it will be through the use of an ionised inert gas. Chapter 4 presents a study on the use of ionised argon gas, created by a PID controlled AC glow discharge, for mitigating excess surface charge on gravitational wave detector optics. It is shown that ionised argon gas created by a glow discharge successfully reduces excess charge on bare fused silica and silica with a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer coating. TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ is the type of coating that will be used in the second generation detector, advanced LIGO. Kelvin probe measurements of the silica samples before and after the charge mitigation procedure are used to estimate the level of charging noise that would be expected in a gravitational wave detector. The glow discharge method detailed in Chapter 4 reduces surface charge to a level that would not limit a second generation gravitational wave detector such as advanced LIGO. It is a concern that some of the charge mitigation processes proposed for future detectors may damage the multi-layer reflective coating on

the detector optics. The transmission of the coated silica sample used in this investigation is measured and compared to a reference sample that was not exposed to the glow discharge to show that it does not cause an observable change to the transmission of a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer coating, within the experimental limits of the instrument.


It is very likely that the optics of gravitational wave detectors will become charged during the cleaning process of the mirrors before they are placed under vacuum in the detector. Chapter 5 discusses different ideas of reducing charging during the cleaning of gravitational wave detector optics. It is likely that the optics of advanced LIGO will be cleaned with a cleaning and protecting product called First ContactTM. Unfortunately, First ContactTM is known to deposit a substantial amount of charge on the surface of silica. It has been suggested that mixing carbon nano-tubes in the First ContactTM solution could reduce its resistivity and therefore might deposit less charge on the optics. The resistivity of First ContactTM with varying amounts of carbon nano-tubes is measured, using both an AC and DC method, to ascertain whether the carbon nano-tubes have any effect on the resistivity of First ContactTM. Unfortunately, due to several problems with the experiment, the resistivity measurements made are inconclusive as to whether carbon nano-tubes have a significant effect on the resistivity of First ContactTM. However, it is shown through surface charge density measurements that carbon nano-tubes do not reduce the amount of surface charge that is deposited on the surface of silica by First ContactTM. A method of reducing the surface charge deposited by First ContactTM by ionising nitrogen gas with a commercially available corona discharge bar is detailed in Chapter 5. This method is shown to reduce the surface charge on bare fused silica and fused silica with a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer coating to a level that would not create excess noise in a future gravitational wave detector such as advanced LIGO. Using the same analysis as described in Chapter 4, it is also shown that this charge mit-

igation method does not effect the transmission of a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer coating, within the experimental limits of the instrument.

A direct measurement of charging noise has still to be made, so it is still unconfirmed whether the assumptions made about the Weiss theory of charging noise are in fact true. Chapter 6 outlines the development of a servo controlled torsion balance experiment for measuring charging noise on a fused silica disc. This chapter discusses the design of the torsion balance and presents initial measurements of surface and polarisation charges on a fused silica disc. The calibration of the instrument in free running mode, i.e. with no servo control, is discussed. This was an important process for expressing the measurements made in terms of a torque noise. The limiting noise sources of the instrument are addressed and it is shown that by making some minor changes to the experimental set up the sensitivity of the instrument can be significantly increased. The improved sensitivity of the instrument is within an order of magnitude of the thermal noise limit at 0.1 mHz, this was deemed adequate sensitivity to begin measurements of charging noise. This chapter also discusses the development of a Mark II version of the torsion balance which should be easier to use and should achieve greater sensitivity than the Mark I instrument. Due to time constraints the performance of the Mark II torsion balance is not presented in this thesis.

The charging noise measurement made with the Mark I torsion balance is presented and discussed in Chapter 7. As well as this measurement, the method of calibrating the instrument while under servo control is detailed. This calibration method is much different from the one described in Chapter 6 because it required measuring the closed loop gain of the torsion balance set up. It is shown that placing a charged silica disc in the vacuum chamber decreases the sensitivity of the instrument significantly. Using the measured offset torque of the torsion bob and the time constant of the surface charge on the silica disc, measured with a Kelvin probe, the level of charging noise that should be

expected is estimated. It is shown that the theoretical estimate of the charging noise and the observed noise level agree quite well. Thus confirming that the assumptions made about charging noise are suitable for predicting charging noise in gravitational wave detectors.



By Bill Watterson.

Chapter 1

The Nature and Detection of Gravitational Waves

1.1 Introduction

Gravitational waves still remain one of the most mysterious and elusive phenomena in physics. Einstein postulated that gravitational waves should manifest themselves as ripples in space-time emitted by asymmetrically accelerating mass in his theory of general relativity [1]. The only detectable sources of gravitational waves are astronomical ones; these include compact object binary systems, supernovae and pulsars. It is also thought that there is a stochastic gravitational wave background that was generated in the very early universe. These waves may have been created by quantum fluctuations in the early universe and then amplified during the inflationary period of the universe. If gravitational waves were detected they may provide more insight into exotic objects such as black holes or even give direct information from the very early universe.

A worldwide network of six interferometric detectors has been set up to detect differential strains in space-time caused by passing gravitational waves generated by astrophysical events. There are three detectors in the US (two with arm lengths of 4 km and one with an arm length of 2 km) that form

the LIGO network [2], one detector located in Italy, Virgo [3], one detector in Germany, GEO600 [4] and one detector in Japan, TAMA 300 [5]. In recent years there has been a worldwide effort to improve the sensitivity of these instruments and upgrade some of them to an advanced detector status [6–8]. This will greatly improve chances of detection. To get the detectors to this improved sensitivity a significant amount of research has gone into investigating how to lower the main limiting noise sources of the detectors, which are laser noise, thermal noise and seismic noise, and other technical noise sources such as charging noise.

It's no wonder that the detection of gravitational waves remains such a tantalising prospect for both physicists and astronomers. Their detection would be a great test of general relativity and open up a new window through which to observe the universe that will no doubt provide some spectacular new discoveries.

This chapter will briefly describe the nature of gravitational waves and how they arise in general relativity. There will also be some discussion of the main astronomical sources of gravitational waves. The remainder of the chapter will focus on the working principles of interferometric detectors which will cover advanced techniques that help improve the sensitivity of the instrument and the detector's main limiting noise sources.

1.2 The Nature of Gravitational Waves

In 1916 Albert Einstein published his revolutionary paper on general relativity [1] which challenged Newton's law of universal gravitation [9]. Einstein's general theory of relativity describes gravity as a consequence of the distortion of the geometry of space-time, whereas Newton described gravity as an instantaneous force. One of the consequences of Einstein's theory is that accelerating asymmetric masses should produce gravitational waves which carry gravitational information and energy away from the source at the speed of light. It is

this key feature of general relativity that presents a solution to the problem of the instantaneous gravitational force described in Newton's theory of gravity.

One of the most important elements of general relativity are Einstein's field equations, which describe the curvature of space produced by matter or energy and can be expressed in the form [10],

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = \frac{8\pi G}{c^4}T_{\mu\nu}, \quad (1.1)$$

where $G_{\mu\nu}$ is the Einstein tensor that describes the curvature of space, $g_{\mu\nu}$ is the metric tensor which encapsulates all the information about the region of space-time of interest, R is scalar curvature, $R_{\mu\nu}$ is the Ricci tensor which represents curved space, G is the Newtonian constant of gravitation, c is the speed of light in a vacuum and $T_{\mu\nu}$ is the stress energy tensor that describes the density and flux of energy and momentum in space-time. If a small perturbation is introduced, $h_{\mu\nu}$ where $|h_{\mu\nu}| \ll 1$, in a flat area of space-time, assuming a weak gravitational field, the metric becomes,

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad (1.2)$$

where $\eta_{\mu\nu}$ is the Minkowski metric which represents flat space-time. With the new value of $g_{\mu\nu}$ substituted into Equation (1.1) it is possible to derive a wave equation that describes the motion of the perturbation $h_{\mu\nu}$. The field equations can be simplified by careful selection of the form of $h_{\mu\nu}$ and by assuming $T_{\mu\nu} = 0$, since we are only concerned with an area of almost flat space-time in the absence of any other energy or mass. Einstein's field equations simplify to a wave equation of the form [10],

$$\left(\nabla^2 - \frac{1}{c^2} \frac{\partial^2}{\partial t^2} \right) h_{\mu\nu} = 0. \quad (1.3)$$

This is an important result as it shows that the metric perturbation of space-time, $h_{\mu\nu}$, has the same mathematical form as a wave that is traveling at the

speed of light in a vacuum. A full and detailed derivation of the gravitational wave equation can be found in the work of C. W. Misner, K. S. Thorne and J. A. Wheeler [10].

The simplest form of $h_{\mu\nu}$ can be obtained by making an appropriate gauge transformation where the 16 components of $h_{\mu\nu}$ can be reduced to just 2 independent components [10]. This simplified form of $h_{\mu\nu}$ is called the transverse-traceless gauge and can be expressed as,

$$h_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_{xx} & h_{xy} & 0 \\ 0 & h_{xy} & -h_{xx} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad (1.4)$$

for a wave propagating in the z direction. This gauge is said to be both transverse and traceless because the perturbation caused by the wave acts perpendicular to the motion of the wave, see *Figures 1.1* and *1.2*, and because the trace of the tensor, given in Equation (1.4), is $h_{xx} - h_{xx} = 0$. The two remaining components, h_{xx} and h_{xy} , represent the two polarisation states of a gravitational wave, which are referred to as the “plus” and “cross” states, where $h_+ = h_{xx}$ and $h_\times = h_{xy}$. Although it may not be initially obvious, the two polarisation states of a gravitational wave are orthogonal to each other. *Figure 1.2* highlights the orthogonality of the wave perturbation for the plus and cross polarisation states.

If a gravitational wave with h_+ or h_\times were to pass through a ring of free particles in the z direction it would cause the ring to be stretched and squeezed as shown in *Figure 1.1*. This effect produces a strain as the test particles are extended and compressed by a distance ΔL . It is the strain produced by gravitational waves that interferometric gravitational wave detectors try to detect. There are currently several long base line interferometers that have been built for the very purpose of measuring the small strains produced by




Figure 1.1: An illustration of the effect of the two polarisation states of a gravitational wave on a ring of free particles. When $h_{xx} \neq 0$ and $h_{xy} = 0$ the wave has plus polarisation and when $h_{xy} \neq 0$ and $h_{xx} = 0$ the wave has cross polarisation. The phase of the wave is shown along the bottom of the figure.




Figure 1.2: This diagram highlights the orthogonality of the plus (left) and cross (right) polarisation states of a gravitational wave.

gravitational waves. The total amplitude of the strain that would be observed by a detector can be expressed as,

$$h = \frac{2\Delta L}{L}, \quad (1.5)$$

where L is the arm length of the detector and ΔL is the fractional change in the arm length of the detector. These detectors will be discussed later in this chapter.

1.3 Sources of Gravitational Waves

It can be easily shown from Einstein's field equations, Equation (1.1), that space-time is in fact very stiff. The constant in Equation (1.1), $8\pi G/c^4 \approx 2 \times 10^{-43} \text{ m}^{-1}\text{kg}^{-1}\text{s}^2$, is a very small number, therefore, it would take a huge energy density to cause a very small distortion of space-time. This means the only sources of gravitational waves that could be considered detectable have to involve an incredibly large amount of mass or energy to create ripples in space-time, and the only sources that are capable of this are astronomical ones. The three main astrophysical sources of gravitational waves will be discussed here. Only a few examples of each source type are given and this is in no way intended to be a comprehensive list of sources. More detailed descriptions on the various types of gravitational wave sources can be found in the references that are cited in the following sections and also in [11].

1.3.1 Burst Sources

If a source emits a short burst of gravitational waves for a duration of about a second or less it is classified as a burst source [2]. Some examples of burst sources that are thought to lie within the detection band of ground-based detectors are coalescing compact object binary systems and supernovae [12].

When a compact object, such as a neutron star or a black hole, is in a binary orbit with another compact object the system will lose energy due to the emission of gravitational radiation. As a result of this the two stars will spiral in towards each other. As the stars pass their last stable orbit they will rapidly begin to fall towards each other and merge. The result of this event is a burst of gravitational radiation which is estimated by Schutz [13] to create a strain of,

$$h = 1 \times 10^{-23} \left(\frac{100 \text{ Mpc}}{r} \right) \left(\frac{M_b}{1.2M_\odot} \right)^{\frac{5}{3}} \left(\frac{f}{200 \text{ Hz}} \right)^{\frac{2}{3}}, \quad (1.6)$$

where r is the distance of the source from Earth, $M_b = \frac{(M_1 M_2)^{\frac{3}{5}}}{(M_1 + M_2)^{\frac{1}{5}}}$ is the mass parameter of the binary system and M_1 and M_2 are the masses of the two compact objects, M_\odot is solar mass and f is the frequency of the wave.

The explosive death of a star, called a supernova, is one of the most violent astronomical events. When a star can no longer sustain nuclear fusion there is no pressure being generated from within the star to halt gravitational collapse. The result of this collapse is an extremely powerful explosion that releases a tremendous amount of energy that could emit gravitational waves if the mass distribution of the supernova is asymmetric. Sathyaprakash and Schutz [14] estimate that for a supernova in the Milky Way at a distance of 10 kpc from Earth, emitting the energy equivalent of $10^{-7} M_\odot$, at a frequency of 1 kHz, and lasting for 1 ms, the gravitational wave strain amplitude would be,

$$h \sim 6 \times 10^{-21} \left(\frac{E}{10^{-7} M_\odot} \right)^{\frac{1}{2}} \left(\frac{1 \text{ ms}}{T} \right)^{\frac{1}{2}} \left(\frac{1 \text{ kHz}}{f} \right) \left(\frac{10 \text{ kpc}}{r} \right), \quad (1.7)$$

where E is the energy of the wave and T is the duration of the gravitational wave burst.

1.3.2 Continuous Sources

As the title of this group suggests, an astronomical source that continuously emits gravitational waves at roughly a fixed amplitude and frequency for at least a few weeks is classified as a continuous source [15]. Some examples of continuous sources that are thought to be within the detection band of ground-based detectors are neutron stars in X-ray binary systems and isolated neutron star systems such as pulsars [12].

An X-ray binary system consists of a neutron star which is accreting mass from a neighboring star. As matter is accreted it releases a huge amount of gravitational potential energy in the form of X-rays. Asymmetries in the accretion could lead to the radiation of gravitational waves. It is thought [16]

that X-ray binaries should produce gravitational waves with a strain of,

$$h \sim 8.5 \times 10^{-27} \left(\frac{r}{10 \text{ kpc}} \right)^{-1} \left(\frac{f}{1 \text{ kHz}} \right)^{-\frac{1}{2}} \left(\frac{\dot{M}}{10^{-8} M_{\odot} \text{ yr}^{-1}} \right)^{\frac{1}{2}}, \quad (1.8)$$

where r is the distance of the source from Earth, f is the gravitational wave frequency and \dot{M} is the accretion rate.

Pulsars are rapidly rotating, highly magnetised neutron stars. Pulsars can emit gravitational waves if they have a mass asymmetry about their axis of rotation. The gravitational waves produced by this type of pulsar will be emitted at twice the rotation frequency of the pulsar because the star spins about its net center of mass, so it effectively has mass excesses on both sides of the star [14]. An estimate for the strain created by such a pulsar can be calculated in the following way [17],

$$h \sim 6 \times 10^{-25} \left(\frac{f_{\text{rot}}}{500 \text{ Hz}} \right)^2 \left(\frac{1 \text{ kpc}}{r} \right) \left(\frac{\epsilon}{10^{-6}} \right), \quad (1.9)$$

where f_{rot} is the rotation frequency of the pulsar and ϵ is the ellipticity.

It was a binary pulsar system, far from coalescence, that provided the first indirect evidence for gravitational waves [18, 19]. Russell Hulse and Joseph Taylor observed the binary system (PSR B1913+16) in 1974 using the Arecibo radio observatory and won the Nobel Prize in Physics for their discovery in 1993. The data from PSR B1913+16 is shown below in *Figure 1.3*. The predicted orbital phase shift calculated from general relativity matches the experimental values to within $\sim 1\%$ [20].

1.3.3 Stochastic Sources

As well as the two classes of sources mentioned previously, there is believed to be a stochastic background of gravitational waves. This is somewhat analogous to the cosmic microwave background and is the result of the superpositions of many gravitational wave sources of astrophysical and cosmological origins




Figure 1.3: A plot of the change in the orbital phase in the binary pulsar system PSR B1913+16 as a function of time. The plot is reproduced from [21].

which cannot be resolved [22]. The stochastic gravitational wave background is predicted [2] to have a strain noise power spectrum of the form,

$$h = 4 \times 10^{-22} \sqrt{\Omega_{GW}} \left(\frac{100 \text{ Hz}}{f} \right)^{\frac{3}{2}} \text{ Hz}^{-\frac{1}{2}}, \quad (1.10)$$

where Ω_{GW} is the gravitational wave energy density per unit logarithmic frequency, divided by the critical energy density needed to close the universe.

One of the major goals in gravitational wave detection would be to detect the stochastic gravitational wave background. Its detection could possibly give some insight into the structure of the Universe in its very early stages of creation and allow the laws of physics to be examined at energies higher than that achievable in a laboratory.

1.4 Gravitational Wave Detection

The first attempts at detecting gravitational waves began in the 1960s when Joseph Weber pioneered the first gravitational wave detector [23]. Weber's detector consisted of a high Q ¹ aluminium bar which would resonate if a gravitational wave of an appropriate amplitude and frequency passed through it. Weber made several claims of coincident detections between the bar detectors that he operated [24, 25], but other bar detector experiments failed to verify these detections [26–29]. Even though bar detectors have improved in sensitivity since Weber's initial efforts they are still limited to narrow detection bands and are still not as sensitive as current interferometric detectors [30].

Modern gravitational wave detectors use interferometry to search for gravitational waves. The idea of interferometric detectors was proposed by Gertsenshtein and Pustovit in 1962 [31], however, the construction of the first kilometre scale interferometric detector did not begin until 1994 [32]. Interferometry has several advantages over bar detectors such as a broader detection frequency band and much higher sensitivity. These detectors are designed to detect low frequency gravitational waves in the range of tens of Hz to several kHz. The basic working principles of interferometry and advanced interferometric techniques are discussed in this section.

1.4.1 Laser Interferometry

Interferometers function by splitting a beam of monochromatic laser light along the two arms of the instrument, which are orthogonal to each other, as shown in *Figure 1.4*. After the light is split at the beam splitter it traverses the length of the arms and reflects off mirrors, labeled Mirror 1 and Mirror 2 in *Figure 1.4*, at the ends of the arms. The beams of light then travel back to the

¹The mechanical quality factor, or Q , is a dimensionless quantity related to the dissipation of an oscillator at a resonant frequency. Q can be defined as the ratio of the stored energy of the oscillating system to the energy dissipated per radian of the oscillation.

beam splitter where they recombine and produce an interference pattern at the instruments output port, which is usually monitored using a photodetector.




Figure 1.4: A schematic of a basic interferometer set up.

Using the simplified example given by Saulson [20] it is possible to show how a gravitational wave can effect the output power of an interferometric detector. Lets consider the situation where a gravitational wave with plus polarisation passes through an interferometer like the one shown in *Figure 1.5*. The wave is traveling in the z -direction, into the page in *Figure 1.5*, and has a strain amplitude of h_+ . As the gravitational wave travels through the interferometer, the arms, of length L , of the detector are stretched and squeezed, increasing or decreasing the amount of time it takes the light to traverse the arms of the instrument. If the detector arm that lies along the x axis is considered first, it is possible to show from special relativity that the space-time interval, ds , between two space-time events linked by a light beam is,

$$ds^2 = (1 + h_+) dx^2 - c^2 dt^2 = 0, \quad (1.11)$$

where dx is the change in displacement along the x axis and dt is the change in travel time of the laser. This equation can be rearranged to give an expression




Figure 1.5: An illustration of the effect of a gravitational wave with plus polarisation traveling through a gravitational wave interferometer. The phase of the wave is shown along the bottom of the figure.

for dt ,

$$dt = \frac{\sqrt{1 + h_+}}{c} dx. \quad (1.12)$$

The change in travel time of the laser light between the beam splitter and the end mirror can be determined by integrating Equation (1.12). This expression can then be simplified by using binomial expansion which gives,

$$\int_0^{\tau_1} dt = \frac{1}{c} \int_0^L \sqrt{1 + h_+} dx \approx \frac{1}{c} \int_0^L \left(1 + \frac{1}{2} h_+ \right) dx, \quad (1.13)$$

assuming that higher order terms resulting from the expansion are negligible. τ_1 in Equation (1.13) is the time it takes the laser beam to travel from the beam splitter to the end mirror. A similar expression can be found for the change in travel time of the laser as the light travels back to the beam splitter. The only difference is that the right hand side of the equation becomes negative due to the change in direction,

$$\int_{\tau_1}^{\tau_2} dt \approx -\frac{1}{c} \int_L^0 \left(1 + \frac{1}{2} h_+ \right) dx, \quad (1.14)$$

where τ_2 is the time it takes the laser beam to travel from the end mirror back to the beam splitter. Normally the light in each arm of the interferometer would take a time of $\tau = 2L/c$ to travel from the beam splitter to the end test mass and back to the beam splitter in the absence of gravitational waves. Taking

into account the disturbance caused by the gravitational wave, τ becomes,

$$\tau = \frac{2L}{c} + \frac{1}{2c} \int_0^L h_+ dx - \frac{1}{2c} \int_L^0 h_+ dx = \frac{2L}{c} + h_+ \frac{L}{c}. \quad (1.15)$$

The gravitational wave will make equal but opposite perturbations in each arm, therefore, the time it takes for the light to traverse each arm is,

$$\tau_x = \frac{2L}{c} + h_+ \frac{L}{c} \quad (1.16)$$

$$\tau_y = \frac{2L}{c} - h_+ \frac{L}{c}. \quad (1.17)$$

The total difference in travel time, $\Delta\tau$ for the light in each arm is then,

$$\Delta\tau = h_+ \frac{2L}{c}. \quad (1.18)$$

The result of this change in travel time is a change in relative phase of the light in both arms of the detector which is measured at the output port as a change in output power of the interferometer. The phase of the wave, ϕ , can be expressed as,

$$\phi = 2\pi f\tau, \quad (1.19)$$

where f is the frequency of the laser light. The phase difference can be expressed simply as,

$$\Delta\phi = \Delta\tau \frac{2\pi c}{\lambda}, \quad (1.20)$$

which then becomes,

$$\Delta\phi = h_+ \tau \frac{2\pi c}{\lambda}. \quad (1.21)$$

The output power of the interferometer can be expressed as,

$$P_{out} = P_{in} \cos^2(\Delta\phi), \quad (1.22)$$

where P_{in} is the input power of the interferometer. As the gravitational wave

passes through the interferometer the light phase builds up in the interferometer and reaches a maximum for $\tau = 0.5\lambda$. If the light stays in the arms of the detector for longer than $\tau = 0.5\lambda$ then the effect of the gravitational wave will start to cancel itself out.

It should be noted that gravitational wave detectors are obviously far more advanced than the description given above for a basic interferometer. In gravitational wave detectors the end mirrors, or test masses as they are commonly referred to, hang from sophisticated suspension systems [33, 34] in order to damp any seismic vibrations and so that they act as freely suspended objects in the presence of a gravitational wave. The suspension system is also connected to a servo control system [35–37] which allows the test masses to be controlled in all degrees of freedom. The test masses are locked in a position that allows a dark fringe to fall on the photodetector, also referred to as a dark port, which is located at the output port of the instrument. In gravitational wave detection it is important that a dark fringe falls on the photodetector². The signal that would be produced by a passing gravitational wave is very small and this very small signal would be impossible to pick out of an already high signal that is randomly fluctuating if a bright fringe were falling on the photodetector. There are also many additional advanced optical configurations employed in real detectors, such as Fabry-Pérot cavities and power recycling mirrors, that are used to increase their sensitivity. Some of these advanced interferometric techniques will be discussed in the following sections.

1.4.2 Fabry-Pérot Cavities

Fabry-Pérot cavities [38] effectively allow the arm lengths of an interferometer to be increased without increasing their physical size and are used in the LIGO, Virgo and TAMA detectors. This would therefore increase the relative optical

²This is not the case in all gravitational wave detectors. Advanced LIGO will use a DC readout which operates at a small power output.




Figure 1.6: A schematic of an interferometer with additional optics that form Fabry-Pérot cavities in each arm of the instrument. The implementation of Fabry-Pérot cavities effectively increases the arm lengths of the detector and therefore increases its sensitivity.

phase shift that would result from a change in displacement of the interferometer test masses. In a Fabry-Pérot cavity the mirrors nearest the beam splitter, see *Figure 1.6*, are partially reflective and the end mirrors are fully reflective. When the light enters the cavity through the partially reflecting mirror the light becomes trapped in the cavity, this allows light to build up inside the arm cavities. The light then traverses the cavity many times before leaving the cavity through the partially reflecting mirror to the detector output. Since the laser beam has to travel through the partially reflecting mirror, the mirror must be of high optical quality to reduce scattering of the laser beam.

1.4.3 Delay Line Interferometry

Delay line interferometry [39] is another method of increasing the optical path of the laser light in the detector. A delay line in an interferometer consists of two curved mirrors. The mirror closest to the beam splitter has an aperture in it to allow the laser beam in and out of the delay line. Once the laser beam enters the delay line through the aperture it bounces back and forth a




Figure 1.7: *Illustrations of two configuration types of delay line interferometer. Configuration (a) is used in the GEO600 detector where the arms are “folded” to increase the optical path length. Set up (b) shows a delay line where the light is reflected between the two mirrors in each arm cavity many times to increase the optical path length.*

predetermined number of times before exiting, thus increasing the optical path of the laser light. A variation of the delay line interferometer is a configuration where the arms are said to be “folded”. In a folded interferometer set up the mirror closest to the beam splitter does not have an aperture in it and the laser light only circulates between the two mirrors once, rather than multiple times. This is the interferometric configuration utilised by the GEO600 detector [4]. *Figure 1.7* illustrates the two different delay line configurations.

1.4.4 Power Recycling

Due to the small signals created by gravitational waves, wasted photons in the detector could be used to increase the overall laser power circulating in the arms detector and therefore increase the power of the detected signal. Advanced interferometric techniques such as power recycling [40] allow the majority of the laser light, that would otherwise be wasted, to be effectively recycled. When a gravitational wave signal is detected, some of the light exits through the dark port of the instrument, however, the majority of the light exits through the input port. Even during normal operating conditions, where no gravitational wave perturbations are present, all the laser light exits the




Figure 1.8: A schematic of an interferometer with a power recycling mirror before the beam splitter. The power recycling mirror “recycles” any light that may be exiting the instrument through the input port.

interferometer through the input port. An extra partially transmitting cavity mirror placed between the laser input and the beam splitter is needed in the interferometer set up for power recycling, as shown in *Figure 1.8*. This extra mirror reflects light, that would be exiting the instrument along the input port, back into the arm cavities in such a way that it is added coherently with the fresh laser light being injected into the detector. This increases the overall laser power in the the interferometer arm cavities and therefore enhances the gravitational wave signal. The power recycling mirror in LIGO increased the overall input laser power by a factor of about 40 [41].

1.4.5 Signal Recycling

Signal recycling [40] is somewhat analogous to the power recycling method described previously. A partially reflecting mirror is placed between the dark port and the beam splitter that reflects light back into the interferometer, as shown in *Figure 1.9*. As mentioned in an earlier section, when a gravitational wave passes through an interferometer there is a resulting phase change in




Figure 1.9: A schematic of an interferometer with a signal recycling mirror before the photodetector. The signal recycling mirror can be tuned to increase the detector's sensitivity at a particular frequency.

the light that is circulating in the interferometer. In frequency space the phase change is represented by sidebands on the carrier signal, which in this case is the laser signal. The main carrier signal is resonant with the Fabry-Pérot cavities, therefore, it remains in the arms of the detector. The sidebands created by a gravitational wave signal are non-resonant with the Fabry-Pérot cavities and so they exit the interferometer through the output port. By including the signal recycling mirror in the interferometer set up an additional resonant cavity is created which keeps the gravitational wave signal circulating in the detector. This can greatly enhance the gravitational wave signal. The position of the signal recycling mirror can be adjusted to improve the sensitivity of the instrument at different frequencies of interest.

1.4.6 Current Detectors

As mentioned in the introduction to this chapter, there is a worldwide network of gravitational wave detectors. The largest of these detectors are the Laser Interferometer Gravitational-wave Observatory (LIGO) detectors in the U.S..

There are two LIGO facilities, one is located near Hanford in Washington state which houses one 2 km and one 4 km detector and the other is located in Livingston, Louisiana, which houses just one 4 km detector. The three LIGO interferometers utilised Fabry-Pérot cavities and power recycling and reached a peak sensitivity of $2 \times 10^{-23} \text{ Hz}^{-\frac{1}{2}}$ during LIGO's fifth science run [30]. LIGO underwent a partial upgrade, to enhanced LIGO, from 2007 to 2009 as a preliminary test of some of the technologies that would be employed in the second generation detectors: advanced LIGO (aLIGO). The upgrades will be discussed briefly in Section 1.4.7.

There are two detectors located in Europe, one is located near Hannover in Germany which has 600 m arms and is a German-British collaborative effort called GEO600. GEO600 does not have Fabry-Pérot cavities in its interferometer set up, however, it was the first kilometre scale interferometric detector to utilise monolithic suspensions [4] and dual recycling [42], which consists of both power and signal recycling. GEO600 also has folded interferometer arms so that it has an effective arm length of 1.2 km. GEO600 has already been upgraded to its second generation status, GEO-HF, and will be discussed in Section 1.4.7.

There is a detector near Pisa in Italy called Virgo which has 3 km arm lengths; this is an Italian-French collaboration. The Virgo interferometer has a similar set up to the LIGO interferometers as it also has Fabry-Pérot cavities and power recycling. One of the unique features of Virgo is its seismic isolation system, called the super attenuator [33], which gives Virgo superior low frequency performance. Virgo is undergoing a partial upgrade, to Virgo+, to test the technologies that will be implemented in advanced Virgo.

There is one detector in Japan, TAMA300, which is in Tokyo and is the smallest of this network of detectors with arm lengths of 300m. TAMA300 is also the least sensitive of the current network of detectors, most likely because it is located in a highly populated area. Like LIGO and Virgo, TAMA300



Figure 1.10: *Photographs of some of the first generation gravitational wave detector sites. Starting from the top left image and going clockwise the LIGO Livingston, Virgo, TAMA and GEO600 sites are shown. The arms of the TAMA detector are highlighted in red as they are difficult to see.*

uses Fabry-Pérot cavities and power recycling in its interferometer set up. TAMA300 will not be upgraded to a second generation status like the other detectors discussed in this section, however, the work conducted at TAMA300 will be used in the construction of the Large-scale Cryogenic Gravitational-wave Telescope (LCGT).

The world wide network of detectors, shown in *Figure 1.10*, has taken part in many collaborative science runs where the results obtained from each detector were compared in order to rule out local noise sources. The advantage of having a worldwide network of detectors is that if a gravitational wave is

detected the signal from each detector can be used to triangulate the position of the source [20]. This information can then be passed on to other observatories that detect electromagnetic radiation to confirm the source of the gravitational waves [43].

Even though the first generation detectors did not have any confirmed detections of gravitational waves, a significant amount of useful astronomy was obtained from the science runs. The data from these science runs were used to publish many articles on the upper limits of the strength of gravitational waves created by individual sources and the gravitational wave background [22, 44]. It also gave scientists the opportunity to better understand the limitations of the instruments. This enabled the gravitational wave community to improve upon certain aspects of the detectors in order to design advanced detectors, or the second generation of gravitational wave detectors.

1.4.7 Second Generation Detectors

The second generation of gravitational wave detectors will mainly consist of first generation detectors which have undergone an upgrade to improve their sensitivity. The LIGO, Virgo and GEO600 detectors will all be upgraded to an advanced detector status and there will be a brand new detector site located in the Kamioka mines in Japan that will house LCGT. These detectors should have a very good chance of detecting gravitational waves as the detection rate of the instruments should increase from less than 1 per year to several detections per year [45].

The upgraded LIGO detectors, aLIGO, will see their sensitivity increase by at least a factor of 10 which will increase the volume of space that the detectors can probe by at least 1000 [6]. The main features of aLIGO, which will be critical for the detector reaching its improved sensitivity, include improved seismic isolation, fused silica monolithic suspensions, improved optical coatings, increased laser power and signal recycling. The aLIGO upgrade is due to

be complete in 2015 and first observations will be taken [46]. The detector is expected to reach its design sensitivity in 2019 [46].

The key features of the upgrade to Virgo, advanced Virgo [7], are similar to aLIGO. Advanced Virgo will make use of fused silica suspensions, signal recycling and higher laser power to improve the sensitivity of the detector. The advanced Virgo upgrade should also be completed in 2014 so that advanced detector science runs can begin.

The main focus of the upgraded GEO600 detector, GEO-HF [8], was to reduce high frequency noise sources such as shot noise. In order to reduce shot noise GEO-HF has a higher power laser and injects squeezed light, discussed in Section 1.5, into the output port of the detector. Once the reduction in shot noise is taken into account, the limiting noise source of GEO-HF at high frequencies is thermal noise from the optical coatings [8]. GEO-HF has the very important task of continuing observations while LIGO and Virgo are being upgraded and will be used as a test bench for novel technologies that are being developed for third generation detectors.

LCGT [47], also known as KAGRA, will be the first kilometre scale baseline interferometer to use cryogenics to cool the test masses and suspensions in order to lower the thermal noise of the detector. LCGT will also be the first large scale gravitational wave detector to be located underground. The optical design of the detector consists of a Fabry-Pérot-Michelson interferometer that will also use power recycling techniques. The cooled test masses will be made of sapphire since the thermal properties of sapphire are superior to that of silica at low temperatures. LCGT will draw on the experience gained from TAMA300 and the Cryogenic Laser Interferometer Observatory (CLIO) test facility. Construction of LCGT is expected to take place between 2010 and 2016 and observations are expected to start in 2018 [47].

1.4.8 Third Generation and Future Detectors

Third generation and future gravitational wave detectors will really push the materials research, interferometric design and data analysis techniques of gravitational wave detectors to their limits. The third generation ground based detector that is currently being planned is the Einstein Telescope (ET) and, further in the future, there are plans to put a detector in space called the Laser Interferometer Space Antenna (LISA). An artist's impression of what these detectors may look like are shown in *Figure 1.11*.

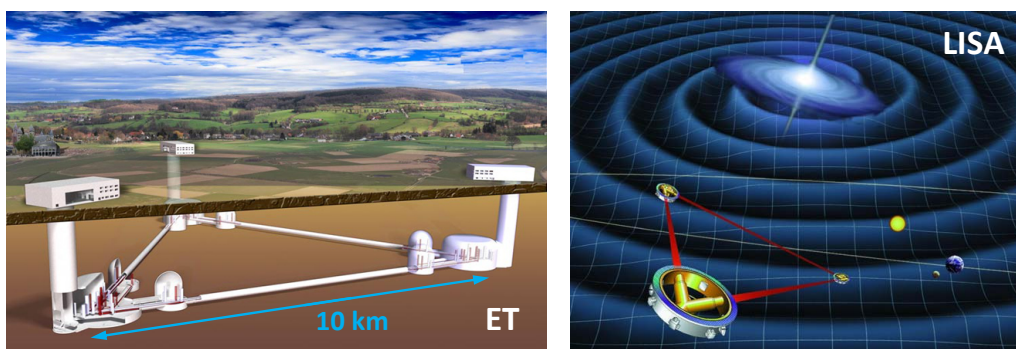


Figure 1.11: Artists' impressions of what ET (left) and LISA (right) could look like.

ET [15] will be a ground based detector that will aim to have a sensitivity at least an order of magnitude better than the advanced detectors discussed in Section 1.4.7. Based on the technology available now, ET will utilise the best features from current and advanced detectors in order to reach the highest possible sensitivity. For example, a similar seismic isolation system to Virgo's super attenuator will probably be used to provide isolation at lower frequencies and squeezed light methods used in GEO-HF may also be used to lower shot noise at high frequencies in ET. One of the most unique features of ET will be its xylophone interferometer configuration, if it is chosen as the final interferometer configuration. This essentially means that ET will consist of two interferometers. One interferometer will have very good high frequency performance, by utilising high laser power and squeezed light to reduce quan-

tum noise at high frequencies. The other detector will have very good low frequency performance, by taking advantage of cryogenics to cool the test masses and suspensions to lower thermal noise, and the detectors proposed underground location to reduce gravity gradient noise. A plot of the estimated ET strain sensitivity, along with other first and second generation detectors, is shown in *Figure 1.12*.




Figure 1.12: A plot showing the increase in strain sensitivity of ground-based detectors with each new generation of detectors. It should be noted that AURIGA [48] is a bar detector, not an interferometric detector. This plot is reproduced from [49].

LISA [50] is planned to be the first space based gravitational wave detector. The main advantage of putting a detector in space is that it is much more sensitive to low frequency sources, such as black hole-black hole binaries, due to the fact that gravity gradient noise and seismic noise are not an issue in space. The arm lengths of terrestrial detectors are hindered by infrastructure costs and the curvature of the Earth, whereas LISA will have arm lengths of an astounding 5×10^9 m. It is thought [50] that LISA may be able to measure strains in space-time as low as $10^{-24} \text{ Hz}^{-\frac{1}{2}}$. Much of the technology that will be used in LISA will be tested in the LISA pathfinder mission which is scheduled

for launch in 2015.

It should be noted that at the time of writing this thesis LISA had changed from being a joint NASA and ESA mission to just an ESA mission due to programmatic and budgetary reasons. This means that there will be some redesign of LISA in order to account for the reduced funding available. Amaro-Seoane et al. [51] describe the major changes from the original LISA mission to the new evolved LISA (eLISA)/New Gravitational Wave Observatory (NGO) mission. eLISA will still reach very good sensitivity at low frequencies, however, it will not be as sensitive as the original LISA design due to reduced arm lengths. Three spacecraft will still be used to form the detector but interferometric measurements will only be made along two arms, and not three arms as shown in the image of LISA in *Figure 1.11*, to reduce cost and save weight for the launch.

1.5 Limiting Noise Sources of Ground Based Detectors

There are numerous noise sources that restrict the sensitivity of gravitational wave detectors, but the three most limiting noise sources are seismic noise (which comprises of both terrestrial seismic noise and gravity gradient noise), thermal noise and quantum noise (which includes shot noise and radiation pressure noise). Seismic noise arises due to low frequency vibrations caused by ground motion and changes in the local gravitational field. Thermal noise occurs due to the fact that the atoms in the detector optics and suspensions have a thermal energy of $k_B T/2$ per degree of freedom and therefore are never truly completely at rest. Quantum noise is comprised of both the statistical fluctuation of photons at the photodetector and the noise created by photons imparting momentum to the detector mirrors. A plot of the main noise sources in aLIGO are shown in *Figure 1.13* to show the combined impact of these

limiting noise sources on a gravitational wave detector.

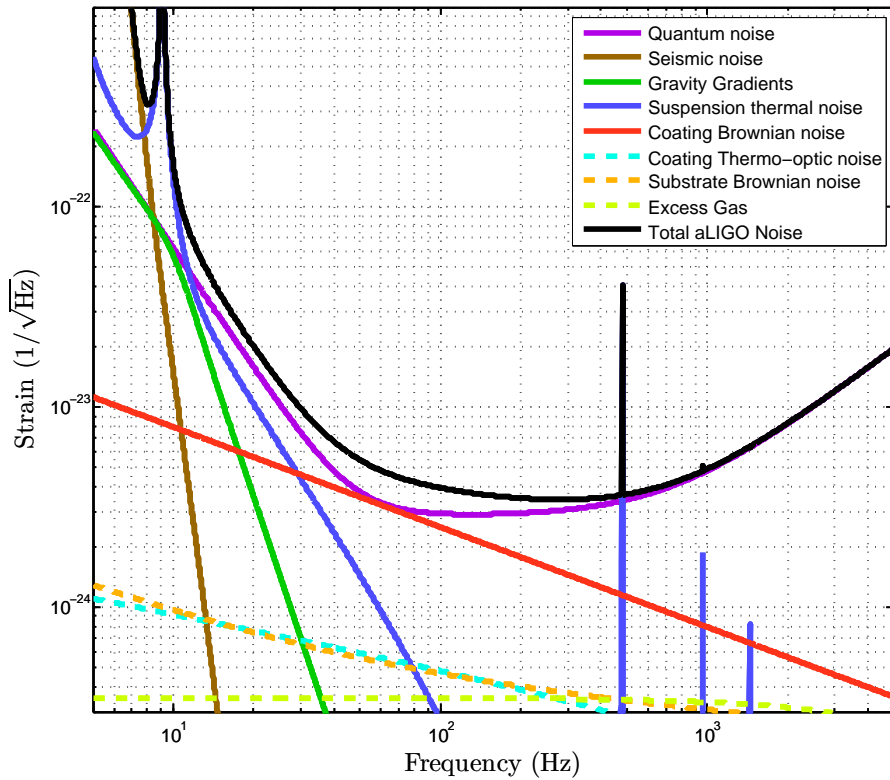


Figure 1.13: A plot of the estimated advanced LIGO sensitivity curve based on calculations of the limiting noise sources. This plot was created using the Gravitational Wave Interferometer Noise Calculator (GWINC) [52].

1.5.1 Gravity Gradient Noise

Gravity gradient noise is the only noise source that cannot be suppressed by making changes, or improving, the interferometer set up. This means that gravity gradients set a limit on the low frequency range (tens of Hz) of ground based detectors. The fact that the density of the ground and air around the detector site is time dependant means that there will be fluctuations in the local gravitational field which can translate into a displacement noise in the detectors. In some cases even the motion of vehicles and people near the detector can induce detectable changes in the local gravitational field [53]. Fluctuations in the density of the ground arise due to seismic waves, which

are always present, in the upper layers of the Earth's crust. The only way to minimise gravity gradient noise is to locate a detector a few hundred metres underground and choose a site that has low seismic activity [15, 54]; this is what is planned for ET. Locating a detector in space, like the LISA detector, would completely remove gravity gradient noise and would allow for the detection of sources in the frequency range 0.1 mHz – 1 Hz [50].

1.5.2 Seismic Noise

At frequencies between 1 Hz and 10 Hz the interferometer noise is dominated by seismic noise. Seismic noise is more unpredictable than other noise sources as it can arise due to earthquakes, stormy weather and human activity such as people walking at the detector site or heavy traffic on a nearby road. It is therefore very important that the interferometer optics are suspended from sophisticated vibration isolation stacks and multiple pendulum suspension systems to damp any seismic motion.

If a gravitational wave detector is to be effectively isolated from seismic disturbances it must have both horizontal and vertical isolation. Horizontal motion is damped through the use of multiple stage pendulum suspensions. Saulson shows [20] that the transfer function of a single pendulum suspension of mass m and resonant frequency f_0 that is connected to the ground through its suspension will have the form,

$$\frac{x}{x_g} \approx \left(\frac{f_0}{f} \right)^2, \quad (1.23)$$

assuming damping is negligible. In Equation (1.23), x is the displacement of the mass, x_g is the ground motion and f is the frequency of interest. This is for the case where $f \gg f_0$. If there are now N cascaded pendulum stages the transfer function becomes,

$$\frac{x}{x_g} \approx \left(\frac{f_0}{f} \right)^{2N}. \quad (1.24)$$

This shows that increasing the number of pendulum stages in a suspension system can increase the attenuation of horizontal seismic motion. Vertical motion is damped using multiple stages of maraging (precipitation hardened) steel blade springs [33, 55, 56]. Attenuation of the vertical motion can also be increased by the use of multiple stages [20].

1.5.3 Thermal Noise

Thermal noise is the most dominant noise source in the frequency range of tens of Hz to a few hundred Hz, which is the detector's most sensitive frequency range. Thermal noise is associated with the thermal excitation of the atoms that make up the mirrors and suspensions, and from internal friction that occurs within the material of the mirrors, coatings and suspensions [57]. The atoms that make up the mirrors and suspensions are never truly at rest and always have a thermal energy of $k_B T/2$ per degree of freedom associated with them. The thermal motion of the atoms can excite resonant modes of the test masses which results in a significant displacement noise in the detector.

In order to reduce thermal noise in gravitational wave detectors, ultra high quality materials such as fused silica are used as the test mass material and, in some cases, as the suspension material [58, 59]. This means that most of the energy of the thermally excited internal resonant modes will be stored near the resonances of the mirrors and suspensions, thus the energy associated with thermal excitations far from the resonance modes, where the operation band of the detector is situated, will be greatly reduced.

With the use of such high quality materials for the test masses and suspensions the main source of thermal noise arises from the highly reflective optical coatings on the test masses. The coating thermal noise can be reduced by doping the coatings with certain materials. For example, thin film multi-layer tantalum pentoxide/silica coatings, which were used in all first generation detectors, have been found to have lower mechanical loss if they are doped with

titania [60]. It is still not well understood why doping the mirror coatings with certain materials has an effect on the mechanical loss of the coating, however, the answer is thought to lie in the effect that the doping material has on the microstructure of the coating [61].

The thermal noise can be reduced further by cooling the mirrors with cryogenics, this is what it planned for LCGT and ET. However, new materials have to be investigated for use in these detectors since mechanical loss can change significantly with temperature [62].

1.5.4 Photon Shot Noise

As described earlier in this chapter, a gravitational wave detector can detect a signal from a passing gravitational wave by measuring changes in the power of the laser light exiting the interferometer at the output port. However, there is a limit to how small a change in optical power can be observed due to the statistical fluctuation of photons at the photodetector of the interferometer. Poisson statistics dictates that if N discrete independent events are counted, such as photons being counted at a photodetector, there will be an error of \sqrt{N} attributed to the measurement.

The strain noise created by photon shot noise can be described by [63],

$$h_{\text{shot}}(f) = \frac{1}{L} \sqrt{\frac{hc\lambda}{\pi^2 P_{\text{in}}}}, \quad (1.25)$$

where L is the detector arm length, h is Planck's constant, c is the speed of light, λ is the wavelength of the laser light and P_{in} is the power of the laser light entering the interferometer. Equation (1.25) shows that by increasing the input laser power it is possible to reduce shot noise. This is why it is of such great importance to utilise advanced interferometric techniques, such as power recycling, and use higher powered lasers to increase P_{in} . Photon shot noise is the dominant detector noise source at frequencies greater than several hundred Hz.

1.5.5 Radiation Pressure Noise

One must be careful when increasing the input laser power of an interferometer as this will increase another component of quantum noise: radiation pressure noise. Radiation pressure noise arises due to the fact that as photons bounce off the mirrors of the interferometer they will impart momentum to the mirrors causing them to be displaced. If the input laser power is increased there will be more photons circulating in the interferometer arm cavities and, therefore, the displacement noise due to radiation pressure will increase. The strain noise created by radiation pressure noise can be described by [20],

$$h_{\text{rp}}(f) = \frac{N}{mf^2L} \sqrt{\frac{\hbar P_{\text{in}}}{2\pi^3 c \lambda}}, \quad (1.26)$$

where N is the number of times a photon travels a distance of $2L$ in the interferometer arm cavities, L being the length of the interferometer arms, and m is the mass of the interferometer mirrors.

1.5.6 The Standard Quantum Limit

The Standard Quantum Limit (SQL) of a gravitational wave detector is defined as the point where the laser power circulating within the detector is such that it optimally minimises both photon shot noise and radiation pressure noise, or the regime where $h_{\text{shot}} = h_{\text{rp}}$. First generation detectors did not reach the SQL as many of them lacked the sensitivity required to reach this limit, however, it is believed that some second generation detectors may be sensitive enough to reach the SQL. This would place a fundamental limit on the sensitivity of future detectors, however, research has shown that the SQL may be able to be beaten by using squeezed light techniques [64].

Sections 1.5.4 and 1.5.5 discussed the quantum noise of gravitational wave detectors, however, this was a very general overview of these noise sources. A more complete picture of how these noise sources actually arise will be

briefly discussed here. Shot noise and radiation pressure noise occur due to vacuum fluctuations [65], which is when photons are randomly created and destroyed in the vacuum field. As vacuum fluctuations enter the interferometer at the output port, fluctuations in the phase quadrature, or the imaginary part of the detected signal, are responsible for shot noise and fluctuations in the amplitude quadrature, or the real part of the detected signal, are responsible for radiation pressure noise. Due to these random fluctuations the phase and amplitude of light in the arm cavities of the detector can't both be measured with absolute accuracy. This is like Heisenburg's uncertainty principle but applied to electromagnetic fields and, just like in the uncertainty principle, it is only possible to reduce the uncertainty of one component while increasing the uncertainty of the other.

It is possible to create what is known as squeezed light that has the uncertainty in one quadrature reduced at the expense of the other quadrature, see *Figure 1.14*. This means that, depending on the frequency of interest, the phase or amplitude uncertainty can be reduced in order to reduce the relevant noise source below its expected noise level without changing the laser power in the detector. Implementation of squeezed light techniques have been successfully demonstrated at the GEO600 detector for the frequency range 10 Hz to 10 kHz [66]. As a result of this work, squeezed light is being considered for future gravitational wave detectors such as aLIGO [6] and ET [15] in order to beat the SQL.

1.6 Conclusions

When gravitational waves are regularly detected at the various gravitational wave observatories around the world they will not only yield some fascinating new physics, but also prove to be a very useful astronomical tool. The first generation of detectors did not directly observe gravitational waves although detection was never guaranteed with these detectors due to low event rates.




Figure 1.14: A diagram of the uncertainty that arises in (a) a normal coherent state compared with states where squeezing is applied to (b) the amplitude quadrature and (c) the phase quadrature. The plots on the left hand side show the uncertainty in the phase and amplitude of an electromagnetic wave and the plots on the right hand side shows the phasor diagrams of the waves on the left hand side. The uncertainty is represented by the blue shaded region on the plots. As stated in the uncertainty principle, the total error, or total area of the blue shaded region, remains the same even though the error of each component changes.

The second generation detectors stand a very good chance of making the first detections of gravitational waves and should reach sensitivities that will allow regular detections to be made. This will begin a new era of gravitational wave astronomy.

In order for future detectors to reach sensitivities capable of making observations of gravitational waves a significant amount of current research is focused on suppressing the limiting noise sources of the detectors. Innova-

tive new technologies that have been investigated and developed should allow future detectors to overcome the limitations of their predecessors.

The research contained within this thesis focuses on investigations of charging noise. Charging noise was not mentioned in Section 1.5 as it is not considered one of the limiting fundamental noise sources in current ground based detectors, however, it will become apparent in Chapter 2 that, if the correct precautions are not taken, charging noise could become a limiting noise source in future detectors. The following chapters will discuss methods of mitigating charge from the detector optics, methods of cleaning the test masses that will minimise the amount of charge transferred to the optic during cleaning and, most importantly, measuring this noise so that the nature of charging noise can be fully understood.

Chapter 2

Charging Noise in Gravitational Wave Detectors

2.1 Introduction

One noise source that is believed to become more prominent in future gravitational wave detectors is noise due to excess electrostatic charges on the detector optics. Although the detector mirrors are made of silica, which is an insulator, the surface will have a very small amount of conductivity, therefore, charge on the surface of the mirrors will redistribute itself across the mirror's surface, albeit very slowly. As the charge moves around on the detector optics it creates fluctuating electric fields which gives rise to time varying forces which act on the detector optics causing them to be displaced. It is thought that noise due to excess charge could potentially limit the sensitivity of future detectors at frequencies less than 100 Hz. For example, the sensitivity of aLIGO could potentially be reduced by over an order of magnitude at 20 Hz if the test masses came into contact with Viton (see Section 2.5).

A confirmed observation of charging noise in a gravitational wave detector has still to be made, however, other problems due to excess surface charges on the detector optics have been encountered in one of the first generation detectors. An incident at the GEO600 detector [67], where a test mass made

contact with a high voltage connection to an electrostatic drive (ESD) on a reaction mass, showed that the positional control of the detector mirrors can also be jeopardised by excess surface charge.

Some observations [68] even suggest that it is possible that the mechanical Q of the detector mirrors could be significantly affected by surface charge. This could therefore increase the thermal noise level in the detector. An investigation carried out by Mortonson et al. [69] showed that the charge itself did not have an effect on the mechanical Q of their fused silica samples, however, it is likely that Coulomb damping from the charge attracting dust to the surface of the samples was the reason for the lower Q observed in the study. Another danger of dust, and other particles, being attracted to the surface of the mirrors by electrostatic charges is that this would reduce reflectance and increase absorption and scattering of laser light. With laser light of such high power circulating in the arm cavities of the detector it is likely that the mirrors would be damaged if dust on them absorbed some of the laser light.

Charge can be transferred onto the detector optics by contact electrification from numerous sources [70]; dust abrasion on the surface of the optics as the vacuum chamber is being pumped out, contact with nearby structures such as earthquake stops and cleaning of the optics. Another potential source of charging is due to cosmic rays hitting the vacuum chamber walls and showering the test masses in electrons. An account of the aforementioned sources of charging will be discussed in this chapter.

The potential of excess surface charge to create noise, and other problems, in future gravitational wave detectors has led to much research into different in situ charge mitigation methods. These methods include using UV radiation [71], combined electron and ion beams [72], conductive coatings [73], ionising nitrogen gas using an electron gun [74] and ionising boil off gas from liquid nitrogen using corona pins [75]. A summary of these methods will be given in this chapter and a detailed discussion of mitigating charge using glow and

corona discharges [76] will be given in chapters 4 and 5.

2.2 Contact Electrification Between Different Materials

The transfer of charge between two different materials is a phenomenon that has been recognised for thousands of years [77]. For a phenomenon that is so well known it may seem surprising that physicists are still unclear about exactly how, or why, two materials actually transfer electrical charges when they are brought into contact. This section will give a brief overview of the physical mechanisms that are involved in the transfer of charge during contact electrification between two metals, a metal and an insulator and two insulators.

2.2.1 Contact Electrification Between Metals

The way in which charge transfers between two metals is a well understood process. The theory of contact electrification between two metals is based on the premise that the transfer of charge is such as to bring the two metals into thermodynamic equilibrium [78, 79]. This essentially means that as the two metals come into thermodynamic equilibrium their Fermi levels equalise.

When two metals are brought into contact electrons flow from the metal with the lower work function of the two metals to the metal with the higher work function. Electrons flow to the metal with the higher work function in order to occupy electron states which are at a lower energy than the electron states that they currently occupy. As the metal with the higher work function gains electrons its potential rises above the potential of the other metal. The difference between the work functions of the two metals is proportional to the potential difference that arises between them, while they are in contact, which is called the contact potential difference, V_c , and can be expressed as,

$$V_c = \frac{(\phi_2 - \phi_1)}{e}, \quad (2.1)$$

where ϕ_1 and ϕ_2 are the work functions of metals 1 and 2, which are shown in *Figure 2.1*, and e is the charge of an electron.

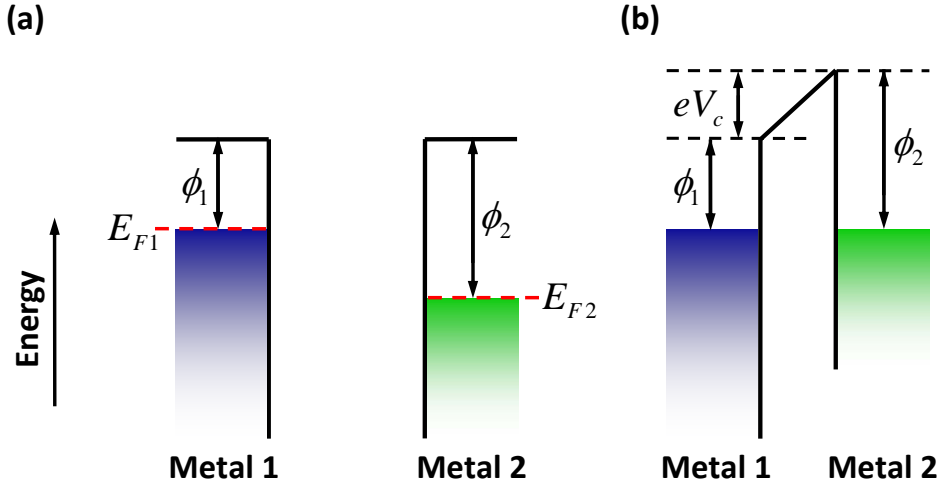


Figure 2.1: A diagram showing how a contact potential arises between two metals in physical contact. The shaded regions are the filled electron states of the metals and the top of these filled states, indicated with dashed red lines, are the Fermi levels of the metals. Diagram (a) shows the metals before contact and diagram (b) shows the metals during contact.

As the two metals are separated the capacitance between them will decrease and therefore the potential difference between them will increase. As the potential difference between the two metals increases away from V_c electrons will tunnel between the metals to try and maintain thermodynamic equilibrium while the separation distance is appropriately small. This highlights that the contact electrification process is not completed until both the metals are separated to a sufficiently large distance at which point electrons can no longer tunnel between the metals.

The residual charge, Q , left on the metals after separation can be calculated, to a good approximation, from V_c and a constant, C_0 , which depends only on the shape of the contacting metals [79],

$$Q = \frac{(\phi_1 - \phi_2)}{e} C_0. \quad (2.2)$$

2.2.2 Contact Electrification Between Metals and Insulators

In the electronic band structure of metals and insulators they both have a valence band, which consists of mainly filled electron states from the valence electrons of the atoms in the material, and a conduction band, which consists of mainly empty electron states which can be occupied by electrons if they have sufficiently high enough energy or all other states below the conduction band are filled. The valence and conduction bands of an insulator are separated by a band gap. Ideally electrons would not occupy the band gap due to the fact that there are no states available there to occupy, however, it is possible for some localised electron states to reside in the band gap. This is shown in *Figure 2.2*. These states are thought to arise due to impurities or defects in the insulating material. It is these states in the band gap that are thought to be what makes contact charging involving insulators possible [78, 80]. These states are distributed over a range of energies with states of lower energy generally filled and states of higher energy generally empty. In this model it is also assumed that an electron occupying a given state cannot transfer to another energy state even if it is thermodynamically favorable [80]. This is due to the fact that these energy states are localised in such a way that the wavefunctions of neighboring states do not overlap.

As with metal-metal contact electrification, metal-insulator contact electrification has also been observed to be correlated with the work function of the metal involved in the process [78]. Evidence suggests that the charge transferred between a metal and an insulator depends on the work function of the metal and some energy level, E_0 , that is characteristic of the insulator. This energy level can be thought of as the Fermi level of the insulator [78, 79]. Like the charge transfer between two metals, if a metal and an insulator are brought into contact electrons will, generally, flow from the material with the higher Fermi level to the one with the lower Fermi level, as shown in *Figure 2.3*.




Figure 2.2: A diagram showing the localised states that arise in the band gap of an insulator. A filled state is represented by a circle with a dash through it and an empty state is indicated with a dash.

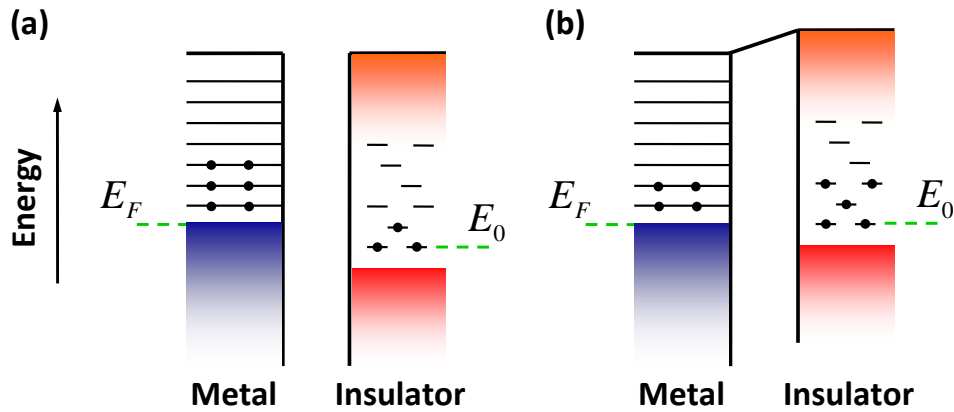


Figure 2.3: A diagram showing the charge transfer between a metal and an insulator. Energy states are indicated with a dash and electrons are indicated with a black circle. Diagram (a) shows the materials before contact and diagram (b) shows the materials after contact.

In a review article by Lowell and Rose-Innes [78] they mention many investigations that show the charge transfer between some insulators and metals to have a linear relationship with the work function of the contacting metal, much like the case for metal-metal contacts. However, they also highlight that there are some cases where different relationships have been observed. The different results that have arisen from different investigations show that there are added complications due to the involvement of insulators. This is due to the fact that

with insulators there may be more charging mechanisms, other than electron transfer, playing a part in the total charge transfer between the two materials. These other charging mechanisms will be discussed in Section 2.2.4.

2.2.3 Contact Electrification Between Insulators

The primary mechanism that dictates how charge is transferred between two insulators is still not well understood. If electrons are the primary charge carriers in the charging process then the charging mechanism described in Section 2.2.2, where localised electron states in the band gap of the insulator are responsible for the transfer of charge, is valid for charging between insulators. For this case the charge transferred is proportional to the difference of the effective Fermi levels of the two contacting insulators. The charge transfer between two insulators is shown in *Figure 2.4*.

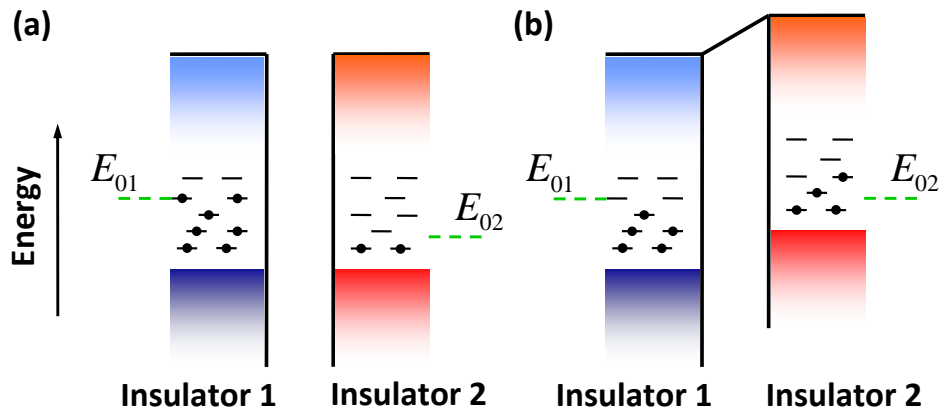


Figure 2.4: A diagram showing the charge transfer between two insulators. Electrons will flow from the insulator with the higher effective Fermi level (E_{01}) to the insulator with the lower effective Fermi level (E_{02}). Diagram (a) shows the insulators before contact and diagram (b) shows the insulators after contact.

If this is indeed the case then it would seem reasonable to assume that materials can be definitively ranked in order of the charge they acquire after being brought into contact with another material. Such a table is often called a triboelectric series, where a material near the top of the series will always acquire a positive charge if it is rubbed with a material lower down in the series.

An example of a triboelectric series is given in Table 2.1 which is adapted from other series given in [81] and [78].

Table 2.1: An example of a triboelectric series.

Positive
Rabbit's Fur
Glass
Wool
Nylon
Silk
Paper
Wood
Amber
Nickel
Copper
Silver
Gold
Polystrene
Acrylic
Polyethylene
Polypropylene
Teflon
Negative

The difficulty in verifying whether materials can be ranked in this manner lies in creating reproducible measurements. Differences in the type of contact (rubbing, sliding, etc), surface contamination, material transfer and sample roughness are all crucial factors that could lead to two samples of the same material occupying different places in the series. At least one investigation has shown that triboelectric “rings” rather than a series can be formed with some materials [82]. According to Shaw and Jex [82] silk was observed to charge glass negatively, and glass deposited negative charge on zinc, however, zinc charged silk negatively.

2.2.4 Complications of Contact Electrification

There is much discrepancy between different studies on contact electrification because there are so many factors to take into account during the process. In some cases ions, rather than electrons, are believed to be the primary charge carrier due to loosely bound ions on the surface of some insulators. Material transfer, type of contact and many other subtle factors have to be taken into consideration during contact electrification experiments.

2.2.4.1 Ion Transfer

One of the biggest unanswered questions in contact electrification research is whether it is electrons or ions that are transferred in the contact electrification process. Harper [79] came to the conclusion that electrons should never be transferred between two materials if those materials are strictly insulators, however, Lowell and Rose-Innes [78] come to the conclusion that it is more likely that electrons are transferred in the contact electrification of insulators. Both works are well respected reviews of contact electrification and even though Harper's work was written nearly 50 years ago, researchers are still no closer to understanding whether electrons or ions are the primary charge carriers in the contact electrification of insulators.

In some situations it has been known for ions to be the charged particles that are transferred during contact electrification instead of electrons [83]. In some insulators there are strongly bound surface ions that are counter balanced by lightly bound surface ions of opposite charge. When the two materials are brought into physical contact the potential wells of the two surfaces merge into one and the lightly bound ions redistribute themselves by diffusion, as shown in *Figure 2.5 (a)*. As the two surfaces are separated a hump arises, shown in *Figure 2.5 (b)*, in the potential well that halts the transfer of charge between the surfaces when the hump is large enough that ions cannot overcome it by thermal excitation. Finally, when the surfaces are at a sufficiently large

distance away the lightly bound ions are trapped in the potential wells of the sample surfaces, *Figure 2.5 (c)*.

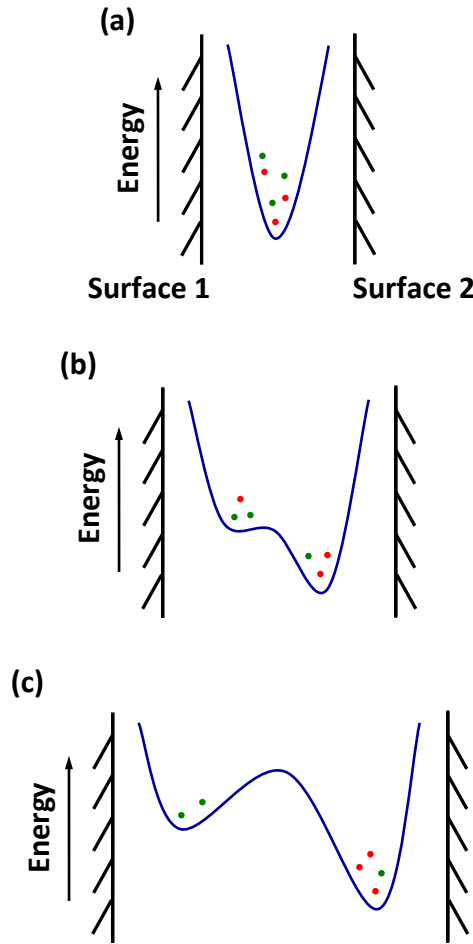


Figure 2.5: A diagram showing how ions are transferred in the contact electrification process. The green dots represent negatively charged ions and the red dots represent positively charged ions. The blue line shows the merged surface potentials of the surfaces. (a) shows the two surfaces in close contact allowing the potential wells to merge and the surface ions of both surfaces to redistribute. (b) shows a hump in the potential well as the surfaces are separated. At this point ions could still overcome the hump by thermal excitation. (c) shows the surfaces at a sufficiently large distance apart that halts charge transfer.

2.2.4.2 Material Transfer

It is possible that as two insulators are brought into contact, material from each of the specimens involved is transferred during the contact process. Other par-

ticles that are foreign to both the specimens, such as dust and other impurities, could also adhere to the surfaces of the samples during the contact process. When material is transferred from one sample to another bonds must be broken to allow this transfer and, therefore, the material being transferred will have charge. Lowell and Rose-Innes [78] state that although there is evidence for the transfer of material during the contact electrification process, the data obtained in investigations of contact electrification is still too reproducible for such a fortuitous mechanism, such as material transfer, to be considered to play a primary role in contact electrification. Material transfer would only become an important factor if there was a large amount of material transferred so that the charge of the transferred material is comparable to the surface charge that is transferred in the process.

2.2.4.3 Other Contributing Factors

Up until this point the most significant mechanisms of charge transfer have been discussed, however, to add even more complexity to the situation there are other contributing factors. It is out of the scope of this thesis to discuss each additional factor in detail but they will be briefly mentioned here for completeness. More detailed information on these additional factors can be found in the works of Harper [79] and Lowell and Rose-Innes [78].

The type of contact (rubbing, sliding, etc) between the two surfaces can have an effect on the charge transferred between the two surfaces. Harder contact can put stress on the material which can lead to surface bonds being broken to produce a negative and positive ion pair which would contribute to the contact electrification process.

Water in the atmosphere can be deposited in a thin layer on the surface of materials. As well as adding a conductive layer onto the surface of a material that may be an insulator, the water will contain H^+ and OH^- ions that could contribute to charge transferred between the two surfaces.

Surface impurities could mean that there are loose bonds present on the surfaces of the contacting surfaces. Loose bonds could contain either electrons or ions that would play a part in the transfer of charge between the two samples.

These added complicating factors make it even more difficult for researchers to come up with a general law for contact electrification.

2.3 Charging Noise

Charge on the surface of insulators does not remain at rest because, in a real life situation, an insulator does not have an infinite resistance. Insulators will always have a very small amount of conductivity and therefore any electrical charge on the surface of the insulator will redistribute across the insulator's surface. When the charge redistributes across the surface of an insulating material the moving charges create a fluctuating electric field. In the situation of a charged gravitational wave detector optic, the fluctuating electric field gives rise to time varying forces that act on the optic creating noise in the detector. If a large enough force is present the sensitivity and control of the instrument could be compromised.

In 1995 Rainer Weiss formulated a theory that describes the behaviour of charging noise [84]. In Weiss's theory it is assumed that the fluctuating forces associated with charging noise can be treated as a Markov process with a single correlation time, τ_0 . A Markov process is a random process in which the future state of a system, X_n , is determined only by the systems most recent previous state, X_{n-1} . The correlation time of the process is the time it takes the system to change from one state to another, i.e. the time for the system to change from state X_n to X_{n+1} . The fluctuating force exerted on the test mass by the moving electrical charges produces a power spectrum, $F_x^2(f)$, given by,

$$F_x^2(f) \approx \frac{2\langle F^2 \rangle}{\pi\tau_0 \left(\frac{1}{\tau_0^2} + (2\pi f)^2 \right)}, \quad (2.3)$$

where f is frequency and $\langle F \rangle$ is the average Coulomb force exerted on the test mass. A full derivation of Equation (2.3) is given in Appendix A. Using Newton’s second law of motion and Equation (2.3) it is possible to derive an equation for the displacement noise that arises from surface charging,

$$x(f) = \frac{1}{m(2\pi f)^2} \sqrt{\frac{2\langle F^2 \rangle}{\pi\tau_0 \left(\frac{1}{\tau_0^2} + (2\pi f)^2 \right)}}, \quad (2.4)$$

where m is the mass of the optic. If an aLIGO optic, which has a mass of 40 kg, had 4×10^{-7} C/m² of charge deposited on its surface by a Viton earthquake stop the displacement noise observed would be $\sim 1 \times 10^{-18}$ m/ $\sqrt{\text{Hz}}$ at 20 Hz. This example uses a τ_0 value of 2.9×10^7 s.

Equations (2.3) and (2.4) show that not only does the displacement noise created by moving surface charge depend on the magnitude of the charge present on the optic but it also depends on τ_0 . It has been shown [70] that τ_0 depends strongly on the cleanliness of the optic. This is likely because surface contamination will introduce a slightly conductive layer to the optic through loose bonds. For a “dirty” optic the value of τ_0 will be small and $x(f)$ will have a frequency dependence of f^{-2} . For a relatively clean optic the value of τ_0 will be large and $x(f)$ will have a frequency dependence of f^{-3} . The highest τ_0 value measured on a clean silica optic was found to be over 8000 hours [85] and is assumed to be a good estimate of the τ_0 value of a clean aLIGO optic.

From measurements of τ_0 [70, 85] the charging noise that one might expect in a gravitational wave detector was estimated and it was found that the noise level caused by charging sits well below the aLIGO noise budget [86]. It is these estimates that have led to the conclusion that charging noise need not be strictly considered a limiting noise source of future gravitational wave detectors, however, it has been shown that if the correct precautions are not taken charging noise could limit aLIGO at frequencies less than 100 Hz.

2.4 Modeling Charging Noise

In order to estimate the level of charging noise that could arise in a gravitational wave detector, for a given surface charge density and correlation time, a simple model based on a technical note written by Lantz [87] was used. In this model the displacement noise contributions from electric charges moving across the surface of the optic and from the time varying force exerted on the optic by a moving earthquake stop are considered. The code which executes the model was written in Matlab.

2.4.1 Defining Surfaces for the Model

The model used to estimate charging noise in chapters 4 and 5 assumes that surface charge on the gravitational wave detector optic is interacting with the earthquake stops of the detector. Earthquake stops sit approximately 1 mm from the surface of the optic in order to limit the excess motion of the mirror during an earthquake. The surface of the optic and the earthquake stop are initially defined in the Matlab model. The tip of an earthquake stop that will be used in aLIGO will be approximately 3.54 mm in diameter and have an area of 9.84 mm². For simplicity, the area of the earthquake stop tip is approximated to that of a square with sides of length 3.14 mm. The earthquake stop surface is assumed to interact with a 10 mm by 10 mm square area on the optic. Any interaction between the surface of the optic and the earthquake stop outside this range will most likely be minimal due to the inverse square nature of the Coulomb force.

The earthquake stop surface and the optic surface have their own coordinate systems. The transverse directions on the earthquake stop are defined as yy and zz and the transverse directions on the optic are defined as aa and bb . The origin of each coordinate system lies at the center of each surface. The axis orthogonal to each of the surfaces will be defined as the x direction for both of the surfaces, or what would be the beam direction in a detector. A diagram

of this set up is shown in *Figure 2.6* for clarity.




Figure 2.6: A diagram of the surfaces of the optic and the earthquake stop defined in the model.

2.4.2 Calculating Noise due to the Movement of Charge on the Optic Surface

Calculating the noise due to the movement of surface charge is fairly straight forward. The total Coulomb force is calculated by integrating it over the surface of the optic. This is done by dividing the surfaces of the optic and the earthquake stop into N smaller square segments. The individual forces exerted by each segment on the earthquake stop are summed to give the total force exerted by the earthquake stop on each individual segment on the optic. The surface charge is assumed to be uniform across both of the surfaces in this model, therefore, the surface charge is divided equally into each of the N segments across the surfaces. The surface charge in each segment of the optic and the earthquake stop is treated as if it were a point charge, for simplicity.

The summed Coulomb force interaction of each point charge in each segment, j , on the optic surface with all the point charges on the earthquake stop tip is calculated using,

$$F_j = \sum_{i=1}^N \frac{Q_j Q_i}{4\pi\epsilon_0 r_{ij}^2} \frac{d_0}{r_{ij}}, \quad (2.5)$$

where Q_i is the point charge on each segment i of the earthquake stop tip, Q_j is the point charge on each segment j of the optic, ϵ_0 is the permittivity of free space, d_0 is the separation distance of optic and the earthquake stop at the origin of the two coordinate systems and r_{ij} is the distance between segment i on the earthquake stop surface and segment j on the optic. r_{ij} is expressed as,

$$r_{ij} = \sqrt{(yy_i - aa_j)^2 + (zz_i - bb_j)^2 + d_0^2}, \quad (2.6)$$

where yy_i is the i^{th} segment along the yy axis on the earthquake stop, zz_i is the i^{th} segment along the zz axis on the earthquake stop, aa_j is the j^{th} segment along the aa axis on the optic and bb_j is the j^{th} segment along the bb axis on the optic.

The average Coulomb force acting on the optic is used to calculate the force noise acting on the optic for a range of frequencies using Equation (2.3), but with an added multiplicative factor of $\sqrt{n_{stops}}$. Where n_{stops} is the number of earthquake stops interacting with the charged optic. For the models presented in this thesis it will always be assumed that the force noise due to charge hopping interacts with all eight earthquake stops on the front and back face of the optic.

It should be noted that for this model only the x component of the force is considered. This is because any displacement noise that is picked up in a detector is due to the displacement of the optics in the beam direction. Any force contributions from the y and z components should be minimal.

2.4.3 Calculating Force Noise from a Moving, Charged, Earthquake Stop

In aLIGO the test mass has eight earthquake stops that face the two flat end surfaces of the optic, four of which face towards the highly reflective surface of the mirror and another four that face towards the opposite surface on the mirror. The four earthquake stops that face the reflecting surface of the mirror are attached to the suspension frame which surrounds the test mass. The suspension frame is affected by seismic motion and therefore the position of the earthquake stops attached to it varies with time. If there is charge on the earthquake stops this will produce a time varying electric field at the surface of the optic which will result in displacement noise. The other four earthquake stops are attached to the reaction mass chain which sits behind the test mass. The reaction mass is seismically isolated in the same way as the test mass, therefore, these earthquake stops should not move significantly in relation to the test mass and will be ignored for this part of the model.

The electric field produced by a uniformly charged earthquake stop at the surface of the optic is calculated by integrating the electric field over the surface of the optic. This is done in the same manner as for calculating the Coulomb force across the surface of the optic, except Gauss's Law, Equation (2.7), is used to calculate the electric field contribution from each segment, i , on the earthquake stop. These contributions are then summed to give the total electric field acting on segment j of the optic,

$$E_j = \sum_{i=1}^N \frac{Q_i}{4\pi\epsilon_0 r_{ij}^2} \frac{d_0}{r_{ij}}, \quad (2.7)$$

where Q_i is the point charge on segment i of the tip of the earthquake stop. Once the electric field acting on each segment, j , on the surface of the optic has been evaluated, the average field acting on the optic surface, E_x , is calculated and used to determine the field gradient.

As mentioned previously, if a charged earthquake stop, which is attached to the suspension frame around the test masses moves, the electric field acting on the surface of the optic will change with time. A simple system will be considered for this part of the model where the surface of the earthquake stop moves from $d_0 = 1 \text{ mm}$ to $d_0 = 1.01 \text{ mm}$ ¹. The field at the two values of d_0 are calculated in order to determine the field gradient at the surface of the optic. This is done by taking the difference of E_x at d_0 and $d_0 + 1 \times 10^{-2} \text{ mm}$ and dividing by the distance between the two positions of the earthquake stop. The field gradient is then divided by the surface charge on the earthquake stop to give the field gradient per Coulomb.

The force noise created by the moving, charged, earthquake stop can be expressed as [87],

$$F_{stops} = \sigma_1 \sigma_2 A_{optic} A_{stop} \frac{\partial E_x}{\partial x Q_1} x_{stops} \sqrt{n_{stops}}, \quad (2.8)$$

where σ_1 is the average surface charge density across the earthquake stop, σ_2 is the average surface charge density across the optic, A_{optic} is the area of the optic surface, A_{stop} is the area of the earthquake stop tip, Q_1 is the total charge on the earthquake stop and x_{stops} is the amplitude spectral density (ASD) of the displacement of the aLIGO suspension frame due to seismic motion. For this model the ASD of the aLIGO suspension frame motion between 1 and 40 Hz was taken as a f^{-1} noise spectrum based on a presentation given by Kissel [88].

Using the same example used to calculate an estimate of the displacement noise of the detector in Section 2.3, the force noise created by a moving earthquake stop can be estimated to be $\sim 6 \times 10^{-16} \text{ N}/\sqrt{\text{Hz}}$. This is small in comparison to the force noise created by the moving surface charge which is $\sim 6 \times 10^{-13} \text{ N}/\sqrt{\text{Hz}}$.

¹It should be noted that this is purely to calculate the electric field gradient. The suspension frame would not normally move such a large distance unless an earthquake was taking place.

2.4.4 Calculating the Total Charging Noise

For the final part of the model the noise created by charge hopping and the movement of the earthquake stops are combined to give the total noise due to charging effects. The force noise due to the moving earthquake stops is added, in quadrature, to the force noise due to charge hopping,

$$F_{total}^2 = F_{stops}^2 + F_{hopping}^2. \quad (2.9)$$

Once this is known it is possible to calculate the total strain noise using Newton's second law, as discussed in Section 2.3. The total strain noise, h_{total} , due to charging is therefore,

$$h_{total} = \frac{1}{L} \frac{F_{total}}{m (2\pi f)^2}, \quad (2.10)$$

where m is the mass of the optic, which is 40 kg for an aLIGO test mass, f is frequency, F_{total} is the total combined noise due to charging and L is the arm length of the detector. For the final strain noise spectrum h_{total} has to be multiplied by two because it will be assumed that the noise is occurring incoherently in all four test masses of the detector.

Following from the examples of the displacement noise and force noise calculations in the previous sections, an estimate of the strain noise that could be observed in aLIGO if the optics came into contact with a Viton earthquake stop is $\sim 3 \times 10^{-22} \text{ } 1/\sqrt{\text{Hz}}$. It should be noted that for this example the noise is slightly less than what is plotted in Section 2.5 because multiple earthquake stops were not considered in this simple calculation.

With any type of finite element analysis, the correct number of mesh elements must be used in order to obtain a realistic model. As the number of elements is increased the model will start to converge. For this analysis it was found that 22.5×10^3 , or 150^2 , elements would give a good estimate of the true charging noise being modeled. To obtain this element value, the model was

executed for different numbers of mesh elements and the strain noise level at 100 Hz was noted. Running the analysis for this many elements takes about 1 hour and 20 mins. A plot of the results from this convergence test is shown in *Figure 2.7*.

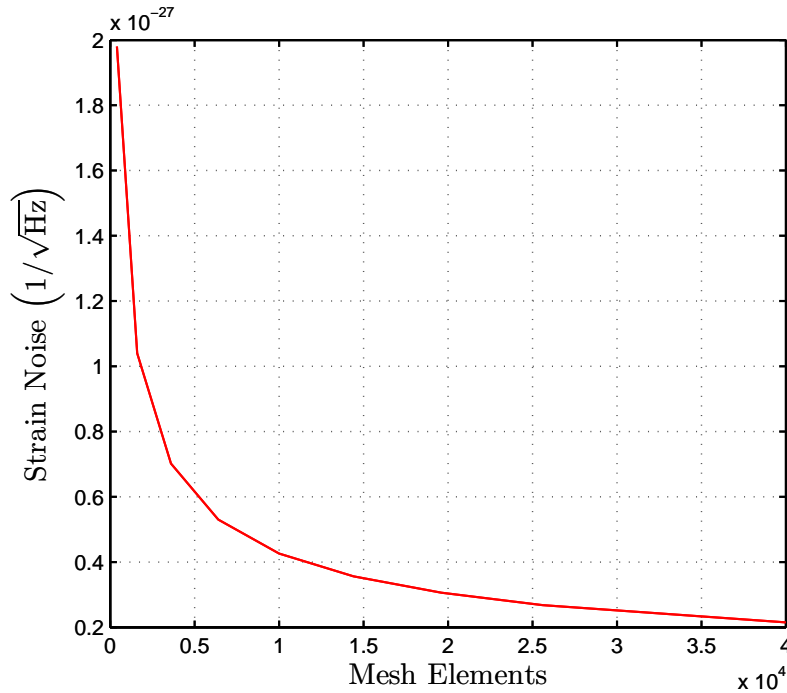


Figure 2.7: A plot of the convergence test data for the charging noise model.

2.4.5 An Example of Modeled Charging Noise

To give an indication of how charging noise compares to the other limiting noise sources of aLIGO, the charging noise from a silica test mass and eight silica tipped earthquake stops, that have just been cleaned with acetone and a clean room cloth, was modeled. A one inch silica disc was cleaned with acetone and a clean room cloth and then placed under vacuum. The charge on a $10 \text{ mm} \times 10 \text{ mm}$ area on the surface of the sample was then measured using a Kelvin probe (see Chapter 3 for more information on the Kelvin probe). A map of the surface charge density on the scanned area of the sample is shown

in *Figure 2.8*.




Figure 2.8: A surface charge density map of a one inch silica disc after it was cleaned with acetone and a clean room cloth. The map was constructed from measurements using a Kelvin probe which will be discussed in *Chapter 3*

The average surface charge density across the surface of the silica sample was found to be $5.36 \times 10^{-9} \text{ C/m}^2$ and this value was taken as the surface charge density across both the modeled silica optic and the silica earthquake stops. A τ_0 value of $1.5 \times 10^7 \text{ s}$ was used for this example. This was the value measured by Ugolini [70] for an optic that has been dry-wiped.

Figure 2.9 shows that, purely from a charging noise perspective, a silica optic and eight silica tipped earthquake stops that have been cleaned with acetone and a clean room cloth will not pose any threat to the sensitivity of aLIGO. However, it will be shown in *Section 2.5* that this noise level could easily increase and compromise the sensitivity of aLIGO if the correct precautions are not taken.




Figure 2.9: A plot of charging noise compared to the total aLIGO noise budget and the main limiting noise sources of the detector. The surface charge density on the surfaces of the optic and the earthquake stops is taken to be $5.36 \times 10^{-9} \text{ C/m}^2$ and τ_0 is taken to be 1.5×10^7 seconds

2.5 Sources of Charging in Gravitational Wave Detectors

2.5.1 Evacuation of the Vacuum Chamber

One mechanism that can deposit charge on gravitational wave detector optics is dust abrasion on the surface of the optics [70]. This occurs when the vacuum chamber that houses the optics is pumped out. Pumping air out of the chamber causes dust to be stirred up and can drag across the surface of the optic depositing charge.

2.5.2 Contact with Earthquake Stops

In initial LIGO the tips of the earthquake stops were made of Viton. This choice in material was very bad from a charging perspective as silica and rubber transfer substantial amounts of charge between each other when the two are brought into contact. *Figures 2.10 and 2.11* show the surface charge density on a silica sample before and after contact with Viton. The average charge density on the sample increases by over two orders of magnitude after coming into contact with a piece of Viton.




Figure 2.10: *Kelvin probe data of the surface charge density across a silica disc before contact with Viton. The average surface charge density across the scanned area before contact is $3.64 \times 10^{-10} \text{ C/m}^2$.*




Figure 2.11: *Kelvin probe data of the surface charge density across a silica disc after contact with Viton. The average surface charge density across the scanned surface after contact is $6.56 \times 10^{-8} \text{ C/m}^2$.*

In 2006 one of the LIGO test masses became stuck against an earthquake stop after some adjustments were made to the LIGO seismic isolation system [89]. Although it was unconfirmed, it was suspected that this was caused by the test mass and the earthquake stop becoming charged after making contact each other. The reason that it is thought that charging had caused the problem is because the test mass became unstuck after air was released back into the chamber. It is possible that the air rushing into the chamber carried excess charge away from the surfaces of the test mass and earthquake stop. Another

reason that it is thought that it is possible that the excess noise was due to charging was because the observed noise had the same frequency dependence as would be expected from charging noise. For a clean optic it is expected that charging noise should have a f^{-3} frequency dependence.

For aLIGO the tips of the earthquake stops will be made of silica since two like materials should transfer minimal charge. There have been investigations into different earthquake stop materials and it was found that silica did transfer the least amount of charge [90]. Silica tipped earthquake stops were incorporated into the set up for enhanced LIGO and were found to be very effective [91]. Using data taken by Ugolini [90], a plot of the charging noise after contact with Viton tipped earthquake stops and silica tipped earthquake stops was made. This is shown in *Figure 2.12*.

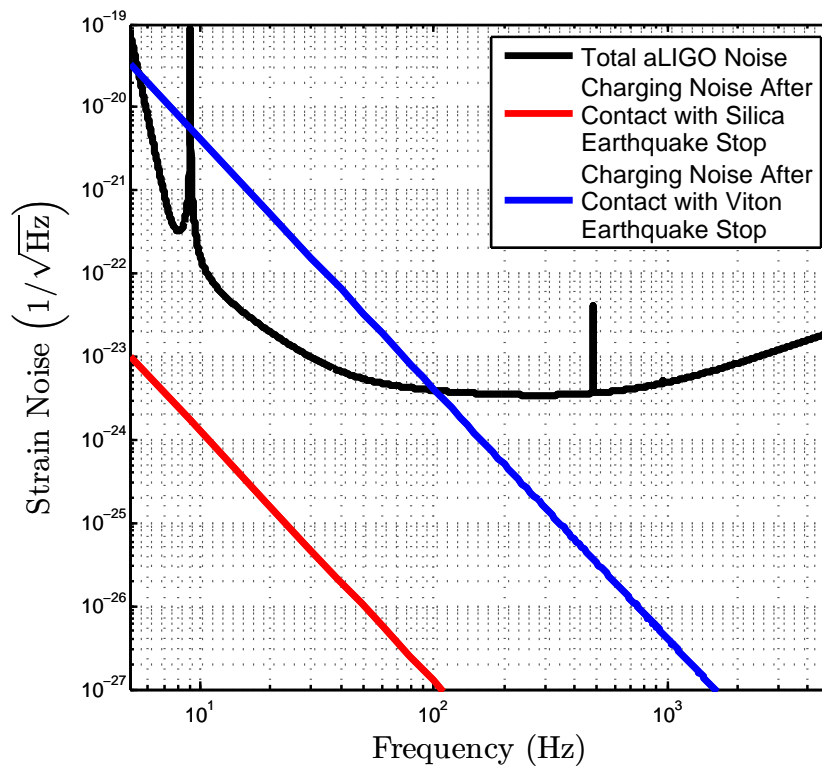


Figure 2.12: A plot of the expected charging noise after an aLIGO test mass has come into contact with a Viton tipped earthquake stop (blue) and a silica tipped earthquake stop (red).

From *Figure 2.12* it can be seen that the charging noise has increased significantly from the noise level shown in *Figure 2.9*. If Viton was to be used in aLIGO as an earthquake stop material the sensitivity of the instrument could be limited at frequencies below 100 Hz. The noise due to contact with the silica earthquake stops is well below the aLIGO noise budget and should not cause any problems in aLIGO.

2.5.3 Cleaning Procedure

It is planned that the aLIGO optics will be cleaned with a product called First ContactTM. First ContactTM is a cleaning and protecting product which is manufactured by Photonic Cleaning Technologies [92]. First ContactTM consists of a polymer which is dissolved in an acetone/ethanol solution. The polymer solution is applied to the surface of the optic and, as the acetone/ethanol solution evaporates, it lifts small particles on the optic into the polymer. When the solution fully evaporates the polymer is left on the optic which dries to form a robust rubbery film. When the polymer film is peeled off it removes excess dirt and dust from the detector mirrors, however, it has been found to deposit a large amount of charge on the mirrors as it is peeled off. The charge deposited can then attract dust and impurities to the surface of the test mass.

A silica sample was cleaned using First ContactTM to demonstrate the increase in surface charge density across the sample. Approximately 0.5 ml was deposited on the surface of the one inch silica disc and spread over the surface using a clean room cloth. Once the sample had dried the First ContactTM was removed from the sample and the sample was then placed under vacuum in order to take a Kelvin probe measurement. This scenario mimics cleaning the aLIGO mirrors with First ContactTM and then pumping out the vacuum chamber containing the optics. *Figures 2.13* and *2.14* show the surface charge density across the silica sample before and after being cleaned with First ContactTM. More results of surface charge deposited by different types

of First ContactTM on silica will be presented in Chapter 5.




Figure 2.13: *Kelvin probe data of the charge on the surface of a silica disc before being cleaned with First ContactTM. The average surface charge density across the scanned area before contact is $8.88 \times 10^{-10} \text{ C/m}^2$.*




Figure 2.14: *Kelvin probe data of the charge on the surface of a silica disc after being cleaned with First ContactTM. The average surface charge density across the scanned surface after contact is $-6.00 \times 10^{-8} \text{ C/m}^2$.*

Using the measured surface charge density after cleaning the silica disc, the charging noise that would be expected in aLIGO was modeled and is shown in Figure 2.15. From Figure 2.15 it would appear that First ContactTM would not limit aLIGO. However, it should be noted that other measurements have shown First ContactTM to deposit at least a magnitude more charge on the surface of silica [93] which would result in an increase in the force noise of two orders of magnitude.

2.5.4 Cosmic Rays

Another possible source of charging is due to cosmic rays hitting off the walls of the vacuum tank and showering the test mass in electrons, as well as causing ionisation of surface atoms on the fused silica test mass [94]. Experiments conducted at Moscow State University showed that a silica mass suspended in vacuum acquired a negative charge at a rate of $10^5 \text{ e}^-/\text{cm}^2/\text{month}$ [95]. The charge was observed to increase in random steps rather than steadily

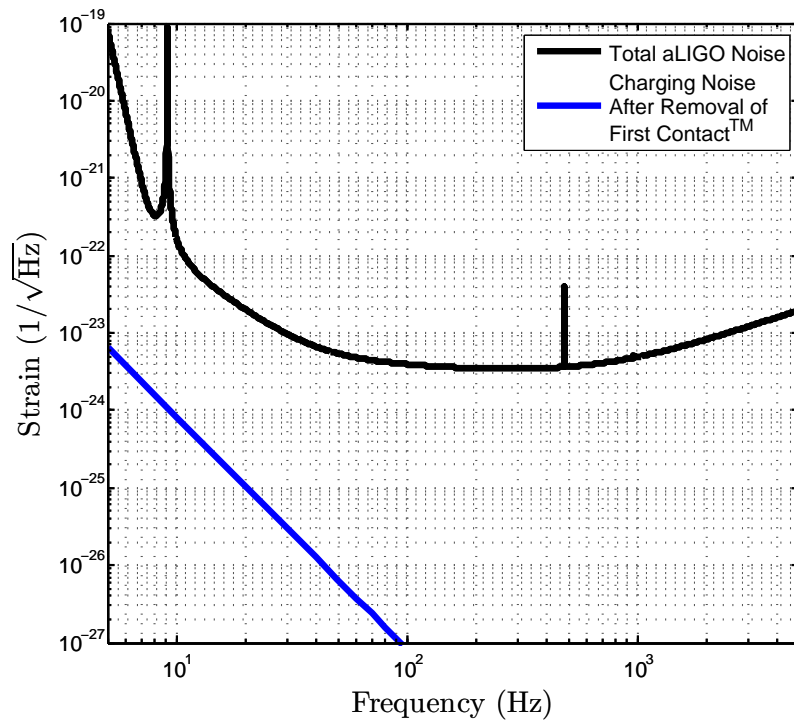


Figure 2.15: *Plot of charging noise after the removal of First ContactTM compared to the total aLIGO noise budget.*

increase over the course of the experiment. These workers also conducted an experiment where they constructed a scintillation detector around the vacuum tank containing their suspended silica mass to see if cosmic ray detections from the scintillator were synchronous with jumps in surface charge on the silica mass [96]. They concluded that there were no statistically significant coincidences between the signals from the cosmic ray detector and the changes in surface charge, however, they point out that there were a few large jumps in surface charge that did coincide with cosmic ray detections which cannot be ignored. This is clearly an interesting area of research which is worth further, more detailed, investigation.

2.6 Possible Surface Charge Mitigation Techniques

2.6.1 UV Illumination

One of the first methods of mitigating surface charge was using UV light since it had been well researched for LISA and LISA Pathfinder [97, 98]. Although this method had been shown to mitigate surface charge effectively on an aLIGO style optic [71], and the method itself would not compromise the pressure in the vacuum chamber, it was found to increase the absorption of the multi-layer reflective coating on the optic [99].

The aLIGO optics have a very strict absorption threshold of less than 1ppm. A low coating absorption is important because if the coatings were to absorb a significant amount of laser light the optics could either be damaged or generate a significant thermal lens [100] which could cause the optical cavity to become unstable.

Further investigation showed that the absorption level could be recovered by annealing the optics [99], however, the optics could not reach the absorption level they had before being illuminated with UV light. It should be noted that in this study the coatings were subjected to far more UV than would ever be used for discharging the optics during a science run.

UV is still under consideration for aLIGO and the most recent development has been to place UV LEDs in the earthquake stops of the detector so that the whole optic isn't exposed to the UV [101]. This would allow the area that has made contact with the earthquake stop to be discharged without illuminating the rest of the mirrors surface with UV.

2.6.2 Conductive Coatings

Conductive coatings could be used to redistribute charge evenly over the mirrors through a thin conductive film on the surface of the mirrors and suspensions [73]. This would provide a passive, in-situ method of charge mitigation.

If the whole of the optic was coated in a thin layer of tin (II) oxide any excess charge would be dissipated quickly through the conductive layer and evenly redistributed over the entire mirror. This method would mean that couplings between localised high density surface charge and nearby grounded surfaces would be greatly reduced. Unfortunately this method would still leave localised charge on the mirrored surface of the optic because the conductive layer would have to sit underneath the highly reflective multi-layer coating. The coating would also have to be of very high quality so as not to increase the thermal noise in the detector. This method of charge mitigation is still under investigation.

2.6.3 Alternating Low-Energy Electron and Ion Beams

A method described by Buchman et al. [72] uses alternating ion and electron beams to neutralise charge on the detector optics. The ion and electron beams would be created by ionising residual gas in the vacuum chamber and would therefore not compromise the vacuum of the detector by introducing more gas into the system. There are a wide variety of technologies available for use as an ion or electron source, however, the investigation undertaken by Buchman et al. [72] suggests that field emission cathodes would be the most suitable for this particular application. Unlike some of the other charge management techniques described in this chapter, the ion and electron beams would be continuously discharging the optics with alternating pulses of the duration of 1 second in order to keep the surfaces of the test masses approximately neutrally charged at all times.

2.6.4 Ionised Gas

Recently there has been a lot of work on discharging silica using ionised gas. One method developed at Trinity University uses an electron gun to ionise a flow of nitrogen gas [74]. The ionised nitrogen then flows over the silica

sample and reduces the charge on the surface of the optic. This method has been shown to reduce surface charge on silica effectively and does not damage the multi-layer reflective coating. The only disadvantage of this method is that the filament of the electron gun only has a finite lifetime. There is the possibility that the filament may break if it is being regularly exposed to gas while hot and may need to be replaced frequently.

A method being developed at MIT uses corona discharge pins to ionise nitrogen gas [75]. The gas is ionised in a high pressure region, by the high electric field generated by the corona pins, and then passes through a small aperture to a lower pressure region where the charged optics are contained. This method has been shown to effectively reduce surface charge on silica and has been tested at the LIGO test facility at MIT, the LIGO Advanced System Test Interferometer (LASTI). This method has also shown that the ionised gas did not effect an aLIGO style multi-layer optical coating at a sensitivity level of 0.2 ppm.

The investigations described in this thesis use a glow discharge to ionise argon gas and a corona discharge to ionise argon and nitrogen gas [76]. The former method would be used in a situation where the detector mirrors have become charged under vacuum, while the latter method would be intended to be used during the cleaning process of the mirrors. These methods will be described in detail in chapters 4 and 5.

The methods described above have been shown to effectively reduce surface charge and have not yet been shown to cause damage to the multi-layer reflective coating on the aLIGO mirrors. There are some concerns that the impact of the charged particles hitting the surface of the mirrors could do some damage, however, this has still to be investigated. The only other disadvantage of these methods is that it would require injecting a small amount of gas into the aLIGO vacuum chambers which would result in some detector down time during a science run.

2.7 Conclusion

A complete theory of how two contacting materials transfer charge between each other is still far from being found. This makes it difficult to determine exactly how much charge may be transferred during contact electrification between a gravitational wave detector optic and an earthquake stop, or during the cleaning procedure of the test masses. However, by taking many measurements of charge transfer between silica and other materials it is possible to obtain a good statistical estimate of what charge would typically be transferred.

To determine the level of charging noise one might expect in a gravitational wave detector, a theory of charging noise, formulated by Rainer Weiss, is used. This theory has never been experimentally verified, however, there have been some, unconfirmed, observations of what is believed to be charging noise in initial LIGO. From analysis of the sensitivity curves of these charging incidents the frequency dependence of the noise observed matched well with what would be expected from Weiss's theory of charging. It is still very important that this theory is experimentally verified so that it is known that the theory is well understood. This will give greater confidence to any charging noise estimates made for aLIGO and future gravitational wave detectors. One possible way of verifying charging noise is by using an instrument called a torsion balance. A torsion balance experiment designed to measure charging noise will be discussed in detail in Chapter 6 and the results of this experiment will be discussed in Chapter 7.

There are many methods of charge mitigation under investigation for aLIGO. At this point it is hard to determine which would be the best method to use as each method has their own distinct advantages and disadvantages. The method chosen for implementation in aLIGO will most likely be the one that poses no risk to the reflective coating on the mirrors and creates minimal detector down time when it is used. The possibility of incorporating multiple charge mitigation methods into aLIGO, and future gravitational wave detectors, should not

be ruled out either.

If the correct precautions are taken, charging noise shouldn't pose too much of a threat to the sensitivity of second generation gravitational wave detectors such as aLIGO. However, these detectors will eventually undergo incremental upgrades which will possibly increase the sensitivity of the detectors by a factor of five, and third generation detectors should expect a sensitivity increase of at least an order of magnitude greater than second generation detectors. At this level, charging noise will likely become a significant limiting noise source. This highlights the importance of investigating and understanding charging noise.

Chapter 3

The Kelvin Probe

3.1 Introduction

The Kelvin probe is a non-contact capacitive device that is widely used to investigate the surface properties of materials. The Kelvin probe is primarily used to measure the work function of metals [102, 103], although it can be used to measure the surface charge on insulators [104, 105]. The idea was first conceived by Lord Kelvin in 1898 [106] and was developed further by W. A. Zisman in 1932 [107] to create the widely used Kelvin-Zisman method of measuring the work function of metals. This method shows that by continuously vibrating a probe electrode near the surface of a metal sample that is in electrical contact with the tip of the probe, and applying a bias voltage to null the measured probe signal, one can obtain the contact potential difference between the probe tip and the metal sample. The work function of the metal under examination can then be calculated from the measured contact potential difference and the known work function of the Kelvin probe tip. This method is still used today, however, developments in atomic force microscopy (AFM) technology in the early 1990s saw the introduction of Kelvin probe force microscopy [108, 109]. With Kelvin probe force microscopy it is possible to measure variations in the work function of materials on an atomic scale.

For the purposes of the investigations presented in this thesis a Kelvin

probe was used to measure the surface charge density on bare fused silica and titania (TiO_2) doped tantala (Ta_2O_5)/silica (SiO_2) multi-layer coated fused silica. TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ is the reflective coating that will be used on the optics of aLIGO [6]. This chapter discusses the working principles of the instrument and details an accurate model of the Kelvin probe output voltage signal. The process of calibrating the Kelvin probe for surface charge density measurements is discussed in detail as this was an important step in estimating the amount of surface charge on the measured samples.

3.2 Experimental Set Up

For the purposes of the experiments described in this thesis a Kelvin probe was used to measure the surface charge density on samples of bare fused silica and fused silica with a doped multilayer reflective coating. The Kelvin probe used in this work was a commercially available Besocke Delta Phi [110] Kelvin probe S which is shown in *Figure 3.1*. The probe consists of a small circular gold tip, that is 2.5 mm in diameter, connected to one end of a thin metal reed. On the other end of the reed, furthest from the probe tip, there is a piezo. As the piezo vibrates, the reed is excited and acts as a vibrating cantilever; this makes the Kelvin probe tip oscillate. A Besocke Delta Phi Kelvin probe 07 control unit was used to apply a high voltage square wave signal to the piezo and to control the oscillation frequency of the piezo, which in turn controls the oscillation amplitude of the Kelvin probe tip. The piezo only oscillates over a very narrow frequency range of 182 – 187 Hz. This is close to the resonant frequency of the metal reed, which is approximately 180 Hz. As the oscillation frequency of the piezo gets closer to the resonant frequency of the metal reed the amplitude of the oscillations of the Kelvin probe tip become larger and the Kelvin probe signal will increase. The AC voltage measured by the Kelvin probe control unit was demodulated to give a DC signal using a Stanford Research Systems SR830 lock-in amplifier. The lock-in amplifier will

also be referred to as a phase sensitive detector throughout this thesis.




Figure 3.1: A photograph of a Kelvin probe similar to the one used for the investigations presented in this thesis. This picture was reproduced from [110]

The Kelvin probe was situated in a vacuum tank, shown in *Figures 3.2* and *3.3*, that can achieve pressures less than 1×10^{-5} Torr using a Varian SH-110 dry scroll pump and Varian Turbo-V 301 navigator pump. Inside the vacuum chamber is an inner structure that houses the Kelvin probe and sample stage motors. The sample stage is attached to a set of Thorlabs NST100 nanostep cross roller bearing stage motors, in an x, y, z configuration, which allows the sample to be moved remotely. The motorised stages each have a dynamic range of 100 mm. To monitor the pressure inside the vacuum chamber a Varian ConvecTorr gauge was used for pressures down to 1×10^{-3} Torr and a Varian IMG-100 inverted magnetron gauge was used for pressures less than 1×10^{-3} Torr.

An automated sample scanning program was developed using LabVIEW so that the Kelvin probe could take a predefined number of scans of a sample automatically. The program is capable of taking scans of an area on a sample's surface or monitoring a single point on a sample's surface for a defined period of time. The first version of this program was developed by the author of this thesis. The program would allow the Kelvin probe to scan across samples




Figure 3.2: A photograph of the vacuum tank where the Kelvin probe is situated.




Figure 3.3: A photograph of the Kelvin Probe and silica sample inside the vacuum tank. The purpose of the copper plate above the probe will be discussed in Section 3.5.2

from left to right, in the x direction, before taking one step upward, in the y direction, and returning back to the left hand side of the sample to start scanning again. However, this was soon found to be a very inefficient way of scanning a sample. One of the stepper motors, the x direction motor, had to be replaced because it had become damaged from being overused during the scanning process. This was most likely caused by overheating while the motor was operating under vacuum. A student, called Noah Hurst, changed the program during the summer of 2009 so that the sample is scanned in a spiral pattern. This meant that the x and y motors were used equally and, therefore, shouldn't become damaged while they are scanning. The previous and current scanning patterns are shown in *Figure 3.4*.

An example of a surface charge density map created from data obtained from Kelvin probe measurements of a silica sample can be seen in *Figure 2.11*. An example of one point on the charged silica disc being monitored over time can be seen in *Figure 3.5*.

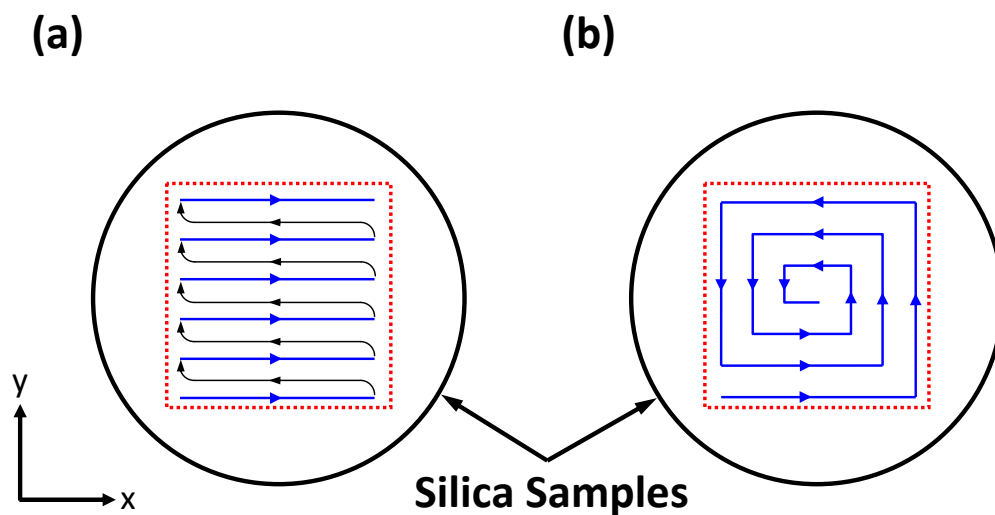


Figure 3.4: A diagram showing the previous (a) and current (b) methods of scanning the surfaces of samples using the Kelvin probe. The area enclosed by the red dashed line represents the scanned area, the blue lines show how the Kelvin probe scanned across the surface of the sample and the black lines on diagram (a) show the additional movement that the x direction motor had to perform in order to get back into position for scanning across the sample again.




Figure 3.5: A plot of the surface charge density on a silica sample changing over time. The Kelvin probe was programmed to take a single measurement every hour for 1000 hours at a point near the centre of the sample.

3.3 Working Principles of the Kelvin Probe

When a sample is brought near the Kelvin probe tip the charge on the sample attracts opposite charge to the tip of the probe, making the probe and the sample act like a pair of capacitor plates. The probe tip oscillates causing a changing capacitance which creates an alternating current that is proportional to the surface potential of the sample. The AC signal is converted to a voltage signal in the Kelvin probe control unit circuitry which is converted into a DC signal using an external phase sensitive detector. A schematic of this process is shown in *Figure 3.6*.




Figure 3.6: A schematic diagram of the Kelvin probe set up. The drive circuit activates the piezo which causes the tip of the Kelvin probe to oscillate. The changing capacitance induces a current which passes through the Kelvin probe preamplifier and the output voltage of the preamplifier is measured by a phase sensitive detector.

If it is assumed that the probe tip and sample are perfect parallel capacitor plates, where the area of the sample being measured by the Kelvin probe is the same area as the probe area, it is easy to show that a time varying capacitance will produce a current, I , that is proportional to the surface voltage of the sample, V_s :

$$\frac{dQ}{dt} = \frac{dC}{dt} V_s, \quad (3.1)$$

$$I(t) = \frac{dC}{dt} V_s, \quad (3.2)$$

where Q is the charge on the probe tip and the surface of the sample, C is the

capacitance between the probe tip and the surface of the sample and t is time.

The capacitance in Equation (3.2) can be substituted for an expression that depends on the area of the probe tip, A , and the distance between the probe tip and the surface of the sample, d_0 . A minor complication is that the distance between the probe and the sample is constantly changing since the probe oscillates by a distance d_1 about the mean probe-sample separation, d_0 . This is easily compensated for by adding a modulation factor of $d_1 \sin \omega t$ to the probe-sample separation distance.

$$C = \frac{\varepsilon_0 A}{d_0} = \frac{\varepsilon_0 A}{d_0 \left(1 + \frac{d_1}{d_0} \sin \omega t\right)}, \quad (3.3)$$

where ε_0 is the permittivity of free space and ω is the angular frequency of the probe oscillation.

Substituting Equation (3.3) into Equation (3.2) gives an expression for the Kelvin probe current which is related to easily measured geometrical quantities, the probe oscillation frequency and the sample potential,

$$I(t) = -\frac{\omega \varepsilon_0 A V_s}{d_0} \frac{\frac{d_1}{d_0} \cos \omega t}{\left(1 + \frac{d_1}{d_0} \sin \omega t\right)^2}. \quad (3.4)$$

Since the model assumes a parallel plate configuration with plates of equal size, which is clearly not the case, the effective area of the plates in the model must be determined for this model to work well. To do this, the capacitance between a circular disc, with the same dimensions as the Kelvin probe tip, and a rectangular sheet of copper, with dimensions much larger than that of the Kelvin probe tip, was modeled in the finite element analysis program ANSYS. Using ANSYS it is possible to calculate the capacitance of the set up purely from the geometries of the two plates. This will also take into account fringing fields. If fringing fields were not taken into account the total capacitance could be significantly underestimated [111]. The values of the capacitances between the probe tip and the copper sheet, as well as the effective parallel plate radii,

at different probe-sample separation distances are given in Table 3.1.

Table 3.1: Values of the capacitances and effective radii of the Kelvin probe and copper sheet for different probe-sample separation distances

Probe-Sample Distance (mm)	Capacitance Calculated Using ANSYS (fF)	Effective Radius (mm)
1	145	2.28
5	93.5	4.10
10	80.5	5.38
15	68.5	6.08
20	63.7	6.77
25	56.1	7.10
30	50.9	7.42

A relation was found between the probe-sample separation distance and the effective radius of the probe by plotting the effective radius against the probe-sample distance on a log-log scale. The function for calculating the effective radius for any probe sample separation was found from the best fit straight line through the data. This plot is shown in *Figure 3.7*. It can be seen on *Figure 3.7* that the data start going non-linear at large probe-sample separations. This non-linearity is ignored because the linear fit should be adequate for the range of probe-sample separation distances considered here.

The final unknown quantity in Equation (3.4) is d_1 . In order to observe how the modulation distance of the Kelvin probe changed with oscillation frequency, the tip of the Kelvin probe was recorded using a high speed camera while the probe was running. A Casio EX-FH20 digital camera was used to record high speed videos of the Kelvin probe at 1000 frames per second. The videos were imported into Matlab movie player where each pixel could be more easily defined using the Pixel Region tool. A 1 mm slip gauge was imaged at the same time as the oscillating Kelvin probe so that once the Pixel Region tool had been used it was possible to assign a distance to each pixel. An image of a single movie frame of the probe tip and slip gauge is shown in *Figure 3.8*,




Figure 3.7: A plot of the effective radius versus probe-sample separation distance. The effective radii of the parallel plates can be calculated for any probe-sample separation from the best fit line through the data.

while Figure 3.9 shows the same frame but after each pixel has been defined with the Pixel Region tool. The slip gauge was 5 pixels across which means that each pixel was 0.2 mm square.




Figure 3.8: A screen shot of a high speed movie of the Kelvin probe oscillating. The probe tip can be seen on the left hand side of the image and the 1 mm slip gauge can be seen on the right hand side of the image




Figure 3.9: A screen shot of a high speed Kelvin probe movie that has each pixel more easily defined. This was done using the Pixel Region tool in Matlab.

Matlab movie player allows the user to skip through the movie frame by frame. This was done until the probe tip reached the maximum point of its oscillation. The pixel row which corresponded to the middle of the probe tip was noted and then the movie was forwarded until the probe tip reached the minimum of its oscillation. The pixels between the maximum and minimum were then counted to determine the oscillation amplitude of the probe tip at different oscillation frequencies. *Figures 3.10 and 3.11* show the measurement of the probe oscillation at a frequency of 183 Hz and a table of the results are shown in Table 3.2. The error on the oscillation amplitude is taken as \pm half a pixel or ± 0.1 mm.




Figure 3.10: A zoomed in image of the Kelvin probe tip at the maximum point of oscillation.




Figure 3.11: A zoomed in image of the Kelvin probe tip at the minimum point of oscillation. The distance between the maximum and minimum point of oscillation in this instance was found to be 0.8 mm.

Equation (3.4) gives the Kelvin probe current signal, however, it is voltage that is actually measured by the phase sensitive detector, therefore the gain of the current-to-voltage conversion that occurs in the Kelvin probe control unit must be calculated. For obvious reasons the company from which the Kelvin probe equipment was bought did not wish to divulge any detailed information about the Kelvin probe circuitry, so a measurement of the gain was made by applying a known AC current from the output of the lock-in amplifier to the input of the control unit. The AC output voltage from the control unit was

Table 3.2: Measured Kelvin probe oscillation amplitudes from the high frame rate movies

Kelvin Probe Oscillation Frequency (Hz)	Kelvin Probe Oscillation Amplitude (mm)
182	0.55
183	0.4
184	0.3
185	0.2
186	0.1
187	0.05

then measured with the lock-in amplifier so that the demodulated DC voltage could be ascertained. A plot of Kelvin probe output voltage versus Kelvin probe input current is shown below in *Figure 3.12*.




Figure 3.12: Plot of the Kelvin probe output voltage versus probe input current. The gradient of the best fit line through the data points gives the gain of the current-to-voltage converter of the Kelvin probe control unit circuitry.

The gain of the Kelvin probe control unit circuitry can be deduced from

the gradient of the best fit line through the data points. The total gain of the Kelvin probe circuitry, on the probe's lowest sensitivity setting, was found to be 2.1×10^8 V/A. The gain of each of the three probe sensitivity settings is available in the manual that came with the Kelvin probe, therefore, it would be easy to change the total gain of the model if another sensitivity setting was being used. This was not confirmed through measurement because the lowest sensitivity setting was the only setting used for the measurements presented in this work.

Using Equation (3.4) it was now possible to accurately predict the change in output probe voltage as the distance between the probe and the surface of the sample was changed. The change in output probe voltage as the oscillation frequency of the probe was altered was also modeled. These two modeled scenarios are shown in *Figure 3.13* and *Figure 3.14*.

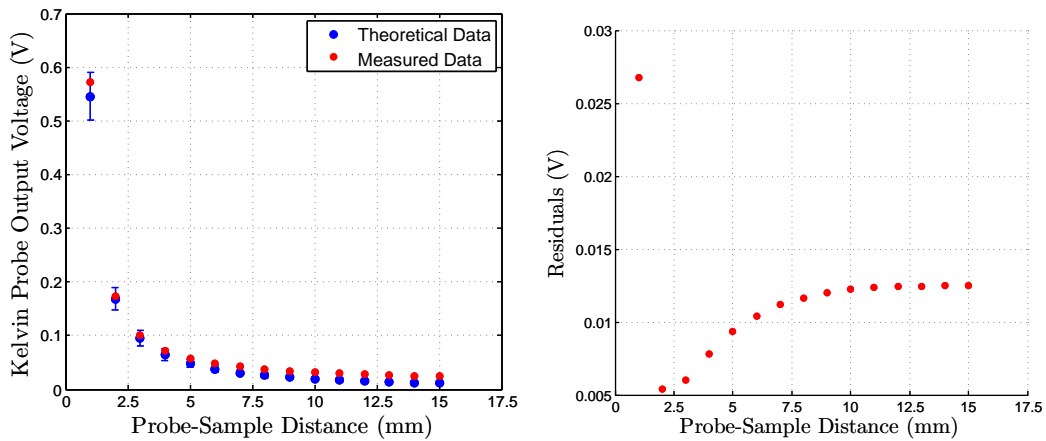


Figure 3.13: A plot of the probe output voltage against the probe-sample distance. The measurements made with the Kelvin probe are represented by the red markers and the predicted probe output voltage calculated from the model is represented by the blue markers.

The data shown in *Figures 3.13* and *3.14* was taken with the Kelvin probe measuring a large sheet of copper held at a potential of 30 V. The error on the modeled data results from uncertainty in the measurements of the probe-sample separation distance and the oscillation amplitude which was $\sim 10\%$. In




Figure 3.14: A plot of the probe output voltage against oscillation frequency. The measurements made with the Kelvin probe are represented by the red markers and the predicted probe output voltage calculated from the model is represented by the blue markers.

Figure 3.13 the measured data and the model start to deviate at large probe-sample separation distances. This is likely due to stray capacitance from other surfaces that may be closer to the probe tip than the copper sheet.

The residuals of both the modeled scenarios are shown below the data plots in Figures 3.13 and 3.14. The residual plots show that the probe signal is always greater than the model. This is likely because of stray capacitances being picked up by the probe as previously mentioned. The residuals in Figure 3.13 highlight the theory that stray capacitance could be a factor at large probe-to-sample distances since the residuals level off to a constant value.

3.4 Method of Calibration

The Kelvin probe has to be calibrated before meaningful surface charge measurements can be obtained. The method of calibration used was the same method that was first used by D. K. Davies in 1967 for electrometers [104]. Davies shows that if the Kelvin probe set up is simply represented as a circuit of capacitors, as shown in Figure 3.15, an equation can be derived that allows

the surface charge density on an insulating surface to be calculated from a given Kelvin probe signal.

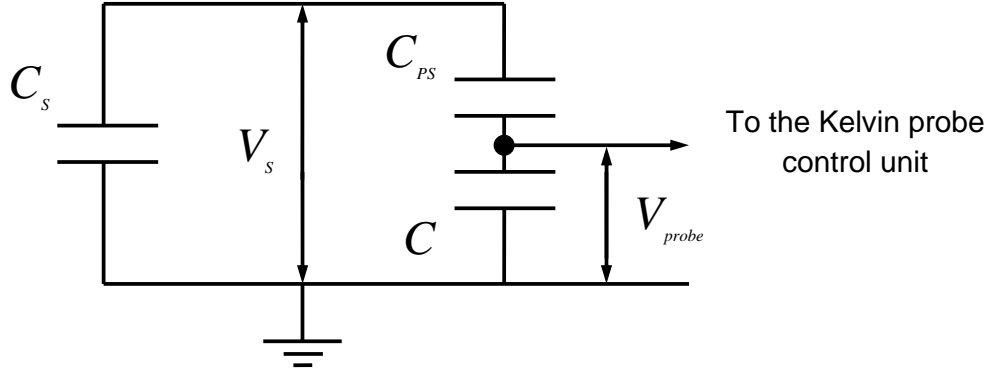


Figure 3.15: A schematic of the Kelvin probe set up represented in a circuit diagram. The probe, sample and ground surfaces can be thought of as the plates of capacitors where C_s is the sample-to-ground capacitance, C_{ps} is the probe-to-sample capacitance and C is the probe-to-ground capacitance.

Using Figure 3.15 it is possible to derive an expression for the measured probe voltage, V_{probe} , that is related to the potential on the surface of a sample, V_s :

$$V_s = I (Z_{ps} + Z), \quad (3.5)$$

$$V_{probe} = IZ, \quad (3.6)$$

$$V_{probe} = V_s \frac{Z}{Z_{ps} + Z}, \quad (3.7)$$

where I is the current in the circuit, Z is the probe-to-ground impedance and Z_{ps} is the probe-to-sample impedance.

Expressions for the impedances in terms of capacitance are then substituted into Equation (3.7). The gain of the Kelvin probe amplifier, g , is also included as a multiplicative factor to give,

$$V_{probe} = gV_s \left(\frac{1}{1 + \frac{C}{C_{ps}}} \right), \quad (3.8)$$

where C is the probe-to-ground capacitance and C_{ps} is the probe-to-sample

capacitance.

It is now possible to derive an equation that gives the calibration coefficient of the Kelvin probe. This is achieved by substituting in a new expression for C_{ps} that gives the capacitance in terms of the capacitor plate area and the distance between the plates, given in Equation (3.9), into Equation (3.8).

$$C_{ps} = \frac{\varepsilon_0 A'}{h} \quad (3.9)$$

$$\frac{C}{A'} = \frac{\varepsilon_0}{h} \left(g \frac{V_s}{V_{probe}} - 1 \right) \quad (3.10)$$

where ε_0 is the permittivity of free space, h is the distance between the probe and the surface of the sample, A' is the effective area of the probe and $\frac{C}{A'}$ is the calibration coefficient of the Kelvin probe.

The calibration coefficient was calculated by measuring a copper plate held at different known potentials, V_s , with the Kelvin probe. This allowed a straight line plot of V_{probe} against V_s to be created and the calibration coefficient was calculated from the gradient of the best fit line through the data. An aluminium plate, a stainless steel plate and a polished copper plate were also measured to ascertain whether there were any large discrepancies between the values of $\frac{C}{A'}$ for different metals. The results obtained from these measurements are presented in *Figure 3.16*. The standard deviation in the values of the gradients from each data set was less than 10 %, therefore, the error on the mean value of the gradient will be taken as negligible compared to other larger sources of error in this calibration. The value of $\frac{C}{A'}$ calculated from the mean gradient was $6.56 \times 10^{-7} \text{ F/m}^2$.

The standard definition of capacitance can be used to relate V_s to the surface charge density on the sample being measured, σ ,

$$V_s = \frac{\sigma A'}{C_T}, \quad (3.11)$$

where C_T is the combined parallel capacitance of C , C_{ps} and C_s . The probe-




Figure 3.16: Plots of V_{probe} against V_s for each of the four scenarios measured. Four values of $\frac{C}{A'}$ were calculated from the gradients of the straight lines and the average value of the four was used in the calibration.

to-ground capacitance, C , is assumed to be negligible [104], therefore, C_T can be expressed simply as,

$$C_T = C_{ps} + C_s, \quad (3.12)$$

which can be expanded to give,

$$C_T = \frac{\varepsilon_0 A'}{h} + \frac{\varepsilon_0 \varepsilon_r A'}{t} = \frac{\varepsilon_0 A'}{h} \left(1 + \frac{\varepsilon_r h}{t} \right), \quad (3.13)$$

where ε_r is the relative permittivity of the sample material and t is the thickness of the sample. Substituting for C_T in Equation (3.11) gives,

$$V_s = \frac{\sigma h}{\varepsilon_0} \left(1 + \frac{\varepsilon_r h}{t} \right). \quad (3.14)$$

Equation (3.14) is then substituted into Equation (3.10) and rearranged to give an expression for calculating the surface charge density on a sample from

the measured probe voltage,

$$\sigma = \left(\frac{C}{A'} \frac{h}{\varepsilon_0} + 1 \right) \left(1 + \varepsilon_r \frac{h}{t} \right) \frac{V_{probe} \varepsilon_0}{gh}. \quad (3.15)$$

Since the term $Ch/A'\varepsilon_0$ is much larger than 1, Equation (3.15) can be simplified further to give,

$$\sigma = \frac{C}{A'} \left(1 + \varepsilon_r \frac{h}{t} \right) \frac{V_{probe}}{g}. \quad (3.16)$$

A calibration plot of σ against V_{probe} was created in order to convert the Kelvin probe signal into a surface charge density. This was done by measuring a $10\text{mm} \times 10\text{mm}$ area on a fused silica sample, after charge had been deposited on the sample, and then converting the average Kelvin probe signal measured over the $10\text{mm} \times 10\text{mm}$ into a surface charge density using Equation (3.16). This measurement was repeated several times for different values of σ . Each time a set of measurements were taken a background measurement was taken with the sample removed from the vacuum chamber. The background measurement was subtracted from the sample measurements because the background signal level is likely due to stray capacitances, this will be discussed in Section 3.5.1. With the background level subtracted the best fit line on the surface charge density calibration plot passes through the origin of the graph. An example of a surface charge density calibration plot is shown in *Figure 3.17*.

It would appear that the calibration plot could be used to calculate the surface charge density for any material just by changing the appropriate values in Equation (3.16). To test this, a sapphire sample was measured in the same way as the fused silica sample so that a calibration plot for sapphire could be created. This is also shown on *Figure 3.17*. The plot of the sapphire data was then compared to a plot that was created by taking the Kelvin probe data of the silica sample and simply changing the values of ε_r and t from those used in the fused silica measurements to those that would be appropriate for the sapphire sample. The value of h was kept the same for both sets of measurements.




Figure 3.17: A Kelvin probe calibration curve that allows the surface charge density of the sample to be estimated from the Kelvin probe signal. A calibration was carried out for both sapphire and fused silica.

The value of ϵ_r was changed from 3.8, used for silica [112], to 9.3 for sapphire [113] and the sample thickness was changed from 5 mm to 8 mm. It was found that both the plots agreed exactly, confirming that our calibration plots can be altered for any material without having to take new measurements.

The biggest source of error in the calibration was due to the fact that the Kelvin probe oscillates and this puts a significant error on h of 50% since the probe tip was 1 mm from the surface of the sample and the oscillation amplitude was 0.5 mm.

3.5 Noise Sources

3.5.1 Stray Capacitance

One of the most dominant noise sources in Kelvin probe measurements is the effect of stray capacitance [114–116]. This is due to the fact that the Kelvin

probe is not an isolated system and can be coupled to grounded surfaces in the chamber or nearby wires. If the contribution that these stray surfaces make to the total charge, Q , detected by the probe is considered, it can be shown that the total charge on the probe is given by [114],

$$Q = C_{12}V_{12} + \sum_{j \geq 3} C_{1j}V_{1j}, \quad (3.17)$$

where C_{12} is the capacitance between the probe (surface 1) and the sample surface (surface 2), V_{12} is the potential difference between the probe and the sample surface and C_{1j} and V_{1j} are the capacitance and voltage between the probe and the j^{th} stray surface.

In order to reduce the effects of stray capacitance all wires and electrical connections are kept out of the way of the probe, all surrounding surfaces were grounded, and all electrical connections to the Kelvin probe were shielded. It is unlikely that the measurements presented in this thesis are affected by stray capacitance because the Kelvin probe output voltage can be modeled accurately without taking stray capacitance into account (see *Figures 3.13* and *3.14*). The only time stray capacitance might be a factor is at probe-to-sample separations greater than a few millimeters, however, the probe-to-sample distance was kept to 1 mm for the majority of the measurements presented here. The only measurements that used a probe-to-sample distance greater than 1 mm are the surface charge density measurements of First ContactTM with and without carbon nano-tubes presented in Chapter 5.

Even though it was thought that the effects of stray capacitance were unlikely to significantly affect the measurements presented in this theses, a quick test for stray capacitance effects was carried out. This test was similar to that of the measurements carried out by D'Arcy and Surplice [114]. A copper plate held at a potential of 10 V was moved around the inside of the vacuum chamber near the walls of the chamber. This was to simulate a high potential difference between the vacuum tank walls and the probe tip to see if there was

a detectable signal from the Kelvin probe. It was possible that the signal did increase by 0.1 mV, or $2.9 \times 10^{-12} \text{ C/m}^2$, however, at this level, it was difficult to determine whether this voltage change was due to the copper plate or just electronic noise. The only precaution that was taken in the experiments presented here was that a background measurement was taken regularly, without a sample in the vacuum chamber, and this value was subtracted from the measurements. This background signal was likely due to stray capacitance from other surfaces, such as wires and the motorised stages, in the vacuum chamber. The background level was typically less than 1 mV which corresponds to a surface charge density of $2.9 \times 10^{-11} \text{ C/m}^2$.

3.5.2 Leakage Current

It was noticed that while taking the surface charge measurements after the discharge procedure discussed in Chapter 4, the Kelvin probe signal had a tendency to drift quickly over short timescales, see *Figure 3.18*. These timescales were too short, typically hundreds of seconds, to be the correlation time of the charging noise and were thought to be from charge on the surface of the sample leaking to the grounded sample stage, perhaps through some monolayer of water on the sample. Due to the electric field which will exist between the surface of the sample and the grounded stage plate, it is possible that a small amount of charge could flow to ground, or be sourced up from ground. To rectify the problem another grounded plate was attached to the Kelvin probe support above the sample [117]. Adding an extra grounded plate nulls the electric field at the surface of the sample so that the surface charge stays on the sample. The copper plate can be seen in the picture of the Kelvin probe shown in *Figure 3.3*. It can be seen from *Figure 3.18* that the additional grounded plate above the probe was very effective at eliminating drift of the surface charge on short timescales.




Figure 3.18: A plot of the surface charge on a silica sample changing over timescales that are much quicker than what would be expected from the correlation time of the movement of the surface charge. When an additional grounded plate is positioned above the probe, at the correct height, the drift stops.

3.6 Conclusion

The Kelvin probe is an extremely useful and well established surface analysis instrument. In this chapter it was shown that the signal that is obtained using the Kelvin probe set up can be modeled accurately when compared to measurements. The model does not take into account external noise sources, such as stray capacitance, therefore, this shows that the measurements taken with this Kelvin probe set up are not greatly affected by spurious noise sources.

Using a calibration method devised by Davies [104] it was possible to convert a measured probe voltage into a surface charge density. This calibration was vital in producing a meaningful result from the Kelvin probe, since a surface charge density can be used to estimate the charging noise that would be expected from a given measurement.

Chapter 4

Mitigating Surface Charge Using an AC Glow Discharge

The potential for excess noise due to surface charge on gravitational wave detector optics has led to promising research into different in situ charge mitigation processes. A brief overview of these methods was given in Chapter 2, but in this chapter a method that utilises an AC glow discharge to reduce excess surface charge on bare fused silica and fused silica with a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer coating will be discussed.

Glow discharge processes have a wide variety of industrial applications such as fluorescent lighting [118], plasma reactors [118], deposition of thin films [119], and static elimination [120]. The advantage of using a glow discharge for static elimination is that both positive ions and free electrons are produced in the process, therefore, it would be possible to neutralise both polarities of surface charge using just one procedure instead of two, hence why it was chosen for this investigation. The flow of the charged particles is controlled using a Proportional-Integral-Derivative (PID) controller. This allows an almost neutral flow of particles to be maintained to ensure the charge on the surface of the sample under test is reduced to a minimum for both polarities of surface charge.

The experimental set up of the glow discharge apparatus will be discussed

in this chapter. This will include a description of how a glow discharge is generated and a description of how the PID controller functions. The results presented include surface charge density measurements of silica samples, taken with the Kelvin probe described in Chapter 3, before and after the glow discharge procedure. The charging noise that would be expected, based on the Kelvin probe measurements, will be modeled to show that the surface charge density on the silica samples is reduced to a level that will not limit a future gravitational wave detector like aLIGO. To show that this charge mitigation procedure does not cause a significant amount of damage to the TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer coating, an analysis on the transmittance of the coated samples after the charge mitigation process is also presented in this chapter.

4.1 Electric Glow Discharges

4.1.1 Introduction to Gas Discharges

Ionised gas, or plasma, can be created by an electrical gas discharge. A classical gas discharge can be created using an evacuated vacuum vessel containing two separated parallel plate electrodes. The vessel is generally filled with a noble gas, with argon being a common choice, at a pressure between 0.1 and 10 Torr. A diagram of the vacuum vessel described above is shown in *Figure 4.1*.

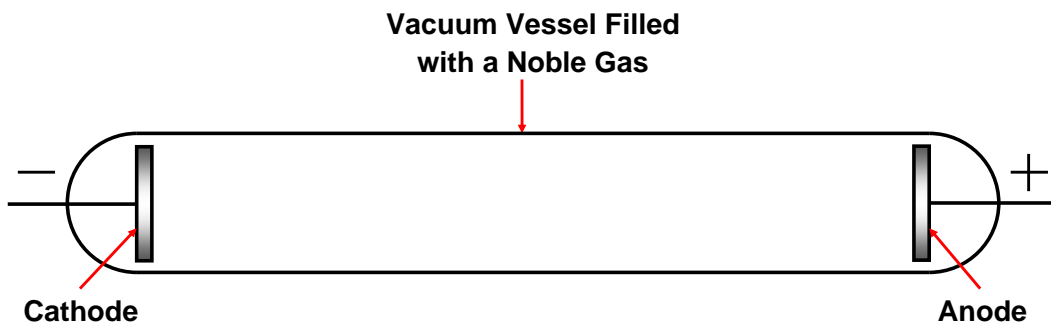


Figure 4.1: A diagram of a classical gas discharge tube. The tube is typically made of glass and usually contains a noble gas at a low pressure.

When a potential difference is applied between the two electrodes an electric

discharge can be created. The type of electrical discharge created depends on the voltage and discharge current characteristics of the set up. A graph of the different gas discharge regimes that can arise is shown in *Figure 4.2*.




Figure 4.2: Voltage and discharge current characteristics of electrical gas discharges. The axes on the plot show breakdown voltage on the y-axis and discharge current along the x-axis. The plot shows that electrical gas discharges go through many different forms as the voltage and current are changed. The letters A-J mark the boundaries of each of the regimes shown on the plot. This plot is reproduced from [118].

As can be seen from *Figure 4.2*, there are three broad regions of low pressure DC electric discharges; dark, glow and arc discharges. The types of discharge that are of interest to the work presented in this thesis are the glow discharge regime, which will be discussed further in this chapter, and the corona discharge regime which will be discussed in Chapter 5.

4.1.2 The Glow Discharge Region

A glow discharge can be created by applying a potential, typically at a few hundred volts, across a volume of gas. As electrons are accelerated from the

cathode they will collide with neutral gas atoms creating ions and free electrons. If the electron energy and electron number density are high enough visible light will be produced by excitation collisions, hence the plasma is called a glow discharge. The newly created positive ions will be driven towards the cathode and the newly created free electrons, and negative ions, will accelerate towards the anode. This new population of charged particles will collide with more neutral gas atoms creating a self-sustaining discharge as this process is continually repeated.

A glow discharge contained within a classic discharge tube, as described earlier, has a distinctive structure. The structure of a typical glow discharge is shown in *Figure 4.3* and the different regions of the glow discharge will be discussed briefly here.




Figure 4.3: A diagram showing the structure of a glow discharge. Glow discharges can have several distinctive dark and luminous regions that arise along the axial length of the discharge.

4.1.2.1 Aston Dark Space

The Aston dark space is a region of the glow discharge with a high electric field and a negative space charge¹ as it is mainly filled with electrons that are in the process of being driven away from the cathode. These electrons have not yet

¹Rather than treating each region of the glow discharge as being filled with distinct point-like charges it is more common to treat the region as a charge distributed over a volume of space called a space charge.

been accelerated to energies that are capable of creating excitation collisions with the neutral gas atoms in the discharge tube, therefore, this area remains dark. The Aston dark space is sometimes difficult to see in a glow discharge as it is often obscured by the cathode glow.

4.1.2.2 Cathode Glow

The cathode glow is a region of high ion density. This is due to the large number of positive ions moving towards the cathode. Some of these positive ions will be excited due to collisions with free electrons moving away from the cathode and will emit a photon which gives rise to the luminosity in this region. The length of the cathode glow changes depending on the type of discharge gas and the pressure of the gas.

4.1.2.3 Cathode Dark Space

The cathode dark space is a dark region that follows the cathode glow. No excitation collisions take place in this region because the electrons being driven from the cathode will have lost energy during the collisions that occurred in the cathode glow region. The moderate electric field in this region is able to accelerate these electrons to a point where excitation collisions can occur in the next section of the glow discharge, the negative glow region. This, again, is an area of high ion density.

4.1.2.4 Negative Glow

The negative glow region is where the most intense ionisation and excitation takes place. Consequently, this is where the luminosity of the glow discharge is most intense. The electron number density is high in this region as this is where most of the electron current is carried.

4.1.2.5 Faraday Dark Space

After the negative glow region the electrons are slowed down sufficiently enough to halt any ionisation by excitation. As a result, the region after the negative

glow region has very little light emission and is called the Faraday dark space. In this region the electron number density tends to decrease as the electrons recombine with ions in this region.

4.1.2.6 Positive Column

The region after the Faraday dark space is the positive column which is the region that most resembles a classical description of a plasma due to the fact that it is quasi-neutral. The electric field in the positive column is just large enough to maintain the level of ionisation necessary to continually generate this quasi-neutral state. Since the positive column is almost neutral it has a very uniform glow.

4.1.2.7 Anode Glow

At the end of the positive column there is a slightly brighter region called the anode glow. The anode glow is a boundary of positive space charge between the positive column and the anode dark space. This arises due to the fact that the region immediately in front of the anode, the anode dark space, has a negative space charge. The anode glow helps draw electron current from the positive column to the anode.

4.1.2.8 Anode Dark Space

The anode dark space is often referred to as the anode sheath. A plasma sheath arises when an excess of a particular polarity of charge on the surface of an electrode attracts particles of the opposite polarity from the plasma. Since the plasma is quasi-neutral this layer of charges attracted to the electrode acts as a barrier between the plasma and the electrode and is often referred to as a sheath. The anode attracts electrons from the positive column and, therefore, there is a negative space charge in this region creating a sheath around the anode.

4.1.3 Paschen's Law

In the late 1800s Friedrich Paschen discovered that the breakdown voltage of a gas, V_b , between two parallel plate electrodes is a function of the product of the gas pressure, p , and the electrode distance, d , [121]. Paschen's work also shows that the breakdown voltage decreases with the product dp up to a point where it reaches a minimum value called the Paschen minimum. If dp is decreased further, beyond the Paschen minimum, the breakdown voltage starts to rise again. The breakdown voltage can be calculated using Paschen's Law [121, 122],

$$V_b = \frac{Bpd}{\ln(Apd) - \ln\left(\ln\left(1 + \frac{1}{\gamma}\right)\right)}, \quad (4.1)$$

where A and B are the Paschen coefficients for the particular gas used and γ is the electron emission coefficient. The electron emission coefficient is simply the number of electrons emitted from the cathode per incident ion or photon. Paschen curves for air (nitrogen) and argon that were calculated using Paschen's law are shown in *Figure 4.4*.

In a glow discharge, the area that encompasses the region from the surface of the cathode electrode to the negative glow is often referred to as the cathode region. At low pressures, the cathode region will adjust its axial length, d_c , in order to establish a minimum value of the product d_cp where,

$$d_cp \approx (dp)_{min}. \quad (4.2)$$

The product $(dp)_{min}$ corresponds to the minimum breakdown voltage of the gas, or the Paschen minimum. At the Paschen minimum the glow discharge maintains itself under conditions of a minimum cathode fall voltage and power dissipation. The cathode fall voltage will be approximately equal to the minimum breakdown voltage of the gas under this condition.

For this investigation a noble gas should, ideally, be used as the discharge gas. This is because noble gases are unreactive and therefore shouldn't damage




Figure 4.4: *Paschen curves for air and argon gas. The Paschen coefficients and the electron emission coefficient required for these plots were taken from [123].*

the multi-layer reflective coating on the coated samples. Argon was chosen as the discharge gas because it is used in the deposition process of the reflective coating [124], therefore, it is very unlikely that ionised argon would damage the coating. Argon is also a sensible choice in terms of trying to minimise the energy of the electrons and ions created from the glow discharge as it has quite a low Paschen minimum compared to other noble gases that were considered, such as nitrogen. Since argon breaks down more easily, the driven electrode voltage can be kept as low as possible and, therefore, reduce the possibility of damaging the sensitive multi-layer coating through the emission of UV light.

4.1.4 AC Glow Discharges

In the set up that will be described in Section 4.2 an AC glow discharge was used. The reason this was used instead of a DC glow discharge was because it was found that when the glow discharge was operated with a DC power supply

there was always a residual charge on the silica samples after the discharge was turned off. This residual charge was of the same polarity as the voltage that was being applied to the driven electrode, with the second electrode remaining grounded. The explanation for this is likely that the plasma potential is slightly negative when a negative potential is applied to the driven electrode and slightly positive when the applied voltage is positive. *Figure 4.5* shows how the potential may vary between the two electrodes with a square wave signal of amplitude V_0 applied to the driven electrode.

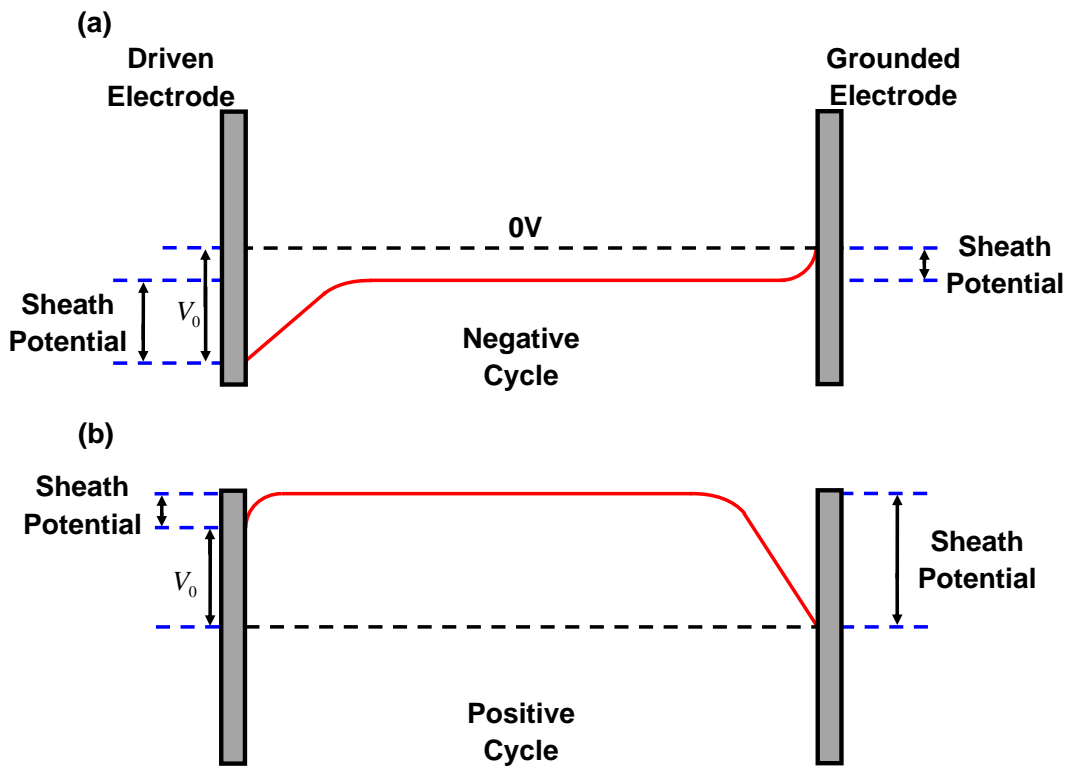


Figure 4.5: Diagram (a) shows how the potential across the glow discharge may vary during the negative cycle of the square wave signal and diagram (b) shows how the potential across the discharge may change during the positive cycle of the square wave signal. V_0 is the amplitude of the square wave signal applied to the driven electrode.

Using an AC square wave signal reduced the residual charge on the silica samples. However, it was extremely difficult to ensure that the plasma voltage was the same magnitude, but opposite in sign, for both positive and negative

cycles of the square wave signal. To get the glow discharge as close to neutral as possible a DC offset was applied to the square wave. Using the DC offset allowed the flow of charges arriving at the sample to be finely tuned so that the flow was almost neutral. Pictures of the glow discharge electrodes used for this investigation can be seen in *Figures 4.6* and *4.7*. *Figure 4.6* shows the glow discharge plates while the glow discharge is not in operation and *Figure 4.7* shows the glow discharge plates in use.




Figure 4.6: *Glow discharge electrodes positioned inside the vacuum tank.*




Figure 4.7: *Glow discharge electrodes during operation.*

4.2 Experimental Set Up

For the charge mitigation experiment discussed in this chapter, the vacuum tank that housed the Kelvin probe described in Chapter 3 was used. Two aluminium plates were bolted to the inner structure of the vacuum chamber to make up the parallel plate electrode arrangement that produced the glow discharge. This geometry was chosen for simplicity. Each plate was $50 \text{ mm} \times 50 \text{ mm} \times 5 \text{ mm}$ and they were spaced approximately 30 mm apart. One electrode was earthed while the other plate, the driven electrode, was connected to a Matsusada Precision Inc. AMT-1.5B40 high voltage amplifier. The high voltage amplifier had a Hewlett Packard 33120A waveform generator connected

to its input that supplied a 1 kHz square wave signal to the amplifier. This is a similar set up to that used by A. Ohsawa [120] for making a precisely balanced ioniser, however, a much simpler electrode geometry is used here and the pressure at which the discharge is generated is considerably lower in this investigation.

It was observed that the high voltage amplifier had a noticeable decay time when it was turned off, probably caused by discharging capacitors in the amplifier, that would cause an undesired residual charge to be deposited on the surface of the sample. This problem was solved by using a Panasonic AQV254 solid state switch to disconnect the voltage to the driven electrode immediately when required. The solid state switch uses back-to-back MOSFETs which allow the AC current to flow through the device when activated. Unlike mechanical switches, this solid state switch uses optical coupling to turn the switch on/off. An LED is activated by an externally applied control voltage, of 5 V, which illuminates a photodiode. The current produced by the photodiode opens the gates on the MOSFETs so that current can flow through the device. When the switch required to be closed the supply current to the LED was grounded, thus closing the gates of the two MOSFETs.

In order to maintain an almost neutral flow of both positive and negative charge carriers, a Proportional-Integral-Derivative (PID) servo controller was implemented using LabVIEW. A Faraday cup [125], whose output was read by a Keithley 2000 multimeter and connected to the PID controller, was used to monitor the current of the charged particles flowing from the glow discharge. The Faraday cup was placed at a similar distance from the glow discharge as the silica sample in order to measure an identical current from charged particles arriving at the silica sample from the discharge. The current arriving at the Faraday cup was too small to measure with appreciable accuracy, therefore, a current-to-voltage converter was used to change the current into a voltage that could be measured more easily. A circuit diagram of the current-to-voltage

converter used in this experiment is shown in *Figure 4.8*.




Figure 4.8: A circuit diagram of the current-to-voltage converter used to convert the current measured by the Faraday cup into a voltage that could be measured with the Keithley multimeter.

If the voltage measured by the Keithley multimeter deviated from zero the PID controller changed the DC offset of the signal produced by the waveform generator to keep the flow of charged particles almost neutral. The algorithm used by the PID controller to calculate the change in the DC offset required will be discussed in Section 4.2.1.

It was noticed that the voltage signal from the unshielded Kelvin probe fluctuated wildly while the glow discharge was in operation. After the glow discharge was turned off the probe signal took a few hours to settle back to a steady value, so a box was constructed from a thin sheet of stainless steel to shield the probe. The box was cut into two sections, a small section that was bolted to an aluminium plate above the Kelvin probe and a larger section that was bolted next to the sample on the sample stage. To shield the probe, the stage was moved so that the large section was positioned underneath the probe and then the stage was moved up, in the z direction, until the small section slotted inside the larger section. This completely shielded the Kelvin probe from the glow discharge. *Figure 4.9* shows the Kelvin probe shielded in the steel box. The Faraday cup and a silica disc sample can also be seen in *Figure 4.9*.

Figure 4.10 shows the unshielded Kelvin probe taking a measurement of the silica disc. A schematic diagram of the entire glow discharge set up is shown in Figure 4.11.




Figure 4.9: A photograph of the sample stage set up for the glow discharge procedure. The Kelvin probe is shielded inside a stainless steel box and the Faraday cup monitors the current produced by the glow discharge near the sample.




Figure 4.10: A photograph of the Kelvin probe unshielded and in position to take a measurement of the surface charge density on the fused silica disc.

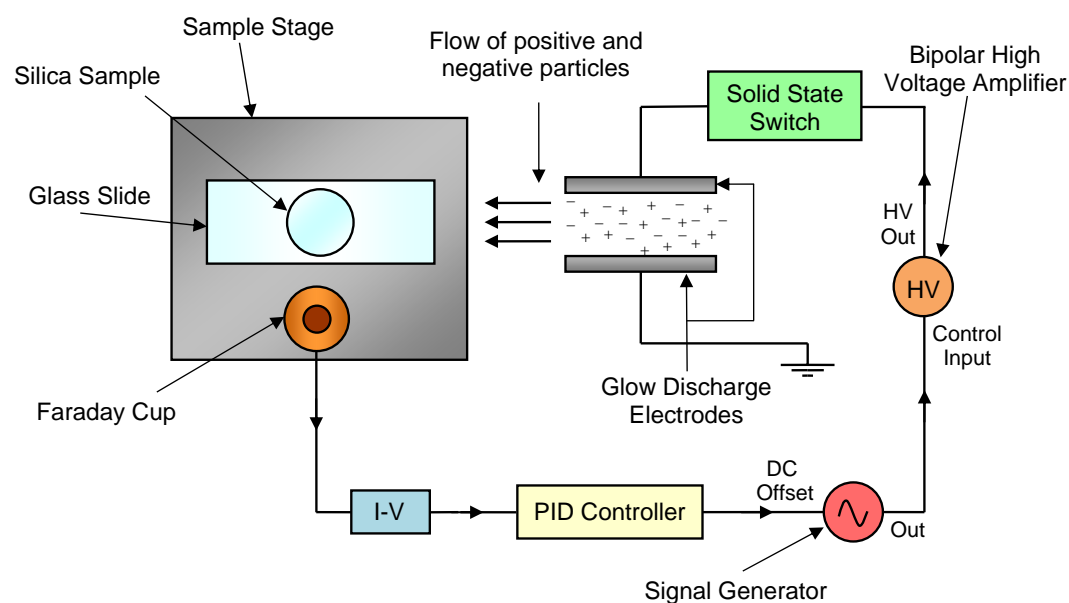


Figure 4.11: A schematic, aerial view, diagram showing the glow discharge experimental set up.

4.2.1 PID Servo Control System

A PID control system is possibly the most common type of feedback control. It is a common choice of controller due to its simplicity, intuitiveness and wide application range [126]. A PID controller changes one of the process variables of a system in order to obtain the desired response from the system. For this experiment a closed-loop PID controller was used to maintain a neutral flow of charged particles from the glow discharge. A PID controller was also used for the torsion balance experiment discussed in chapters 6 and 7, however, a full explanation of the general PID control algorithm used for both experiments will be given here.

Initially, the controller calculates the difference between the measured output of the system, M , and the desired output of the system, N , in order to evaluate the deviation from the desired system output. The desired system output is set by the user and is hereafter referred to as the set point. The deviation from the set point is called the error, e , where,

$$e = M - N. \quad (4.3)$$

Once the error is calculated, the controller can then calculate the appropriate P, I and D action to apply to the process of the system by using a basic PID algorithm.

4.2.1.1 The P, I and D Action

The proportional term affects the instantaneous error of the system and is used to shorten the system response time. The proportional term applies control action, u_p , that is proportional to the error of the system at a given instant. Thus, the proportional term is just a constant, called the proportional gain, K_p , multiplied by the error at time, t , where,

$$u_p(t) = K_p e(t). \quad (4.4)$$

A controller that only uses a proportional control term is adequate for systems that only encounter small deviations from the desired set point [126], however, this type of controller will usually generate a steady offset. In order to remove this offset an integral term can be added to the controller to create a PI controller. The integral term corrects for the summation of the instantaneous error over a time, τ , by applying control action, u_i ,

$$u_i(t) = K_i \int_0^t e(\tau) d\tau, \quad (4.5)$$

where K_i is the integral gain.

A PI controller is also adequate for some situations, however, it is often useful to include a derivative action for fuller control of the system. The derivative action of a PID controller helps account for predicted future changes of a system by determining the rate of change of the error. This helps to reduce overshoot caused by the integral term. The derivative action, u_d , can be determined from multiplying the rate of change of the error by the derivative gain, K_d ,

$$u_d(t) = K_d \frac{de(t)}{dt}. \quad (4.6)$$

When the different types of action have been evaluated they are simply summed to give the complete control action, u , where,

$$u(t) = K_p e(t) + K_i \int_0^t e(\tau) d\tau + K_d \frac{de(t)}{dt}, \quad (4.7)$$

in the time domain or,

$$u(s) = K_p + \frac{K_i}{s} + K_d s, \quad (4.8)$$

in the Laplace domain. The complete PID process is shown in the form of a schematic in *Figure 4.12*.

It was found that for the charge mitigation work the derivative action could be set to zero because the controller worked well enough without it. A plot of




Figure 4.12: A schematic diagram showing the PID feedback controller process.

the voltage read by the Keithley multimeter during the glow discharge process is shown in *Figure 4.13*. During this time the PID controller changes the DC offset of the waveform generator in order to achieve the set point desired.

It can be seen that in *Figure 4.13* there are some spikes that arise in the Faraday cup data. These spikes are mainly due to the waveform generator automatically changing its current limiter settings during the glow discharge process. As the current limiter changes, there is an audible click from the waveform generator and the driven electrode receives a small current surge. This caused a burst of extra charged particles to be created in the glow discharge which resulted in a spike in the current measured with the Faraday cup. Changes in gas pressure also led to occasional discharges between the driven electrode and other nearby grounded surfaces which also caused spikes in the Faraday cup measurements.

4.3 Results

For this investigation a one inch fused silica disc and a silica sample with a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer reflective coating were used in the glow discharge procedure. The silica disc was used in order to show that the glow




Figure 4.13: An example plot of the voltage read by the Keithley multimeter from the current-to-voltage converter during the glow discharge process. The PID controller changes the offset of the square wave voltage signal being applied to the driven electrode. This changes the current that flows from the glow discharge. The set point was -1.32 mV because this was the offset of the current-to-voltage converter when the glow discharge was turned off.

discharge could reliably reduce the surface charge density on silica and the coated sample was used to show that the glow discharge would not damage the multi-layer reflective coating. These samples are shown in [Figure 4.14](#).

It is also possible that the rear face of the test mass could become charged by touching the ESD on the reaction mass. A reliable charge mitigation procedure would have to work in the small gap between the test mass and the reaction mass to reduce the surface charge density on the optic. The aLIGO test masses are wedged, and so the distance between them and the reaction masses varies slightly. The smallest distance between the reaction masses and the test masses is 5 mm, therefore, the glow discharge process was tested for this separation distance. To simulate this scenario, a scale model of this set up was created using two three inch fused silica discs spaced approximately 1.1 mm apart.




Figure 4.14: A photograph of a silica disc (left) and a silica sample with a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer coating (right).

This set up is shown in Figure 4.15.




Figure 4.15: A photograph of the set up required to show that the glow discharge procedure was capable of mitigating charge between two closely spaced optics.

4.3.1 Surface Charge Measurements

In this section the surface charge density measurements of the one inch silica disc sample and the coated sample will be considered first. The samples were charged negatively by rubbing them with a section from a rubber o-ring or

charged positively by rubbing them with a small, soft bristled, paint brush and then placed under vacuum ($\sim 10^{-4}$ Torr). A $10\text{ mm} \times 10\text{ mm}$ scan of the disc sample was then taken with the Kelvin probe, while a $6\text{ mm} \times 6\text{ mm}$ scan was taken for the coated sample due to smaller sample size, so that the charge on the surface of the sample before the discharge process could be estimated. Once this had been completed the Kelvin probe was shielded with the grounded steel box to prevent it being affected by the glow discharge. The turbo pump was then turned off and a small flow of argon gas was released into the chamber using a needle valve. Once the chamber had reached a steady pressure of approximately 10 mTorr the glow discharge was turned on. Whilst the glow discharge was on, the graphical output of the PID controller was monitored to identify when the servo effort had reached a constant value. The set point in this case was the voltage measured by the Keithley multimeter when there was no glow discharge active in the chamber. When a stable set point had been achieved the solid state switch was deactivated to cut the power supply to the driven electrode, the gas flow was turned off and the turbo pump was re-started. Once the turbo pump was at full pumping speed the Kelvin probe was unshielded and another scan of the sample was taken. Examples of Kelvin probe scans of the coated silica sample before and after the charge mitigation process are shown in *Figures 4.16 and 4.17*.

The total time for setting up the process was 20-30 mins, however, the actual discharging part of the process only took around 2 mins on average. This procedure was repeated ten times for the silica sample and twenty times for the coated sample. The magnitude of the average surface charge density, $|\bar{\sigma}|$, before and after the charge mitigation process, the standard deviation between each Kelvin probe scan and the systematic error from the Kelvin probe calibration are presented in *Table 4.1*.

To put these results into context an estimate of the resultant charging noise in aLIGO was predicted using the method described in *Chapter 2*. A plot of




Figure 4.16: *Kelvin probe scan of the surface charge density across the coated silica sample before the glow discharge procedure. The average surface charge density across the scanned area is $2.42 \times 10^{-8} \text{ C/m}^2$.*




Figure 4.17: *Kelvin probe scan of the surface charge density across the coated silica sample after the glow discharge procedure. The average surface charge density across the scanned surface is $-9.55 \times 10^{-10} \text{ C/m}^2$.*

Table 4.1: Results of the Kelvin probe measurements of the silica samples before and after the glow discharge procedure.

Sample	$ \bar{\sigma} _{before}^a$ (C/m^2)	$ \bar{\sigma} _{after}$ (C/m^2)	Stand. Dev. (C/m^2)	Systematic Error (C/m^2)
Silica	$> 5.8 \times 10^{-8}$	9.45×10^{-10}	3.03×10^{-10}	8.22×10^{-10}
Coated Silica	$> 5.8 \times 10^{-8}$	1.68×10^{-9}	2.72×10^{-10}	1.46×10^{-9}

^a Only a lower limit is given for $|\bar{\sigma}|_{before}$ because the Kelvin probe signal was saturated on most of the measurements

the total aLIGO noise budget, the charging noise after the glow discharge procedure and the charging noise before the glow discharge procedure are shown in [Figure 4.18](#). Only a lower limit estimate of the charging noise predicted before the charge mitigation procedure could be given since the Kelvin probe signal had saturated during these measurements. As can be seen in [Figure 4.18](#), the resulting charging noise from the glow discharge procedure investigated here would not limit the low frequency performance of aLIGO.

The results of mitigating excess surface charge between the two three inch




Figure 4.18: Resulting charging noise after the glow discharge process, red line, compared to the aLIGO noise budget, black line. A lower limit estimate of the charging noise before the charge mitigation procedure is also shown, green dashed line.

silica discs will now be discussed. The smallest spacing between the aLIGO test masses and the reaction masses is just 5 mm, therefore, an in situ charge mitigation procedure for advanced detectors would have to be able to reduce the charge in this gap. In order to show that the glow discharge procedure investigated here was capable of this, a scale model of this set up was created using two three inch silica discs. The test masses and reaction masses of aLIGO are 340 mm in diameter, therefore, the set up used in this investigation was approximately 4.5 times smaller than the real aLIGO set up. This meant that the gap between the two three inch discs was set to 1.1 mm.

The silica disc attached to the sample stage was charged by rubbing it with a section of a rubber o-ring while no excess charge was deposited on the second disc. A 28 mm \times 28 mm scan of the sample was taken before and after the glow

discharge process with the Kelvin probe. Unfortunately, due to obstructions in the vacuum tank, this was the largest scan size that was possible. To get as large a scan of the surface as possible, the scan was also offset from the centre of the disc. Kelvin probe scans of the charged silica disc before and after the discharge process are shown in *Figures 4.19 and 4.20*.




Figure 4.19: *Kelvin probe scan of the surface charge density across the three inch silica sample before the glow discharge procedure. The average surface charge density across the scanned area is $7.25 \times 10^{-8} \text{ C/m}^2$.*




Figure 4.20: *Kelvin probe scan of the surface charge density across the three inch silica sample after the glow discharge procedure. The average surface charge density across the scanned surface is $-3.59 \times 10^{-9} \text{ C/m}^2$.*

The same procedure, described earlier in this section, was followed for the two three inch silica disc samples except the glow discharge was left on for 30 mins. The reason for this was because it was found that the surface charge density on the charged silica disc did not diminish as quickly using discharge times of a few minutes.

By comparing *Figures 4.19 and 4.20* it can be seen that this charge mitigation procedure does reduce charge between the two silica discs. It should also be noted that there are high patches of localised positive and negative charge before the charge mitigation procedure, but after, the charge distribution across the scanned area is a lot more uniform. This is highlighted by looking at the highest peak and lowest trough values in *Figures 4.19 and 4.20*.

The highest peak in *Figure 4.19* deviates from the mean by $3.95 \times 10^{-8} \text{ C/m}^2$ and the lowest trough deviates from the mean by $9.10 \times 10^{-8} \text{ C/m}^2$. However, in *Figure 4.20* the highest peak only deviates from the mean by $1.03 \times 10^{-8} \text{ C/m}^2$ and the lowest trough only deviates from the mean by $1.32 \times 10^{-8} \text{ C/m}^2$. The fact that the maximum deviations from the mean have decreased after the glow discharge procedure, and that the maximum deviations are closer to being the same value, shows that the surface of the silica sample is more uniform than it was before the discharge procedure. *Figure 4.21* shows a rescaled Kelvin probe scan of the silica disc after the glow discharge procedure. This shows that even on the higher resolution scale the surface charge density is a lot more uniform across the surface of the optic. This part of the investigation was not as extensive as the previous part because this was purely to show that it was possible to discharge in a gap between two silica masses.




Figure 4.21: A rescaled plot of the Kelvin probe scan of the surface charge density across the three inch silica sample after the glow discharge procedure. On a higher resolution it can be seen that the surface charge across the surface of the disc is still reasonably uniform.

These results show that the glow discharge charge mitigation process pre-

sented in this chapter is a consistent and reproducible way of reducing the surface charge density on fused silica and fused silica with a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer coating. The method is also robust since the parts that actually create the discharge only consist of two aluminium plates, so it is highly unlikely that the apparatus should fail. This is a distinct advantage over other methods of charge mitigation that use electron and ion guns as an ioniser since they contain filaments that could break under vacuum and would have to be replaced.

4.3.2 Transmission Measurements

When creating a process to discharge the detector optics it was of critical importance that the process did not damage the reflective coating in any way. To ascertain whether the glow discharge did affect the transmission of the optical coating the transmission of each sample was measured to see if any significant change could be observed. The transmission of a coated sample that had undergone the glow discharge process 20 times was compared to the transmission of a coated sample that had not been exposed to the glow discharge.

The transmission measurements were taken using a Sentech SE800-PV spectroscopic ellipsometer, shown in *Figure 4.22*. The ellipsometer shines light of varying wavelengths between a light source and a detector in order to measure the intensity, and polarisation, of the light over a range of wavelengths. The transmission of a sample is then measured from the ratio of the intensities of the light detected when a sample is not between the light source and the detector and when a sample is placed between the light source and the detector.

The wavelength range scanned was 850 – 1200 nm with 1064 nm being the wavelength of interest since this is the wavelength of the laser that will be used in aLIGO. For a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ coating the reflectivity should




Figure 4.22: A photograph of the ellipsometer used to measure the transmission of the coated silica samples used in this investigation. The arm which contains the light source and the arm which contains the detector are highlighted in the picture.

be approximately 99.9 % at 1064 nm, therefore, if the sample is undamaged by the discharge process a transmission of about 0.1 % should be measured. The transmittance of the two samples were measured 40 times each in order to obtain two adequately sized data sets to analyse. A transmission measurement of a coated sample is shown in *Figure 4.23* and the full results of the transmission measurements are shown in *Table 4.2*. These results include: the average transmission of the samples at 1064 nm, the standard deviation and the error on the mean transmission values. All the data taken for these measurements are represented in the histograms shown in *Figure 4.24*.




Figure 4.23: An example of the transmission data taken using the ellipsometer. The plot on the left shows the transmission data for the full wavelength range measured. The plot on the right shows the data for a smaller wavelength range so that the transmission at 1064 nm can be more easily seen. The transmittance of a coated silica sample at 1064 nm is almost zero, as expected.

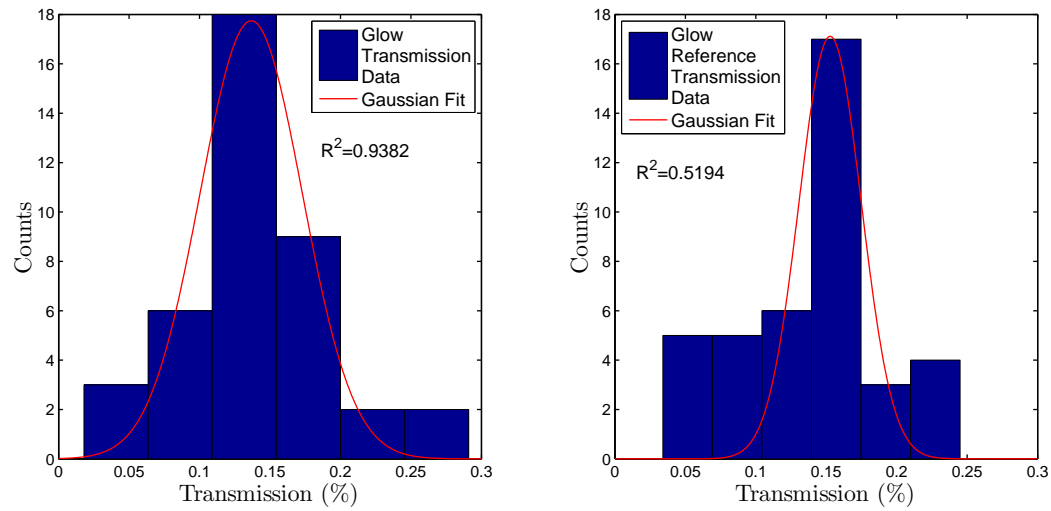


Figure 4.24: A histogram of all the transmission measurements of the glow discharge sample and the reference sample. The histogram on the left fits well to a Gaussian distribution, however, the histogram on the right does not fit as well to a Gaussian distribution. It is likely that some more data points would help increase the R^2 of the fit as the transmission data of the reference sample taken for the corona measurements (see Chapter 5) gives a good fit to a Gaussian distribution.

Table 4.2: Results of the transmission measurements of the reference sample and the glow discharge sample.

Sample	Average Transmission (%)	Standard Deviation	Error on Mean
Reference	1.23×10^{-1}	5.27×10^{-2}	8.34×10^{-3}
Glow	1.36×10^{-1}	5.46×10^{-2}	8.63×10^{-3}

4.3.2.1 Student's t-test

The transmission data in Table 4.2 was analysed using Student's t-test [127]. Student's t-test is a standard statistical test for comparing the means of two independent samples, in order to test the null hypothesis that there is no significant difference between the means of the populations from which the two samples were drawn.

Under the null hypothesis, and the assumption that the two populations are each well characterised by a normal distribution, the t-test statistic allows the computation of a confidence interval and a p-value for the difference of the population means. The confidence interval indicates the interval in which the difference of the population means should lie. Ideally the confidence interval should contain zero for the measurements presented here as this would indicate that the means of the two data sets are the same. A p-value which measures the probability, given that the null hypothesis were true, that we would obtain a difference in the sample means as large or larger than that obtained with the actual data is also acquired from Student's t-test. The test is carried out at a confidence level of 95%, thus, a p-value less than 0.05 would indicate poor support for the null hypothesis while a p-value greater than 0.05 would indicate that the data were consistent with being drawn from populations with the same mean value.

For this analysis the Student's t-test function in Matlab was used to compare the mean transmissions of the coated sample which had undergone the charge mitigation procedure, the Glow data set, and the coated sample which had not been exposed to the glow discharge, the Reference data set. The results of this analysis are presented in Table 4.3.

Table 4.3: Results from the Student's t-test analysis carried out on the transmission data of the reference sample and the glow discharge sample.

Samples	p-value	Confidence Intervals
Reference and Glow	0.39	$-3.40 \times 10^{-2}, 1.33 \times 10^{-2}$

From the results presented in Table 4.3 it can be seen that a fairly high p-value was obtained when the means of the Glow and Reference data sets were compared, thus indicating that the samples are consistent with being drawn from populations with the same mean value. Consequently, we also see that the confidence interval for the difference of means includes the value zero. This also suggests that the samples are consistent with being drawn from populations with the same mean value. This analysis strongly suggests that the glow discharge process did not cause an observable change in the optical properties of the multi-layer coating.

4.3.2.2 Kolmogorov-Smirnov test

To serve as a check to the Student's t-test analysis a Kolmogorov-Smirnov test (K-S test) [128] was also carried out on the data. The reason for this is that the K-S test does not make any assumptions about the distribution of the data sets whereas Student's t-test assumes that the data sets are normally distributed.

The K-S test is a statistical test for comparing the empirical cumulative distribution functions of two samples in order to test the null hypothesis that there is no significant difference between the populations from which the two data sets were drawn. The cumulative distribution function of a data set describes the probability that we would obtain a measurement less than or equal to that already obtained with the actual data.

As well as calculating a p-value, the K-S test calculates a test statistic, the K-S test statistic, which quantifies a distance between the empirical cumulative distribution function of the coated sample used in the glow discharge

experiment and the empirical cumulative distribution function of the reference sample. If a large value for the K-S test statistic is obtained then it is unlikely that the two data sets are drawn from the same data population and, therefore, a low p-value will be calculated. As with Student's t-test, the K-S test was carried out at a confidence level of 95%, hence p-values greater than 0.05 would indicate that the test has failed to reject the null hypothesis that the reference sample data and glow discharge sample data are drawn from data populations with the same mean value.

Table 4.4: Results from the Kolmogorov-Smirnov test analysis carried out on the transmission data of the reference sample and the glow discharge sample.

Samples	p-value	K-S Statistic
Reference and Glow	0.724	0.15

From the results presented in Table 4.4 it can be seen that, again, a high p-value was obtained when the Glow and Reference data sets were compared, thus reinforcing the hypothesis that the samples are consistent with being drawn from the same population. This is a strong indication that the glow discharge procedure did not cause substantial observable damage to the TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer reflective coating. To give a clearer indication of how similar the two data sets are, a plot of the empirical cumulative distribution functions of the Reference and Glow data is shown in *Figure 4.25*.

It should be noted that these measurements were taken only to check for large changes in the transmission of the coated sample. With these measurements the analyses could only detect changes in transmission of approximately 200 ppm. The sensitivity of approximately 200 ppm was calculated by creating mock data sets that had a similar standard deviation to the real data but a different mean value. The analyses were performed with the mock data and the reference sample data and a detectable difference was registered by the rejection of the null hypothesis. The empirical cumulative distribution func-




Figure 4.25: A plot of the empirical cumulative distribution functions of the transmission data obtained from transmission measurements of the reference sample, blue, and the glow discharge sample, red. The mock data, black, shows a cumulative distribution function that would cause the K - S test to reject the null hypothesis. The plot shows the probability, $F(x)$, of obtaining a transmission less than or equal to the data already obtained, x .

tion of the mock data is shown in Figure 4.25. The aLIGO mirrors require an absorption level of less than 1 ppm and with these measurements it was not possible to achieve this level of sensitivity. Therefore, the effect of this charge mitigation procedure will have to be investigated further in order to confirm whether it would be appropriate for use in future detectors.

4.4 Discussion

The results presented in this chapter show that, in principle, a glow discharge using argon gas could be used to effectively reduce electrostatic charge on the test masses of a future gravitational wave detector like aLIGO. However, there are some concerns that must be addressed before this procedure can be deemed

appropriate for use in aLIGO.

The largest source of error in this procedure is getting the discharge gas to a steady pressure in the vacuum chamber. If the pressure fluctuates then it is much harder to get a steady glow discharge. This problem could possibly be solved by having the discharge gas flow into the vacuum chamber through a flow regulator. In this investigation the gas was released into the vacuum chamber through a needle valve which made it difficult to maintain a constant pressure in the vacuum tank.

With a glow discharge there is the possibility that UV could be produced which, as mentioned earlier, can cause damage to the multi-layer coating. However, the voltage on the driven electrode is kept as low as possible in order to reduce this possibility. Improvements to this procedure were also considered which would eliminate the possibility of UV emission. If the glow discharge could be contained in some way that prevented any UV light illuminating the test mass, and yet still allow the ionised gas to flow to the surface of the test mass, the threat of UV would be eliminated. To give an idea of what such a device may look like, a photograph of an ioniser designed by Dr. L. Cunningham is shown in *Figure 4.26*. This ioniser was initially designed to create ionised gas using a corona discharge between two closely spaced plate electrodes contained within a white teflon cup. The electrodes had small holes drilled in them so that neutral gas could flow into the cup at one end and ionised gas could flow out of the cup at the other end. A similar device could be made where a glow discharge ionises the discharge gas inside the cup and the resulting ionised gas flows out of the device to mitigate charge on the detector mirrors. This would also give some directional control to where the discharge gas should flow rather than filling the whole chamber with ionised gas. This idea was not pursued due to time constraints.

Another problem that could be encountered in a gravitational wave detector is that the driven electrode can also begin to create a discharge between other




Figure 4.26: *A photograph of a corona discharge gas ioniser designed by Dr. L. Cunningham.*

grounded surfaces inside the vacuum chamber. This, again, makes it more difficult to obtain a stable glow discharge. The contained glow discharge idea could solve this problem or at least help reduce it.

One concern that had been raised by a few members of the LIGO scientific collaboration is the use of argon as the discharge gas. Argon can be harder to pump out of a vacuum system compared to some other gases [129] and could increase the down time of a detector. For example, the pumping speed for a Varian Turbo V3K-G vacuum pump is 2050 litres/second for argon and 2200 litres/second for nitrogen below 1×10^{-6} mbar. This is why other charge mitigation procedures being investigated elsewhere, that use ionised gas, use nitrogen as it is easier to pump out of a vacuum system. Nitrogen was not tried in this investigation because it has a higher Paschen minimum than argon, see *Figure 4.4*, therefore, using nitrogen would mean using larger voltages on the driven electrode and increasing the likelihood of producing UV. Again, if the contained glow discharge method described previously was used, nitrogen gas could be used as the discharge gas since the UV would not escape the ioniser device.

4.5 Conclusion

The charge mitigation processes investigated in this chapter was shown to reduce the surface charge density on a bare fused silica disc and a fused silica sample with a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer reflective coating to a level that would not limit a future gravitational wave detector such as aLIGO. It was also shown that it would be possible to mitigate charge between the test mass and reaction mass of a gravitational wave detector in the event that the test mass had touched the ESD on the reaction mass.

Using a glow discharge has the distinct advantage that it eliminates the use of filaments which could break in the vacuum chamber and would have to be replaced. If anything goes wrong with the glow discharge process it is most likely that a problem has occurred with the equipment outside the tank which means the vacuum of the detector wouldn't need to be broken. The parts and instruments needed for this procedure are also relatively easy and inexpensive to set up.

Even though the method with which the transmission of the coated samples was measured was only able to detect differences in transmittance at a level of 200 ppm, the results show that the glow discharge procedure did not cause damage to the TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ coating. One possible method of measuring differences in transmission at much higher levels of sensitivity is photothermal common-path interferometry (PCI) [130]. This method has been used to determine the increase in absorption caused by exposure of coated optics to UV radiation [131] and could be used to determine whether the charge mitigation method presented here causes an increase in absorption in TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ coatings at levels below 200 ppm.

The work in this chapter could be furthered by developing a device that constrains the glow discharge to eliminate the possible danger of illuminating the test masses with UV. This would also allow nitrogen gas to be used which is easier to pump out of a vacuum system. A great test of this discharging process

would be to try it at LASTI to see how it performs with optics, and a complete set up, similar to that of aLIGO. This would really show how appropriate the method is for an advanced gravitational wave detector. This research could also be continued by investigating charge mitigation on materials that will be used in third generation detectors such as silicon and sapphire. Some future detectors plan to cool the optics of the detector to cryogenic temperatures, so there is the need to investigate charging and charge mitigation at cryogenic temperatures as well.

Chapter 5

Reducing Surface Charge During the Cleaning of Future Gravitational Wave Detector Test Masses

5.1 Introduction

As mentioned earlier in Chapter 2, First ContactTM (FC) may be used to clean and protect the aLIGO test masses. If FC is going to be used, a procedure must be investigated that removes the excess charge that it deposits on the detector optics and can be easily incorporated into the cleaning process. A group of researchers at Caltech have had success reducing excess surface charge using a N₂ ioniser gun [132], however, there has been no formal publication of this work. The ioniser gun is used to blow nitrogen ions, created using a corona discharge, over the mirror's surface as FC is removed. Another possibility that has been considered is to make FC more conductive by mixing carbon nanotubes (CNT) into the FC solution. The reasoning behind this idea is that the CNT may increase the conductivity of the FC and therefore some of the deposited charge may be carried away in the FC film as it is removed.

In this chapter it will be determined whether FC mixed with CNT would be

appropriate to use on future gravitational wave detector optics. It is possible that the FC mixed with CNT will deposit carbon on the surface the optics which would increase absorption. Scanning electron microscope (SEM) images of a silica sample before and after FC with CNT is applied are taken to deduce whether any carbon is deposited.

Measurements of the resistivity of FC with varying concentrations of CNT will be attempted, using AC [133, 134] and DC [135] methods, to determine whether the CNT actually have an affect on the resistivity. The AC method uses a lock-in amplifier to measure the impedance of the FC samples and the DC method uses a custom made printed circuit board (PCB), composed of a chain of low noise op-amps, to measure the resistivity of the samples.

Even if the CNT affect the resistivity of the FC, it does not guarantee that this will actually reduce the surface charge deposited on fused silica. Kelvin probe measurements of a one inch silica disc after being cleaned with an FC solution with no CNT and an FC solution with a high concentration of CNT will be presented. A comparison between the two sets of measurements will be made to determine whether FC mixed with CNT actually affects the surface charge density deposited on fused silica.

Measurements of the surface charge density after coated and bare fused silica samples had been exposed to a corona discharge in nitrogen gas are also presented in this chapter [76]. These measurements are then used to calculate the resulting charging noise that would be expected to be observed in aLIGO to show that the charge has been reduced to a level that would not limit aLIGO. Transmissivity measurements of a fused silica sample with a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer reflective coating, that had undergone 20 exposures of the corona discharge, were taken using the ellipsometer set up described in Chapter 4. These measurements were then compared to a reference sample that had not been exposed to the corona discharge to show that there was no measurable change in transmission after exposure to the

corona discharge. The analysis carried out on the samples was the same as the analysis carried out in Chapter 4 for the glow discharge procedure.

5.2 SEM Images of Silica Discs After Being Cleaned with First Contact

It is a concern that the FC that is mixed with CNT may deposit carbon on the surface of the silica optics. This could increase the absorption of the mirrors which could lead to the optical coatings being damaged or cause thermal lensing problems. SEM images of the surfaces of two silica discs were taken, using a Hitachi TM-1000 Tabletop SEM, before and after being cleaned with different types of FC. FC with multi-walled CNT mixed in the FC solution was used on one disc while ordinary FC, with no CNT, was used on the other disc. Both discs were deliberately scratched at the edge so that the same area could be imaged before and after. Images were taken in an area approximately $5\text{ mm} \times 5\text{ mm}$ on both of the discs.

Figures 5.1 and 5.2 show the surface of a silica disc before and after it had been cleaned with ordinary FC at a magnification of 800. The debris from the scratch made on the sample is clearly visible in *Figure 5.1*, however, there is no debris shown in *Figure 5.2* because all of the loose silica particles have been removed by the FC. It can be clearly seen from *Figure 5.2* that the FC has left the surface of the sample very clean.

Figures 5.3 and 5.4 show images of the surface of a silica disc before and after it had been cleaned with FC mixed with CNT, also at a magnification of 800. The concentration of CNT is given in percentage of the overall mass of the solution. The mass percentage of the solution used here was $4.02 \times 10^{-2} \%$. The images show that, once again, the FC has removed the loose silica particles from the surface of the sample and left it very clean. It was noticed that a lot of carbon was trapped in the FC film that was peeled off the sample. These specs of carbon were visible to the naked eye which suggests the CNT might have a




Figure 5.1: *An SEM image of a silica disc before ordinary FC was applied. The disc was scratched so that there would be some debris present on the surface of the disc and also so that the same location on the disc could be imaged after removing the FC.*




Figure 5.2: *An SEM image of a silica disc after ordinary FC has been removed from its surface. The debris has been removed by peeling off the film of FC deposited on the surface of the disc.*

tendency to clump together. From these SEM images there certainly doesn't seem to be any large deposits of carbon left on the surface of the silica sample, however, whether individual strands of CNT are deposited on the surface of the sample will need to be investigated further. Since clumps of carbon can be seen randomly scattered through the FC, it can be assumed that the CNT are not distributed evenly throughout the FC. It is possible that in the small area where the images were taken, no CNT have been observed because there wasn't a very high density of CNT at that particular place on the sample. For more concrete confirmation, more images should be taken all over the sample to ensure that patches of carbon are not found. These images should be at a higher magnification since individual CNT usually have diameters which are nanometers in scale but have lengths that are microns in scale. This was not carried out for this investigation due to time constraints. The other SEM images taken for this investigation can be seen in [Appendix B](#).




Figure 5.3: *An SEM image of a silica disc before FC with multi-walled CNT was applied.*




Figure 5.4: *An SEM image of a silica disc after FC with multi-walled CNT has been removed from its surface. The debris has been removed and there doesn't appear to be any CNT deposited on the surface of the disc.*

5.3 Resistivity Measurements of Different Types of First Contact

The resistivity of FC solutions with varying concentrations of CNT were measured to ascertain whether mixing CNT in the FC solutions has any effect on the conductivity of the FC. The resistivity was measured using both an AC and a DC method to check that the results were consistent.

5.3.1 The First Contact Samples

In order to measure the resistivity of the different types of FC, a sample was made by taking a piece of blank Veroboard, approximately $50 \text{ mm} \times 50 \text{ mm}$, and scoring a line down the middle of the board. This created a board with two copper pads, that were electrically isolated from each other, that could be used to solder wires. The FC was poured across the gap in the board and left to dry so that an electrical connection was made between the two copper pads through the FC. When the FC sample was dry, a potential difference could then be applied across the sample creating a current which could be measured.

Four types of FC were used in the measurements presented in this chapter; FC with a CNT mass percentage of $4.02 \times 10^{-2} \%$, FC with a CNT mass percentage of $1.11 \times 10^{-2} \%$, FC with a CNT mass percentage of $2.67 \times 10^{-3} \%$ and FC with no CNT. A picture of a sample is shown in *Figure 5.5*.




Figure 5.5: A photograph of one of the FC samples used in the resistivity measurements.

As well as the measured resistance of the sample, R , the resistivity, ρ , of the FC sample is determined by the gap between the two copper pads, l , and the area of the electrodes at each end of the sample, A ,

$$\rho = R \frac{A}{l}. \quad (5.1)$$

It will be assumed that the copper electrodes are parallel and have an area equal to the thickness of the copper layer multiplied by the length of the FC sample. To measure the gap between the copper pads for each sample an instrument called Form Talysurf 50 was used. The Talysurf instrument is primarily used for measuring the surface roughness of objects, however, it can be easily used to measure the gap distance between the copper pads in the samples used here. The Talysurf drags a probe along the surface of the object being measured and determines changes in height along the sample's surface from the force acting on the probe. An example scan of one of the FC samples

is shown in *Figure 5.6*.




Figure 5.6: An example of a Talysurf scan of a Veroboard sample used in the resistivity measurements of FC.

To determine the thickness of the copper layer, a piece of Veroboard was measured that had tracks of copper wet etched. This measurement is shown in *Figure 5.7*. The copper layer was measured to be approximately $30\text{ }\mu\text{m}$ thick. The length of the FC strip was measured using digital calipers. Table 5.1 shows the gap distance measurements and FC strip measurements for each of the FC samples.

To estimate an error on the Talysurf measurements the thickness of a 1 mm slip gauge was measured. The thickness of the slip gauge should be very accurate so any deviation from the true thickness of the slip gauge will give the error of the Talysurf instrument. Several measurements of the slip gauge gave values between $\sim 0.9\text{ mm}$ and $\sim 1.1\text{ mm}$, therefore, an error of $\pm 10\%$ will be assigned to the Talysurf measurements. It will be assumed that the uncertainty in the length of the FC sample is negligible compared to the Talysurf measurements.




Figure 5.7: A Talysurf scan of the surface of a PCB board with two tracks of copper etched off. The thickness of the copper layer was taken to be approximately $30\mu\text{m}$

Table 5.1: Measured gap distances between the copper pads of each First ContactTM sample and the length of each First ContactTM sample

Mass Percentage of CNT (%)	Gap Distance (mm)	FC Sample Length (mm)
4.02×10^{-2}	0.237	13.9
1.11×10^{-2}	0.373	11.8
2.67×10^{-3}	0.220	13.4
0	0.245	12.0

5.3.2 DC Resistivity Measurements

To measure the DC resistivity of the FC solutions a high voltage was applied across the FC sample to produce a DC current. However, the FC is so resistive that the current produced was very small even though a high potential difference was used. The current had to be converted into a voltage, using a high impedance amplifier, and then amplified, using a custom made PCB, so that a measurable signal could be produced.

5.3.2.1 Experimental Set Up

Due to the small currents produced, a PCB had to be created, which was designed using Eagle PCB design software, that minimised noise from leakage currents. A method of reducing leakage currents is to create a ground ring [136], or guard ring, around the tracks of the circuit. Using this method means that any external currents will flow to ground rather than into the circuit.

The circuit has three stages of signal amplification. The first is a current-to-voltage converter which, as well as amplifying the detected signal, converts the current to a voltage using a TC7652 chopper stabilised op-amp. The chopper amplifier has minimal drift due to temperature changes, therefore, any drift in the signal should only be due to the resistivity of the sample changing with temperature. The voltage signal is then amplified further at two low gain stages in the circuit using OPA37 op-amps. Both of these stages each has a gain of 100. This amplified the signal enough so that it could be measured using a Keithley 2000 multimeter which was connected to a PC, through a general purpose interface bus (GPIB) connection, to record measurements using LabVIEW. A circuit diagram is shown in *Figure 5.8* and a photograph of the PCB is shown in *Figure 5.9*.

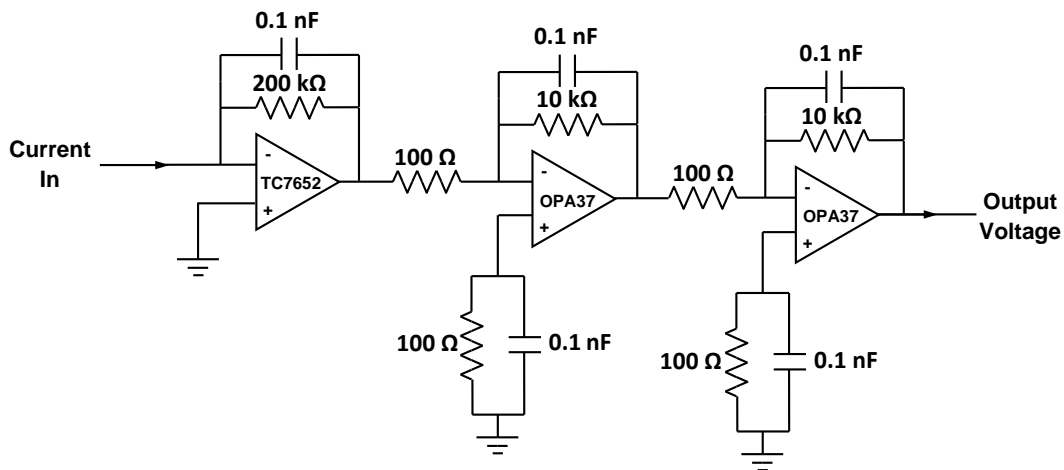


Figure 5.8: A schematic of the circuit used to measure the DC resistance of different types of FC.




Figure 5.9: A photograph of the complete PCB used to measure the DC resistance of different types of FC.

Since resistivity is dependant on temperature, the temperature was monitored during measurements using a PT-100 temperature sensor. The sensor was also monitored using the Keithley multimeter mentioned previously. A Matsusada Precision Inc. AMT-1.5B40 high voltage amplifier was used in the experimental set up to apply a voltage across the samples.

The circuit was tested by measuring the resistance of two $5\text{ G}\Omega$ thin film resistors connected in series. The test measurement was left running overnight to see if the signal had a significant drift. The average resistivity over the course of the measurement was $9.87\text{ G}\Omega$. The average temperature over the course of the measurement was $21.6\text{ }^{\circ}\text{C}$. The resistivity value measured is slightly outside the 1% error quoted on the resistors and there is some drift in the data. This could indicate that there is some leakage current in the set up since the resistivity seems to decrease over time. *Figure 5.10* shows the resistance measured over the course of the measurement and *Figure 5.11* shows a plot of the temperature against resistance. There does not appear to be an obvious temperature correlation from *Figure 5.11*.




Figure 5.10: A plot of the measured resistance of two $5\text{ G}\Omega$ thin film resistors, connected in series, over time. The average resistance measured is $9.87\text{ G}\Omega$.




Figure 5.11: A plot of the measured temperature over the course of the test measurement against the DC resistance measured over the course of the measurement. There is not a clear correlation between temperature fluctuations and changes in resistance at this level.

5.3.2.2 Results

The resistivity of four different types of FC were measured for this investigation. Three of these had CNT mixed in the FC solution and the fourth did not contain any CNT. For each measurement, the sample was placed in a metal box which was then sealed air tight as well as possible. The box had a gas hose attached so that dry nitrogen could be released into the box. This was to reduce any surface moisture that could be present on the samples and also to reduce the humidity of the atmosphere inside the box. The nitrogen flow was turned off after 15 mins and 75 V was applied across the sample. The sample was left for about half an hour while the output voltage from the PCB circuit was monitored. The results of the resistance values measured are shown in Table 5.2. Any measurements in Table 5.2 that are labeled *Null* are measurements where the measured signal decayed to zero before the end of the measurement. For one of the samples a plot of the resistance measured is shown in Figure 5.12 and a plot of temperature against resistance is shown in Figure 5.13.

Table 5.2: Measured resistivities of First ContactTM with varying concentrations of CNT using DC method

Mass Percentage of CNT (%)	Resistance (Ω)	Resistivity (Ω m)	Ave. Temperature ($^{\circ}$ C)
4.02×10^{-2}	513	0.90	22.50
1.11×10^{-2}	<i>Null</i>	<i>Null</i>	22.69
2.67×10^{-3}	<i>Null</i>	<i>Null</i>	22.76
0	262	0.39	22.44

From the DC resistivity measurements it would appear that the resistivity of the ordinary FC is less than the type mixed with CNT. Seeing as the resistance changed so much during the experiment, the measurements can only really be taken as an order of magnitude estimate of the resistivity. It could be




Figure 5.12: A plot of the measured resistance of the FC sample with a CNT mass concentration of $4.02 \times 10^{-2} \%$.




Figure 5.13: A plot of temperature against resistance during the measurement of the FC sample with a CNT mass concentration of $4.02 \times 10^{-2} \%$. The plot seems to indicate that the changing resistivity observed could be due to changes in temperature.

that with a smaller current being measured the resistance is genuinely changing with temperature. However, since the test measurement did not show such an obvious correlation with temperature, it is possible that this is just a coincidence. It is also likely that there is significant error in the resistivity values calculated due to the assumptions made in this experiment, for example the electrode area is likely not to be as perfect as assumed. As shown in *Figure 5.6*, the thickness of the board appears to increase towards the edge of the gap, although it is unknown if the extra material is in fact copper or other debris. The inclusion of CNT to the FC solution would have to have a significant effect on the resistivity of FC to be considered in aLIGO. From these measurements it can only be assumed that CNT have little to no effect on the resistivity of FC.

5.3.3 AC Resistivity Measurements

The resistivity of the different types of FC were also measured using an AC method to check that the measurements were consistent. For these measurements an AC voltage signal was applied across the FC sample and the resulting current was measured using a lock-in amplifier. If it is assumed that the FC sample consists of a resistor and a capacitor in parallel then the impedance, Z , of the set up can be expressed as,

$$Z = \frac{R}{\omega^2 R^2 C^2 + 1} - \frac{j\omega R^2 C}{\omega^2 R^2 C^2 + 1}, \quad (5.2)$$

where R is the resistance of the FC sample, C is the capacitance of the FC sample and ω is the angular frequency of the applied voltage signal. As the frequency tends to zero, the impedance should be equal to the resistance of the FC sample.

5.3.3.1 Experimental Set Up

To create a high voltage AC signal a Hewlett Packard 33120A waveform generator was connected to the high voltage amplifier mentioned in Section 5.3.2.1.

A program was written in LabVIEW that swept through different frequencies on the waveform generator and recorded the resulting current measured by the lock-in amplifier. The temperature was monitored in the same way as described in Section 5.3.2.1.

As with the DC measurements, a test of this method was carried out on circuit components with known values. A $5\text{ G}\Omega$ thin film resistor was connected in parallel with a 22 pF capacitor and the impedance of the circuit was measured for the frequency range $1 - 30\text{ Hz}$. As the frequency decreases the impedance of the circuit becomes more dominated by the resistor. Therefore, as the frequency tends to zero the impedance should level out at the resistance of the resistor in the circuit. A plot of the impedance measured for this set up and the theoretical impedance, calculated using Equation 5.2, is shown in Figure 5.14.




Figure 5.14: A plot of the measured impedance of a $5\text{ G}\Omega$ resistor and a 22 pF capacitor connected in parallel, red, and the theoretical impedance, blue, over a range of frequencies. The impedance at 1 Hz is measured to be $4.13\text{ G}\Omega$.

Figure 5.14 shows that the resistance starts to level off to the value of the

resistor in the circuit as expected. The value of the impedance at 1 Hz is measured to be $4.13 \text{ G}\Omega$ which is significantly less than the true value of the resistor. The resistor is quoted as only having a 1 % error in its resistance value. There could be other capacitances in the set up that haven't been considered that could be affecting the measurements slightly, for example the capacitance of the coaxial cables used. Based on this test measurement it will be assumed that the uncertainty in the measurements of the resistance using this method could be as high as 20 %.

5.3.3.2 Results

The same four FC samples used for the DC measurements were used here. The method of preparing the samples with dry nitrogen for measurements was the same as described in Section 5.3.2.2. A sine wave voltage was applied across the sample with an amplitude of 90 V ($63.6 \text{ V}_{\text{rms}}$) and the frequency range swept through was from 30 Hz down to 1 Hz. The results of these measurements are shown in Table 5.3 and a plot of the data for one of the FC samples is shown in Figure 5.15.

Table 5.3: Measured resistivities of First ContactTM with varying concentrations of CNT using AC method

Mass Percentage of CNT (%)	Resistance ($\text{G}\Omega$)	Resistivity ($\text{G}\Omega \text{ m}$)	Ave. Temperature ($^{\circ}\text{C}$)
4.02×10^{-2}	34.2	0.058	22.16
1.11×10^{-2}	61.2	0.059	22.46
2.67×10^{-3}	48.7	0.089	22.37
0	33.5	0.049	22.40

From the results presented in Table 5.3 it appears as if the ordinary FC is still the least resistive FC sample. However, when comparing just the FC samples mixed with CNT, the sample with the highest mass percentage of




Figure 5.15: A plot of the measured impedance of the FC sample with a CNT mass concentration of $4.02 \times 10^{-2} \%$. The impedance at 1 Hz was measured to be 34.2 GΩ. The faint red lines on the plot show the region of uncertainty.

CNT is the least resistive.

Clearly there needs to be a lot of improvement in this resistivity experiment to achieve consistent results. It is slightly concerning that the data from the two measurements do not compare well and it is unknown why there is such a deviation in the results. It is very difficult to measure the impedance in this instance since it does not appear from *Figure 5.15* that the measurement starts to level off at low frequency. This would indicate that it is very likely that the value taken as the impedance for the samples underestimates the true impedance of the samples. Measurements were not taken below 1 Hz because it was found that measurements at these frequencies were generally very noisy. It is likely that poor sample preparation may have been the biggest problem with this experiment as both methods seemed to agree well when measuring known circuit components. Just taking the data of each measurement method

individually, it appears as if CNT does not have a significant effect on the resistivity of FC. However, seeing as the measurements of the AC and DC method differ so greatly when compared, all that can be said from these measurements is that they are inconclusive.

The statistical tests used in the previous chapter were not used here because there was not enough data to carry out the tests successfully and the quality of the measurements were not as good as they could have been. If there was more time available it is possible that better measurements could have been taken to allow thorough analysis.

5.3.4 Surface Charge Density Measurements

Since the resistivity measurements did not give a clear indication of the effect of mixing CNT in FC, Kelvin probe measurements were taken of a one inch fused silica disc after it had been cleaned with FC with and without CNT. Even if the CNT does affect the resistivity of the FC polymer, the important point for aLIGO is that the CNT actually causes a significant decrease in the deposition of surface charge.

A one inch silica disc was cleaned with ordinary FC solution and a FC solution mixed with multi-walled CNT. The mass percentage of the CNT was $4.02 \times 10^{-2} \%$. A 0.5 ml volume of FC solution was deposited on the surface of the silica disc and spread over the surface of the optic with a cotton bud. Care was taken not to touch the surface of the optic with the cotton bud. Once the surface of the silica disc was evenly covered with FC solution it was left for 24 hours to dry. Once the FC film was completely dry, a scalpel blade was run along the edge of the disc to remove any FC that may have dripped over the edge of the sample face and onto the barrel of the disc. A piece of Kapton tape was placed on the FC film and used to peel off the FC film. The sample was then moved under the Kelvin probe for a measurement. It should be noted that all these measurements were conducted at atmospheric pressure

since this would mimic a more realistic cleaning scenario for aLIGO rather than performing the measurements under vacuum. This measurement was repeated 10 times for both types of FC. The results are shown in *Figure 5.16*.




Figure 5.16: A plot of the Kelvin probe measurements of the charge deposited by FC with and without CNT. The charge deposited by FC with and without CNT peeled off a silica disc was measured ten times for each FC type.

The Kelvin probe voltage was not turned into a surface charge density because the probe-sample separation distance had to be increased from 1 mm to 6 mm to avoid the probe signal being saturated for every measurement. At this probe-sample separation distance some of the assumptions made to obtain Equation (3.16) do not hold up and the values calculated would likely be inaccurate. Converting the Kelvin probe signal into a surface charge density is not essential for these measurements and a relative measurement is sufficient.

All the measurements of the surface charge on the FC mixed with CNT caused the Kelvin probe signal to saturate. Only a couple of the ordinary FC Kelvin probe measurements saturated. The measurements in *Figure 5.16* show that lower Kelvin probe voltages, and therefore lower surface charges, tended to be measured from FC without CNT. Thus indicating that there is no benefit

from adding CNT to FC.

There are no errors shown on the plot because the uncertainty in the Kelvin probe voltage is likely small ($< 10\%$) and the effects of stray capacitance on the Kelvin probe signal, as shown in Chapter 3, are most likely negligible as well.

5.4 Mitigating Surface Charge Using an AC Corona Discharge

Another proposed method of reducing surface charge during the cleaning procedure of the aLIGO optics is to use ionised gas. As mentioned previously, researchers have already had success removing excess charge using an N_2 ioniser gun. Ioniser guns use a sharp point, or emitter, to ionise whatever gas is flowing over it. The flowing gas carries the ions out of the device and they are drawn towards the charged surface to be neutralised. In this section it will be shown that a commercially available anti-static corona bar could be used to mitigate charge during the cleaning procedure of a gravitational wave detector optic. The physical mechanism that ionises the gas in the anti-static bar is the same as for the ioniser gun, therefore, the results of the investigation presented here should be a good indication of the performance of an ioniser gun.

5.4.1 A Corona Discharge

In Chapter 4 it was mentioned that there are several types of electric gas discharge that can be created. In Chapter 4 the focus was on glow discharges, but in this chapter a corona discharge will be discussed. A corona discharge is created in areas of high electric fields, usually near sharp edges and corners, in electrically stressed gases preceding electrical breakdown of the gas. The corona discharge is classified as a dark discharge because the currents used for this type of discharge are very low, however, if higher currents are used then a slight glow can be seen around the discharge electrodes.

The reason that such high electric fields are generated at sharp points and edges is because the electric field increases with decreasing surface area. If Gauss's law is considered, given in Equation (5.3), it is clear that as the radius of curvature, r , of a sharp point decreases, the electric field, E , will increase as the inverse square of r .

$$\int E \, dA = \frac{Q}{\varepsilon_0} \quad (5.3)$$

$$E = \frac{Q}{2\pi\varepsilon_0 r^2} \quad (5.4)$$

In the equations above ε_0 is the permittivity of free space, Q is charge and it is assumed that the sharp point has the surface area of a half sphere ($2\pi r^2$).

Corona discharges are widely used in many different types of industries, however, they are primarily used for static elimination [118]. In the manufacturing industry where there is significant contact in the production process, for example between sheets of paper and rollers during paper printing, there will be some form of corona anti-static system present to neutralise items on the production line. In industries that use plastic injection molding or printing, failure to reduce excess charge can lead to machines on the production line being damaged, or extremely high voltages can build up on the production machines or the products themselves. Therefore, since corona anti-static bars are so common, it was decided that for this investigation an off the shelf anti-static device would be used.

5.4.2 Experimental Set Up

The anti-static bar consists of a row of equally spaced sharpened pins that have a high voltage AC signal applied to them. This creates a very high electric field which can effectively ionise the gas which surrounds the pins. The anti-static bar that was used for this investigation was a Static Clean International [137] BR2200 rectangular static eliminator bar with an additional gas feed attachment. The gas feed attachment was for directing a flow of either

nitrogen or argon gas to the pin electrodes so that the air around them was displaced. A Static Clean TSN75 AC power supply was also used to supply the pin electrodes with 7.5 kV. For this procedure the samples were discharged at atmospheric pressure on the lab bench. The static bar was secured to a clamp stand so that the bar could be moved across the sample at a fixed height of approximately one inch above the sample surface.




Figure 5.17: *The corona discharge experimental set up.*



Figure 5.18: *A close up of one of the pins that creates the corona discharge.*

5.4.3 Results

5.4.3.1 Surface Charge Measurements

For this investigation the same two sample types were used as in Chapter 4; a one inch bare fused silica disc and a fused silica sample with a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer reflective coating. The silica disc was used to show that the corona discharge could consistently reduce the surface charge density on silica and the coated sample was used to show that the corona discharge would not damage the sensitive multi-layer reflective coating.

The bare silica sample was charged positively or negatively as described in Chapter 4, while the coated sample was charged using FC. The recommended method of cleaning the aLIGO test masses with FC, as specified by the LIGO Scientific Collaboration [138], was followed as closely as possible. A scan of the sample was taken with the Kelvin probe, using the same scan sizes as

stated in the previous chapter, so that the charge on the surface of the sample before the discharge process could be calculated. Once this had been done, the sample was moved under the corona bar. The gas flow was turned on for a few seconds before the static bar was turned on in order to displace any air at the pin electrodes. The bar was then turned on and moved across the sample in a single sweeping motion. The surface charge density of the sample was then measured again. Examples of Kelvin probe scans of the coated silica sample before and after the charge mitigation process are shown in *Figures 5.19 and 5.20*.




Figure 5.19: *Kelvin probe scan of the surface charge density across the coated silica sample before the corona discharge procedure. The average surface charge density across the scanned area is greater than $-5.80 \times 10^{-8} \text{ C/m}^2$.*




Figure 5.20: *Kelvin probe scan of the surface charge density across the coated silica sample after the corona discharge procedure. The average surface charge density across the scanned surface is $3.33 \times 10^{-9} \text{ C/m}^2$.*

This discharge process only takes a few minutes to complete and is easily incorporated into the LIGO Scientific Collaboration’s aLIGO test mass cleaning procedure. This charge mitigation procedure was repeated ten times for the bare silica sample, using both argon and nitrogen, and twenty times for the coated sample using nitrogen. The coated sample was discharged with nitrogen because it was found to be the more effective discharge gas from the silica disc measurements. The results obtained are presented in Table 5.4.

Table 5.4: Results from the corona discharge investigation using argon and nitrogen gas.

Sample	$ \bar{\sigma} _{before}^a$ (C/m ²)	$ \bar{\sigma} _{after}$ (C/m ²)	Stand. Dev. (C/m ²)	Sys. Error (C/m ²)
Silica (Ar)	$> 2.8 \times 10^{-8}$	1.60×10^{-9}	9.30×10^{-10}	$\pm 1.39 \times 10^{-9}$
Silica (N ₂)	$> 2.8 \times 10^{-8}$	5.65×10^{-10}	5.62×10^{-10}	$\pm 4.92 \times 10^{-10}$
Coated Silica	$> 5.8 \times 10^{-8}$	3.59×10^{-9}	3.66×10^{-9}	$\pm 3.13 \times 10^{-9}$

^a Only a lower limit is given for $|\bar{\sigma}|_{before}$ because the Kelvin probe signal had saturated

As in Chapter 4, a plot of the charging noise that could be expected in aLIGO is plotted using the surface charge density measurements obtained. A plot of the total aLIGO noise budget, the charging noise after the corona discharge procedure and the charging noise before the corona discharge procedure are shown in Figure 5.21. Only a lower limit estimate of the charging noise predicted before the charge mitigation procedure could be given since the Kelvin probe signal had saturated during these measurements. As can be seen in Figure 5.21, the resulting charging noise from the corona discharge procedure investigated here would not limit the low frequency performance of aLIGO.

The two flat circular faces of the aLIGO test masses could be coated in FC in order to take preliminary test mass positioning tests. This means that there would be a lot of charge deposited on the surface of the test mass which faces the reaction mass. As discussed in Chapter 4, in order to effectively reduce surface charge on the non-reflecting face of the aLIGO test masses, the ionised gas will have to pass through the small gap between the test mass and reaction mass. To simulate this scenario, a scale model of this set up was created using two three inch fused silica discs spaced approximately 1.1 mm apart. This was the same set up as described in the previous chapter. The corona bar was held approximately 10mm from the edge of the two discs and left running for about 20 mins. This was carried out on a lab bench at atmospheric pressure like the




Figure 5.21: Resulting charging noise after the corona discharge process, blue, compared to the aLIGO noise budget, black. A lower limit estimate of the charging noise before the charge mitigation procedure is also shown, green.

corona discharge process carried out on the one inch silica disc sample and the coated silica sample. Scans of the charged three inch fused silica sample before and after the corona discharge procedure are shown in *Figures 5.22 and 5.23*.

It can be seen that most of the surface charge is reduced in the two silica disc set up. It is possible that leaving the corona bar running longer will reduce the charge completely across the whole sample surface, however, it was not possible to optimise the procedure for this particular set up due to time constraints. If it is not possible to reduce charge between the optics in this manner, it is likely that the reaction mass/test mass chain will be separated to allow the FC to be removed, and the charge to be mitigated, more easily.

The results show that the corona discharge charge mitigation process described is a consistent and reproducible way of reducing surface charge on bare fused silica and fused silica with a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer coating.




Figure 5.22: *Kelvin probe scan of the surface charge density across the three inch silica sample before excess surface charge is reduced using the corona bar. The probe signal had saturated, so the average surface charge density across the scanned area is greater than $1 \times 10^{-7} \text{ C/m}^2$.*




Figure 5.23: *Kelvin probe scan of the surface charge density across the three inch silica sample after using the corona bar to mitigate excess charge. The average surface charge density across the scanned surface is probably slightly higher than $3.1 \times 10^{-8} \text{ C/m}^2$ because the probe signal was still saturated at some points on the sample.*

The procedure is also very simple and should not be difficult to include in the cleaning process of the aLIGO optics. Inconsistencies between each run of the process probably arise from the fluctuating gas flow around the pin electrodes and from sweeping the bar across the sample at different speeds on each run.

Further investigation will be required in order to optimise the process for mitigating surface charge between the test mass and the reaction mass of future gravitational wave detectors. Here, the anti-static bar was left at one position for 20 mins, however, the charge might be reduced more effectively if the bar is slowly moved around the optics over a longer period. It is hard to recreate this situation in a lab environment with such a rudimentary set up as used here. In a real detector it might not be easy, or possible, to move an anti-static bar around a fully suspended optic. The ioniser gun might be more effective for this procedure since both the emitter electrode and the gas flow will be pointing in the same direction. When using the corona bar, the gas flow is

dissected across the pins and therefore it is difficult to direct the gas between the two optics.

5.4.3.2 Transmission Measurements

To make sure that the corona discharge did not cause a significant amount of damage to the TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer coating, the same transmission measurements that were made in the previous chapter were made for the coated sample used in the corona discharge investigation. As before, the transmission of the coated sample that had undergone the charge mitigation procedure 20 times was compared to the transmission of a coated sample that had not been exposed to the corona discharge to determine whether any statistically significant changes had occurred. The reflectivity of the TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ coating should be approximately 99.9 % at 1064 nm, therefore, if the sample is undamaged by the discharge process the transmission should be 0.1 %. The transmittance of the two samples were measured 40 times each in order to obtain two adequately sized data sets to analyse. The results of these measurements are shown in Table 5.5 and all the data are represented in the histograms in *Figure 5.24*.

Table 5.5: Results of the transmission measurements of the coated silica sample after the corona discharge procedure and a reference sample.

Sample	Average Transmission (%)	Standard Deviation	Error on Mean
Reference	1.23×10^{-1}	5.27×10^{-2}	8.34×10^{-3}
Corona	1.18×10^{-1}	7.40×10^{-2}	1.17×10^{-2}

The same Student's t-test and K-S test analyses used in the previous chapter were carried out on the transmission data. The results of the Student's t-test analysis can be seen in Table 5.6. This analysis shows that a high p-value was obtained which indicates that there is no statistically significant

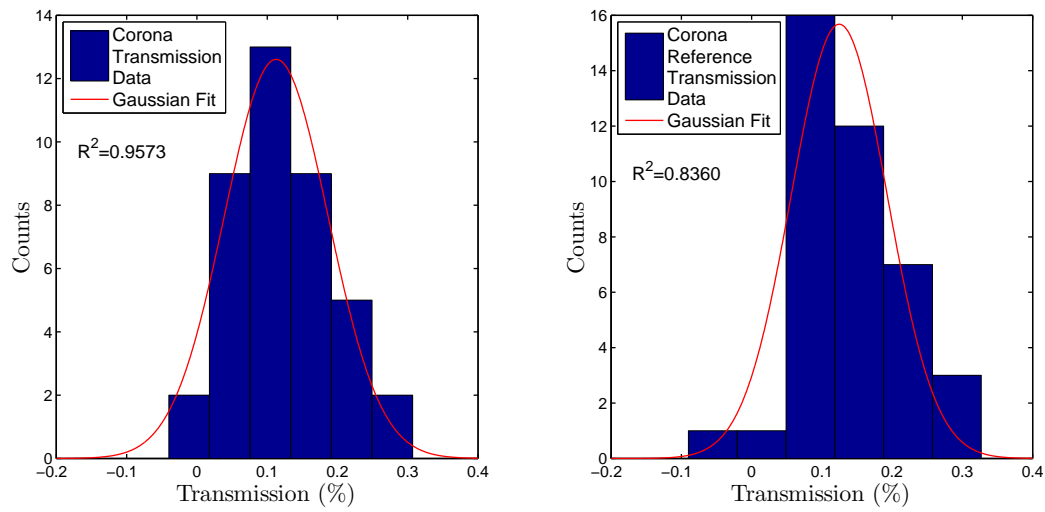


Figure 5.24: A histogram of all the transmission measurements of the corona discharge sample and the reference sample. The R^2 values of both the above histograms show that they fit well to a Gaussian distribution.

difference between the sample that was used for the corona discharge investigation and the reference sample which had not been exposed to the corona discharge. This suggests that the corona discharge procedure did not cause observable damage to the multi-layer reflective coating.

Table 5.6: Results from the Student's t-test analysis carried out on the transmission data of the reference sample and the corona discharge sample.

Samples	p-value	Confidence Intervals
Reference and Corona	0.65	-2.69×10^{-2} , 4.25×10^{-2}

The results of the K-S test are shown in Table 5.7. A high p-value was also obtained in this analysis when the transmission of the Corona sample and the Reference sample data sets are compared. This result is consistent with the one obtained from the Student's t-test analysis that the corona discharge procedure did not cause substantial observable damage to the TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer reflective coating. A plot of the empirical cumulative distribution functions of the Reference and Corona data is shown in Figure 5.25.

Table 5.7: Results from the Kolmogorov-Smirnov test analysis carried out on the transmission data of the reference sample and the corona discharge sample.

Samples	p-value	K-S Statistic
Reference and Corona	0.983	0.10




Figure 5.25: A plot of the empirical cumulative distribution functions of the transmission data obtained from measurements of the reference sample and the sample used during the corona discharge investigation. The y axis of the plot is the empirical cumulative distribution function, F , of the Corona and Reference data sets, x .

As stated in Chapter 4, these measurements were taken only to check for large changes in the transmission of the coated sample. The aLIGO mirrors require an absorption level less than 1 ppm, however, the measurements made in this investigation were only sensitive to changes in transmission of approximately 200 ppm.

5.5 Discussion

From the results presented in this chapter, it seems that CNT do not have a significant effect on reducing the amount of charge deposited by FC. For FC mixed with CNT to be considered for aLIGO, or any other future gravitational wave detector, it would have to deposit significantly less charge on silica than ordinary FC, which doesn't seem to be the case.

It is likely that the largest source of error in the resistivity measurements is from the design of the sample. A few other types of sample were tried but didn't work effectively. For example, some FC sandwiched between two copper plate electrodes was attempted but the FC would not dry properly even after being left for weeks. The FC requires a large surface area to be exposed so that the solvents in the FC evaporate and this is why the FC did not dry properly. Electrode plating techniques and silver paint were also used to try and deposit electrodes on the FC once it had dried, however, since FC is designed to be very unreactive it was found impossible to deposit electrodes on the surface of the FC film. Four point probe measurements were also attempted using facilities available in the particle physics group at the University of Glasgow. Unfortunately this set up is designed for measuring resistivities much lower than those expected in the work presented here. As a result, the measurements obtained from the four point probe were often inconsistent. The question of whether CNT effects the resistivity of FC was not answered in this work, however, it was clearly shown that CNT mixed in FC do not have an effect on the amount of surface charge that is deposited on silica.

It seems like the most sensible solution to mitigating charge deposited by FC is to use an ionised gas. It has been shown that ionised nitrogen, created from a corona discharge, effectively reduces excess surface charge on silica. Measurements of the transmissivity of a fused silica sample with a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer reflective coating showed that there was no observable change in the transmittance of the coating at a sensitivity of 200 ppm. Thus

indicating that no significant damage was caused to the coating. There is also other work that indicates that nitrogen gas ionised by corona discharge does not effect an aLIGO style coating [75]. The only part of the procedure that needs more optimisation is mitigating charge between a detector test mass and reaction mass.

5.6 Conclusion

Unfortunately it was not possible to show whether CNT reduces the resistivity of FC. This is likely due to poor samples and the fact that it is very difficult to measure highly resistive materials. However, the point of adding CNT to FC was to see if it reduced the amount of surface charge deposited on silica. It was shown through Kelvin probe measurements that mixing CNT in FC does not decrease the amount of surface charge deposited on silica.

In this chapter it was shown that a procedure that utilises a device called an anti-static bar could be used to mitigate charge on gravitational wave detector optics. The method effectively reduces the surface charge density on fused silica with a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer reflective coating to a level that would not effect the low frequency performance of an advanced gravitational wave detector such as aLIGO. The method also partly reduced the surface charge between two three inch silica samples, however, this part of the procedure would have to be optimised for a real detector.

It was also shown that the corona discharge procedure does not cause any observable change to the transmission of a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer coating at a level of 200 ppm. As discussed in Chapter 4, there are other techniques that can measure changes in the absorption of optical coatings to a much higher sensitivity, however, it has already been shown in other work [75] that nitrogen ionised by corona discharge does not damage a TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer coating at a sensitivity level of less than 1 ppm.

Chapter 6

Developing a Torsion Balance for Measuring Charging Noise

6.1 Introduction

In order to make accurate estimates of the charging noise that could be expected in future gravitational wave detectors, a confirmed observation of it must be made. Charging noise is believed to have been detected during an incident at one of the LIGO detectors [139]. It was thought that the test mass of one of the detectors may have come into contact with a rubber earthquake stop when adjustments were being made to one of the vibration isolation systems. The resulting noise had the same frequency dependence as would be expected of charging noise and seemed to disappear once the vacuum chamber was brought back up to atmospheric pressure. However, since there was no measurement of the charge deposited on the test mass there was no way to accurately model the noise to see if it coincided with the level observed.

This chapter will detail the development of a torsion balance apparatus designed for measuring charging noise. A torsion balance is an instrument capable of measuring very weak forces, and so, should be capable of measuring the small electromagnetic fluctuations caused by surface charge moving around on the surface of a silica disc. With the incorporation of a Kelvin probe into

the torsion balance set up it will be possible to measure the surface charge density on a silica sample and, therefore, estimate the level of charging noise that would be expected to be observed.

In this chapter a description of the experimental set up will be given, which will include the design of the torsion bob, capacitive sensor circuit and the PID controller, and some initial measurements of surface and polarisation charge on a fused silica disc. The different noise sources which affect the low frequency sensitivity of the instrument will also be discussed and the method used to remove these noise sources from the raw data will be described. There will also be a discussion of some changes that were made to the initial set up in order to increase the sensitivity of the torsion balance. The chapter will conclude with a discussion of the design of a Mark II version of the instrument which should achieve greater sensitivity than the Mark I instrument.

6.2 The Torsion Balance

Torsion balances have been used in physics to measure extremely weak forces for over 200 years [140]. Torsion balances have classically been used to measure the force of gravity between masses in order to accurately calculate the Newtonian constant of gravitation [141, 142], however, some have been used to measure electrical phenomena [143], detect neutrinos [144] and to measure the expected residual acceleration of the LISA test masses [145].

Even though many different designs of the torsion balance have been documented, nearly all of them share some common design features. A typical torsion pendulum consists of a long thin fibre of tungsten, or a high Q material such as fused silica, which supports the bob of the pendulum. The bob is typically a dumbbell or cross like shape made of a light metal such as aluminium. Attached to the ends of the arms of the bob are test masses. When an external mass, which has a much larger mass compared to the bob, is placed near one of the test masses on the torsion balance, the bob will rotate as the gravitational

attraction pulls the test mass towards the larger external mass. The torsion balance has not been restricted to the conventional design described above. There are many modern torsion balance designs that use magnets [143] and even superconducting materials [146] to suspend the torsion balance bob.

The torsion balance's motion can be described as a damped oscillator which follows the equation,

$$I \frac{d^2\theta}{dt^2} + b \frac{d\theta}{dt} + \kappa\theta = \Gamma, \quad (6.1)$$

where I is the moment of inertia of the bob, θ is the angle through which the bob is displaced, t is time, b is the damping constant, κ is the stiffness of the fibre and Γ is the applied torque. The external torque acting on the test mass causes the bob to be rotated to a point where it is cancelled by the restoring torque of the suspending fibre. The fibre acts like a spring and therefore should have a small spring constant, or restoring torque, which is represented by κ in Equation (6.1). The external force being applied to the test mass can be calculated from the angle that the torsion balance rotates through since it is proportional to the applied external force. A schematic of the basic torsion balance set up described above is shown in Figure 6.1.




Figure 6.1: A schematic of a classic torsion balance set up. The torsion bob is suspended by a thin fibre and when an external torque, Γ , is applied to the test masses, the bob rotates.

6.3 Experimental Set Up

6.3.1 The Vacuum Tank and Inner Structure

The torsion balance used for the work presented in this thesis is housed in a vacuum chamber, shown in *Figure 6.2*, that can achieve pressures of less than 5×10^{-5} Torr using a BOC Edwards XDS-10 dry scroll pump and a Varian Turbo-V 301 navigator pump. A Varian ConvecTorr gauge was used to monitor the pressure inside the vacuum tank down to 1×10^{-3} Torr and a Varian IMG-100 inverted magnetron gauge was used for pressures less than 1×10^{-3} Torr.




Figure 6.2: A photograph of the vacuum tank and inner structure. Also labeled in the photograph are the torsion balance translation and rotation stages, the eddy current damper and the torsion bob. These will be described in the following sections.

Inside the vacuum chamber there is an inner structure from which the bob is

suspended. The position of the bob can be adjusted using two Thorlabs PT1/M translation stages, in an x-y configuration, and a Thorlabs CR1/M continuous rotation stage¹. The translation stages allow a 25 mm range of motion and the rotation stage can be continuously rotated. The translation stages have to be positioned correctly before the air is pumped out of the chamber, however, the rotation stage can be controlled using a stepper motor attached to the rotation stage's adjusting arm whilst under vacuum. The stepper motor is controlled using a Saia-Burgess Samotronic 101 stepper motor driver which is connected to a PC through a peripheral component interconnect (PCI) 6229 ADC so that it can be programmed using LabVIEW. A picture of the stepper motor set up can be seen in *Figure 6.3*.




Figure 6.3: A photograph of the stepper motor set up.

Each step of the stepper motor is 1.8° , but in order to obtain a finer resolution of the rotational position of the torsion balance, a gear box is used to reduce each motor step by a factor of 100. The gear box is connected to the adjusting arm of the rotation stage via a universal joint. If the rotation stage

¹The rotation stage shown in *Figure 6.2* is a Thorlabs PR01 high precision rotation stage. This was replaced with the CR1/M stage because the PR01 could not be rotated continuously.

adjusting arm is rotated 1° the resulting rotation of the stage is 0.042° . In order to rotate the adjusting arm of the stage by 1° the stepper motor must take approximately 56 steps, therefore, one step of the stepper motor results in a $7.5 \times 10^{-4}^\circ$, or $13.1 \mu\text{rad}$, rotation of the stage.

There is a fixed platform near the bottom of the tank onto which two Thorlabs MTS50-Z8E motorised translation stages, in an x-y configuration, are attached. The motorised translation stages are used to move charged silica samples near the bob of the torsion balance in order to make a measurement. Both of these stages have a range of 50 mm. For the measurements made in this chapter a small copper rod was fixed to the motorised translation stages instead of a silica sample. The copper rod was connected to an analogue output channel so that sinusoidal voltage signals could be applied to the rod using LabVIEW. These motors can be seen in *Figure 6.4*.

Also attached to the fixed platform are two columns. One column has two sensor plates attached to it and the other has two actuator plates attached to it, shown in *Figure 6.4*. The sensor plates are used to determine the position of the bob and the actuator plates are used to apply electrostatic servo control to the bob. The position sensing and servo control of the torsion balance will be discussed in detail in later sections.

6.3.2 Torsion Balance Design

The torsion balance used for the work presented in this thesis consists of a light bob, shown in *Figure 6.4*, made of aluminium, a tungsten fibre, an eddy current damper, a thicker tungsten wire and a clamping block which is bolted to the rotation stage.

The torsion bob is made up of a small clamping block for the tungsten fibre and four arms. There are two copper plates attached to the ends of two opposite arms which are used for position sensing and applying servo control. There are two silica discs on the other arms that are used for charge measurements.




Figure 6.4: A photograph of the torsion balance bob.

The bob is suspended by a tungsten fibre $50\ \mu\text{m}$ in diameter and $\sim 0.41\ \text{m}$ in length. The restoring torque of the fibre is proportional to the fourth power of the fibre's diameter, therefore, the sensitivity of the instrument increases significantly if the diameter of the fibre is kept to a minimum [143]. The minimum radius, r , of tungsten wire needed to support the bob was deduced from the yield tensile stress, σ , and the mass, m , supported by the fibre,

$$r = \sqrt{\frac{mg}{\pi\sigma}}. \quad (6.2)$$

The yield tensile strength of drawn tungsten wire is taken as $450\ \text{MPa}$ [147]. The radius of the fibre should be greater than $15\ \mu\text{m}$ to avoid breaking the fibre or deforming it.

The tungsten fibre was attached to the bottom of a copper damper disc, used to damp the simple pendulum mode of the torsion balance, which itself was suspended from the main clamping block, at the rotation stage, by a thicker

tungsten fibre $100\mu\text{m}$ in diameter and $\sim 0.13\text{m}$ in length. The reason for using the shorter, stiffer, wire is because a torsionally stiffer wire should have a simple pendulum mode, and torsional mode, that are higher in frequency than the longer, less stiff, wire used to support the bob. This means that the pendulum and torsional modes of the thicker wire won't couple with the pendulum and torsional modes of the thinner fibre causing unwanted oscillation.

An eddy current damper [148], positioned near the top of the pendulum, is used to damp the pendulum mode of the torsion balance. The eddy current damper is composed of a copper damper disc which has fibre clamps on the top and bottom. The copper disc is enclosed in a housing made of iron which has two neodymium iron rare earth magnets positioned directly above and below the damper disc. As the pendulum mode of the torsion balance is excited, eddy currents are generated in the copper disc which damps the pendulum motion of the bob while the torsional motion of the bob remains unaffected due to the symmetry of the damping disc. A photograph of the eddy current damper is shown in *Figure 6.5* and a cross sectional diagram is shown in *Figure 6.6*.




Figure 6.5: A photograph of the eddy current damper.




Figure 6.6: A cross sectional diagram of the eddy current damper.

6.3.3 Capacitive Sensor Design

The angular position of the bob is measured using a capacitive sensor which was designed by Dr. G. D. Hammond and built by Mr. I. McCrindle.

The copper plates that are used for position sensing and applying actuation, located on the bob, sit between two other fixed copper plates, shown in *Figure 6.4*. On one side of the bob, the two fixed copper plates are for capacitively sensing the position of the bob. The two fixed plates on the opposite side are for applying servo control. The pair of sensing plates create what is commonly known as a capacitive bridge.

An AC drive signal is applied down the fibre of the torsion balance, using a Stanford SR830 lock in amplifier, in order to provide a signal that can be used for capacitive sensing. This voltage signal has an amplitude of 1 V and a frequency of 100 kHz, which is much greater than the resonant torsional frequency of the bob.

The AC drive signal flows through a 1:1:-1 transformer that consists of two secondary coils and one primary coil. The positive terminal of the first secondary coil is attached to one sensor plate and the negative terminal of the other secondary coil is attached to the other sensor plate. The drive signal being applied down the fibre of the torsion balance allows current to flow through the two secondary coils of the transformer, this is defined by the capacitance between the bob and the sensor plates. The position of the bob can be deduced from the difference of these currents, the result of which is an induced EMF in the primary coil. The voltage induced in the primary coil is amplified by two low noise inverting op-amps, an OPA37 and an AD797, and the amplified voltage signal is read into a Stanford SR830 lock-in amplifier where it is demodulated into a DC signal. A schematic of the capacitive sensor circuit is shown in *Figure 6.7*.

The benefit of this capacitive sensor design is that the impedance to ground, through the primary transformer windings, is much less than that due to stray capacitance. If it is estimated that the stray capacitance for a 1 m length of wire is 10 pF then it can be shown that the impedance to ground should be,

$$\frac{1}{\omega C} = \frac{1}{2\pi \cdot 100 \text{ kHz} \cdot 10 \text{ pF}} \approx 160 \text{ k}\Omega. \quad (6.3)$$




Figure 6.7: A schematic of the capacitive sensor circuit.

However, the impedance to ground through the primary coil is estimated to be,

$$\omega L = 2\pi \cdot 100 \text{ kHz} \cdot 560 \mu\text{H} \approx 350 \Omega, \quad (6.4)$$

given a measured primary inductance of $560 \mu\text{H}$. Thus the majority of the current flows through the windings. The noise of the capacitive sensor, connected to the torsion balance but with no drive signal, at 1 mHz is $1.62 \times 10^{-4} \text{ V}/\sqrt{\text{Hz}}$ which translates into a torque noise of $6.02 \times 10^{-14} \text{ Nm}/\sqrt{\text{Hz}}$.

6.3.4 PID Servo Control System

A PID servo control is used to control the position of the torsion balance. In physical terms, proportional control shortens the natural period of the torsion balance, integral control ensures that the torsion balance maintains the set-point and derivative control damps out oscillations of the torsion balance.

The AC voltage signal from the capacitive sensor is demodulated by the phase sensitive detector and read onto a PC running a custom PID controller program, written in LabVIEW, using a PCI 6229 ADC. The program samples a predetermined number of data points from the voltage signal at a selected frequency and fits a straight line to those data points [149, 150]. The center point of the fit gives the most probable angular position of the torsion balance,

which is used in the proportional and integral part of the control program. The gradient of the fit gives the angular velocity of the torsion balance which is used in the derivative part of the control program.

In order to calculate the y intercept, a , and the gradient, b , of the sampled data, an ordinary least squares fit analysis [151] is used. If it is assumed that for each measured data point, i , the sensor voltage, V , can be related to time, t , using,

$$V_i = a + bt_i + \epsilon_i, \quad (6.5)$$

where ϵ_i is the residual of the i^{th} data point, the constants a and b can be calculated from minimising the sum of the squares of the residuals, S , where

$$S = \sum_{i=1}^n (V_i - (a + bt_i))^2. \quad (6.6)$$

Equation (6.6) is minimised when $\frac{\partial S}{\partial a} = 0$ and $\frac{\partial S}{\partial b} = 0$. By solving the two partial differential equations, it is possible to obtain the best fit values of the constants a and b , which will be denoted with a subscript “LS”. The expressions used for calculating a_{LS} and b_{LS} are given below.

$$a_{LS} = \frac{\sum V_i \sum t_i^2 - \sum V_i t_i \sum t_i}{n \sum t_i^2 - (\sum t_i)^2}, \quad (6.7)$$

$$b_{LS} = \frac{n \sum V_i t_i - \sum V_i \sum t_i}{n \sum t_i^2 - (\sum t_i)^2}. \quad (6.8)$$

After the least squares fit has been performed on the sampled data, the correct output voltages that should be applied to the actuator plates in order to hold the torsion balance at its set point are calculated. The servo control voltage, V_{servo} , is given by,

$$V_{\text{servo}} = \alpha (a_{LS} - V_{\text{set}}) + \beta \Sigma (a_{LS} - V_{\text{set}}) \Delta t + \gamma b_{LS}, \quad (6.9)$$

where α , β and γ are the gains of the P, I and D control respectively and V_{set} is a programmable set point. A bias voltage is applied to each actuator plate and

V_{servo} is added to the bias voltage of one actuator plate while it is subtracted from the bias voltage of the other actuator plate. This creates the appropriate potential difference needed to hold the torsion balance at its set point.

6.4 Calibrating the Capacitive Sensor

The torsion balance must be calibrated in order to accurately determine the torques acting on it. To do this, the bob was rotated through a known angle and the resulting change in output voltage from the capacitive sensor was measured. The torsion balance was moved in steps of $\sim 890 \mu\text{rad}$ and the change in the output voltage of the capacitive sensor was measured. This was repeated for both clockwise and anti-clockwise directions until the bob had been rotated through $\sim 1780 \mu\text{rad}$ in both directions. A plot of measured sensor voltage against rotation angle is shown in *Figure 6.8*. The gradient of the best fit line through the data gives a value of $96.1 \pm 7.0 \mu\text{V}/\mu\text{rad}$. Dividing the measured capacitive sensor data by this calibration coefficient allows the sensor voltage to be converted into an angle.

The largest source of error in the calibration process is likely from the stepper motor taking more, or less, steps than intended. If it is assumed that the error in counting the stepper motor steps is a simple Poisson error then the uncertainty can be estimated as the square root of the number of steps taken. This is likely to over estimate the uncertainty as during operations it was found that the stepper motor was generally quite accurate. However, there was not enough time to do a significant number of calibrations to get a more accurate estimate of what the uncertainty in the calibration actually was. Therefore, the Poisson error will be used in this instance. The standard deviation of the capacitive sensor data is only $140 \mu\text{V}$ which is negligible.

Once the sensor voltage is changed into an angle, it has to be converted from an angle into a torque. In order to do this, the data must be divided by




Figure 6.8: A plot of capacitive sensor output voltage against rotation angle. The gradient of the best fit line through the data points was used to calibrate the torsion balance in terms of an angle. The error on the gradient is 7.3 %.

the response function of the torsion balance [152],

$$r(f) = \frac{1}{\kappa \left(\left(1 - \frac{f^2}{f_0^2} \right) + \frac{i}{Q} \right)}, \quad (6.10)$$

where κ is the stiffness of the fibre, f is frequency, f_0 is the resonant frequency of the torsion balance and Q is the quality factor of the torsion balance.

The stiffness of the fibre can be determined from the shear modulus of the suspending fibre, G , the radius of the suspending fibre, r , and the length of the suspending fibre, L ,

$$\kappa = \frac{G\pi r^4}{2L}. \quad (6.11)$$

The stiffness of the tungsten fibre was calculated to be $1.6 \times 10^{-7} \text{ Nm/rad}$ when taking the value of G to be 156 GPa [147].

The quality factor of the torsion balance was determined to be 1449 from

the decay in the amplitude of the oscillations of the torsion balance after it had been excited. The amplitude of the oscillations, A , follow that of a decaying exponential of the form,

$$A = A_0 e^{-t/\tau}, \quad (6.12)$$

where A_0 is the initial amplitude, t is time and τ is the characteristic time of the ring down. The mechanical Q factor can be related to τ using,

$$Q = \pi f_0 \tau. \quad (6.13)$$

A typical ring down of the torsion balance is shown in *Figure 6.9*.




Figure 6.9: A plot of the decaying oscillations of the torsion balance after it had been excited. The Q factor of the torsion balance was determined from fitting an exponential decay, shown in red, to the data.

The resonant frequency, f_0 , was found through calculation of the natural period, T , of the torsion bob,

$$T = 2\pi \sqrt{\frac{I}{\kappa}}. \quad (6.14)$$

The inertia of the bob was calculated using the computer aided design software,

SolidWorks, in which the parts of the torsion balance were designed. When an assembly of an object is constructed in SolidWorks, the software is able to calculate the moment of inertia from the material properties assigned to parts in the assembly. The moment of inertia of the bob from SolidWorks was given to be $8.27 \times 10^{-5} \text{ kg m}^2$. This gives a natural period of 142.8 seconds which compares well to the observed natural period of 143 seconds. The observed natural period was determined from looking at the FFT of the torsion balance data.

It should be noted that this calibration was purely for the capacitive sensor, therefore, it could only be applied to measurements that were made with the torsion balance in free running mode or, in other words, with no servo control. A method of calibrating the torsion balance while under servo control will be discussed in Chapter 7.

6.5 Initial Results

6.5.1 Detecting Surface and Polarisation Charges on a Silica Disc

The torsion balance was set up such that a small copper rod was positioned near one of the silica discs on one arm of the torsion balance. The rod was orientated so that one of its flat circular ends was facing one of the flat circular surfaces of the silica disc. The copper rod had an AC voltage signal of amplitude 1 V and a frequency of 1 mHz applied to it. Most of the surface charge was mitigated from the silica disc using the static eliminator bar used in the experiments described in Chapter 5, however, some residual charge still remained on the disc which interacted with the copper rod.

Figure 6.10 shows an FFT of the time series data from the capacitive sensor. The peaks at 1 mHz and 2 mHz are the 1ω and 2ω components of the drive signal respectively, while the signal at 7 mHz is the resonant frequency of

the torsion bob. Two peaks arise because the force acting on the torsion balance consists of two contributions. The first is due to the force between the residual surface charge on the silica disc, Q_{surface} , and the charge on the copper rod, Q_{copper} , and the second is due to the force interaction between Q_{copper} and polarisation charges, Q_{pol} . Using Gauss's law it is possible to obtain an approximation for Q_{copper} that is related to the voltage applied to the rod, V ,

$$Q_{\text{copper}} = \frac{V\pi r^2 \varepsilon_0}{d} \cos \omega t, \quad (6.15)$$

where r is the radius of the copper rod, d is the distance between the surface of the silica disc and the rod, ω is the angular frequency of the 1 mHz voltage signal and t is time. Equation (6.15) is then substituted into Coulombs law to obtain an expression for the force exerted on the torsion balance,

$$F = \frac{(Q_{\text{surface}} + Q_{\text{pol}}) Q_{\text{copper}}}{4\pi \varepsilon_0 d^2}. \quad (6.16)$$

If it is assumed that $Q_{\text{pol}} = \alpha Q_{\text{copper}}$, where α represents a ratio related to the relative permittivity of silica and ε_0 , Equation (6.16) becomes,

$$F = \frac{Q_{\text{surface}} V r^2 \cos \omega t}{4d^3} + \frac{\alpha V^2 \pi r^4 \varepsilon_0 \cos^2 \omega t}{4d^4}. \quad (6.17)$$

The two terms in Equation (6.17) represent the 1ω and 2ω components that were observed in *Figure 6.10*. The first term is obviously the 1ω part of the signal as it contains a $\cos \omega t$. The $\cos^2 \omega t$ in the second term can be rearranged to give $0.5 + 0.5 \cos 2\omega t$ which is the 2ω part of the signal.

If it is assumed that Q_{pol} is small compared to Q_{surface} an estimate of the surface charge on the silica disc can be made. Using the value for F calculated from the output voltage of the capacitive sensor a value of $\sim 5 \times 10^{-14}$ C, which gives an estimated surface charge density of approximately 1×10^{-10} C/m², was obtained. This value agrees well with Kelvin probe measurements, made in Chapter 5, of residual charge on fused silica after being discharged with a

corona bar.




Figure 6.10: *Amplitude spectral density of the capacitive sensor output voltage showing a signal due to surface charge interaction at 1 mHz and a signal due to polarisation charge interaction at 2 mHz. The third peak at 7 mHz is the resonant frequency of the torsion bob.*

Unfortunately, the Kelvin probe was not incorporated into the experimental set up for this measurement, therefore, the actual surface charge density on the silica disc could not be confirmed. Since the copper rod was attached to the sample stage there wasn't enough room to fix the Kelvin probe to the stage. Even though the surface charge density could not be confirmed through Kelvin probe measurement, this was still a good preliminary test of the torsion balance.

6.6 Improving the Sensitivity of the Torsion Balance and Locating Noise Sources

To make an observation of charging noise with the torsion balance the sensitivity of the instrument had to be increased. After working with the set up

for several months, and reading through the appropriate literature, the main limiting noise sources of the instrument soon became clear. In this section the main noise sources of the instrument will be discussed as well as changes that were made to the set up in order to increase the sensitivity of the torsion balance. There will also be a description of the multiple regression least squares fit analysis that was used to remove some of the low frequency noise from the capacitive sensor data.

6.6.1 Decreasing the Fibre Diameter

The 50 μm tungsten fibre was replaced with a 40 μm tungsten fibre. As mentioned earlier in this chapter, decreasing the diameter of the fibre can increase the sensitivity of the instrument significantly. With the fibre diameter decreased, the stiffness of the fibre also decreases while the natural period of the torsion balance increases. The new calculated stiffness of the fibre is $6.5 \times 10^{-8} \text{ Nm/rad}$ and the new natural period of the torsion balance is 178 seconds, or 5.62 mHz. Suspending the bob by a fibre 25 μm in diameter was also attempted, however, the bob was notably more difficult to control, and to set up, due to the longer period of the bob.

6.6.2 Changing the Sensor/Actuator Configuration

The sensor/actuator set up was changed so that on each of the columns, on which the sensor and actuator plates are mounted, there is a sensor and an actuator plate instead of two sensors on one column and two actuators on the other. In this configuration the capacitive sensor rejects simple pendulum motion (common mode) of the bob but remains sensitive to differential rotation of the bob. A diagram of this set up is shown in *Figure 6.11*.

6.6.3 Shielding the Bob and Sensor Plates

It was thought that the exposed sensor plates could be picking up stray capacitance from exposed wires or even from surface charge on the insulating




Figure 6.11: A diagram of the new sensor and actuator configuration of the torsion balance set up. This set up rejects the first order pendulum mode of the torsion balance in the x and y directions but is still sensitive to the rotational motion of the bob. This is because the transformer is 1:1:-1 and the pendulum motion is common mode, whereas the torsional motion is differential.

coating on the coaxial wires inside the vacuum tank. A shield was fabricated to cover the top half of the columns on which the sensor and actuator plates were mounted, and almost the entirety of the bob. The shield was made out of a thin sheet of copper which was $20\text{ }\mu\text{m}$ thick. The copper was easily bent, cut and soldered into the necessary shape to create the shield. Screening the sensor plates and the torsion bob from dielectric surfaces increased the sensitivity of the torsion balance significantly. A plot of the torque noise of the torsion balance before and after the shield was put in place is shown in *Figure 6.12*.

6.6.4 Contact Potentials

As discussed in Chapter 2, contact potential differences can arise between two metals with different work functions when they are brought into contact. The torsion balance is composed of different metals which could create various contact potentials. These contact potentials can change with temperature which would give rise to a fluctuating voltage acting on the torsion balance




Figure 6.12: A plot showing the increase in sensitivity after the bob and sensor plates were shielded.

[152]. Contact potentials can also cause problems when trying to maintain servo control of the torsion balance. Sometimes when the bob was under servo control it would appear as if the servo actuation was having little or no effect. This is due to the fact that in order for the servo control to have any effect, it must first cancel out the contact potential. To get around this problem the servo voltages had to be kept above the contact potential difference between the torsion balance and the actuator plates. This was achieved by having the bias voltage of the actuator plates, discussed in Section 6.3.4, set appropriately high.

The contact potential between the bob and the servo plates was measured by rotating the bob through a known number of steps, while the bob was under servo control, and monitoring the change in the servo voltages. The torque being applied to the bob by each actuator plate, Γ , can be calculated using,

$$\Gamma = \frac{1}{2} \frac{dC}{d\theta} (V - V_{cp})^2 + \Gamma_0, \quad (6.18)$$

where $\frac{dC}{d\theta}$ is the capacitance gradient between the servo plate and the bob, V is the applied DC bias, V_{cp} is the contact potential and Γ_0 is the initial offset torque. When Γ is plotted against V the data follow a parabola where the voltage at the minimum of the parabola, or maximum depending on the sign of the torque, is equivalent to the contact potential. The parabola follows the form,

$$\Gamma = \frac{1}{2} \frac{dC}{d\theta} V^2 - \frac{dC}{d\theta} V V_{cp} + \left(\frac{1}{2} \frac{dC}{d\theta} V_{cp}^2 + \Gamma_0 \right). \quad (6.19)$$

Plots of the torque, calculated from the rotation of the bob and the stiffness of the torsion fibre, against applied bias voltage for each of the actuator plates are shown in *Figure 6.13* for each of the servo plates. The values of the contact potentials measured were 1.56V and 1.40V and the average capacitive gradient value was calculated to be 1.2×10^{-11} F/rad. As stated previously, the error in the rotation of the bob is estimated to be the square root of the number of stepper motor steps. The error on the number of stepper motor steps taken is $\sqrt{8}$. If $\sqrt{8}$ steps is converted into an uncertainty in the angular displacement of the bob we get $\pm 3.7 \times 10^{-4}$ rad which gives a torque uncertainty of $\pm 2.4 \times 10^{-12}$ Nm.

No precautions were taken in order to reduce any excess noise caused by contact potentials. There was some discussion of getting parts of the bob, the servo plates and the actuator plates coated in gold, using a thermal evaporation process, to try and reduce this noise. Gold has a very uniform work function and therefore should minimise any adverse effects of contact potentials. Unfortunately the thermal evaporation apparatus available was not suited to coating large pieces, such as the bob. Most of the bob is made of aluminium which is also very difficult to coat due to the oxide layer which is present on its surface. If parts of the bob were to be coated in gold, they would, most likely, have to be sent to an external company which would be costly and time consuming.




Figure 6.13: Plots showing the torque applied to the bob by the actuator plates against servo voltage. The data points follow a parabola and the contact potential differences between the bob and the actuator plates can be deduced from the parabolic fit.

6.6.5 Natural Unwinding of the Fibre

When the torsion balance is in free running mode, the signal measured from the capacitive sensor tends to drift over time. Part of this drift is due to the fibre naturally unwinding under the load of the bob. This unwinding is a result of the method used to wind the tungsten wire onto a spool during the manufacturing process. If the bob is suspended by a new fibre, it has to be left for a few days until the fibre settles. This drift is also noticeable for at least a day after pumping out the vacuum tank, even with a fibre that has been used for several weeks. The fibre will also drift due to its anelasticity. This is when the fibre slowly relaxes back to its original state [153]. There are methods of annealing the fibre in the literature [154] that claim to remove, or reduce, the drift from the fibre, however, after annealing one fibre it was found that there was not a significant reduction in the drift.

6.6.6 Temperature Drift

The main culprit of low frequency drift in the torsion balance data is due to changes in temperature. As the temperature in and around the tank changes, the tungsten fibre and metal structure expands or contracts causing the bob to change position. Three PT-100 platinum resistance thermometers were used to measure the temperature inside and outside the vacuum tank. One was located outside the tank near the capacitive sensor circuitry and the other two were located inside the tank; one at the top of the inner structure and one near the bottom of the inner structure. The room in which the torsion balance apparatus was located did not have any temperature control. The temperature variation typically observed is shown in *Figure 6.14*.

From *Figure 6.14* it can be seen that the temperature variation outside the vacuum tank is greater than the temperature variation inside the tank. This is likely due to gusts of air entering the laboratory from the lift shaft that is immediately outside. The vacuum tank filters the temperature variations and makes them appear a lot smoother inside the tank.

Using two power resistors, one located at the top of the set up and another near the bottom, it was possible to change the temperature inside the tank. Using the power resistor attached to the top of the set up, the temperature was increased by a small amount and then left to decay back to equilibrium. A linear correlation was found between changes in the output voltage of the capacitive sensor and the changing temperature in the vacuum tank. This correlation is shown in *Figure 6.15*.

Torsion balance facilities that achieve very high sensitivities often have their instrument contained within a room where the temperature is monitored and maintained at a constant temperature [155]. This helps to reduce thermally induced drift in the fibre. Unfortunately, there was not enough time, or funds, to develop such a sophisticated set up for this experiment. In order to reduce the low frequency noise created by changing temperature, the thermal fibre




Figure 6.14: *Three plots showing the typical temperature variations measured by the three thermometers used in the experimental set up. The top plot shows the temperature outside the vacuum tank and the other two show the temperature inside the vacuum tank.*

drift was fitted out of the data using a multiple regression least squares fit program. The fit program will be discussed in more detail in a later section.

6.6.7 Tilt

Ground tilt is another major noise source that plagues some torsion balance experiments [156]. If the main clamping block at the top of the set up tilts




Figure 6.15: A plot showing the linear correlation between changes in temperature and changes in the capacitive sensor voltage.

relative to the ideal position of the torsion bob, this motion can couple into the torsional motion of the bob [143]. If the apparatus tilts it will also change the position of the torsion balance, which always hangs along the vertical axis of the set up, relative to the sensor and actuator plates. This can lead to pseudo torques being detected in the torsion balance data.

An Applied Geomechanics Inc. 755 miniature tilt sensor was used to monitor ground tilt. Figure 6.16 shows the tilt of the vacuum tank in the x and y directions over approximately 42 hours. It was found that tilt variations of the vacuum tank were less than a few μrad , therefore, it is unlikely that the torsion balance is limited by noise due to tilt. The tilt data is used to fit out any low frequency noise created by tilt, however, this will likely be a very small effect.




Figure 6.16: A plot of the x and y tilt of the vacuum tank.

6.6.8 Magnetic Field Fluctuations

An AD22151 magnetic field sensor was used to monitor changes in the background magnetic field that might occur in the laboratory. When the torsion balance was first set up there were some problems with magnetic fields and magnetic materials. The bob was initially bolted together using stainless steel bolts, and it was found that by moving a strong neodymium iron rare earth magnet near the bob it would rotate significantly. The stainless steel bolts were then replaced with titanium ones so that now the bob does not obviously appear to be influenced by magnetic fields. *Figure 6.17* shows the background magnetic field over the course of a few days in the lab where the torsion balance is located. There is no obvious correlation between the torque noise and the magnetic field sensor output and therefore shouldn't be a limiting noise source.




Figure 6.17: A plot of the magnetic field recorded over a weekend in the room which contains the torsion balance. The average field is ~ 0.29 G

6.7 Using Multiple Regression Fitting to Remove Known Noise Sources

A multiple regression fitting program, that uses the least squares criterion, was used to fit out some of the noise sources, discussed in previous sections, from the torsion balance data. Fibre drift, thermal effects and tilt were fitted out of the raw capacitive sensor data. It has been shown in Section 6.6.6 that small temperature changes will cause a linear change in the bob position. Therefore it was assumed that the tilt and fibre drift effects also had a linear correlation with changes in the bob position if these effects are small.

A multiple regression fitting program taken from “Numerical Methods Using Matlab” by Lindfield and Penny [157] was used for this analysis. This method is much like the linear least squares fit analysis described in Section 6.3.4 because it uses least squares criterion. This means that the best fit parameters will be determined from minimising the sum of the square of

the residuals of the equation below,

$$V_i = a_0 + a_1 t_i + a_2 T_{\text{IN}}(t_i) + a_3 T_{\text{OUT}}(t_i) + a_4 \text{tiltx}(t_i) + a_5 \text{tilty}(t_i), \quad (6.20)$$

where V_i is the measured capacitive sensor voltage of the i^{th} data point, t_i is the time at which the i^{th} data point was taken, a_0 is the DC offset of the capacitive sensor measurements from zero, T_{IN} is the normalised average temperature inside the tank, T_{OUT} is the normalised temperature outside the tank, tiltx is the normalised tilt in the x direction and tilty is the normalised tilt in the y direction. The parameters a_1 to a_5 are unknown coefficients.

The fit program finds an estimate, b_k , of the unknown coefficients in Equation (6.20) using initial parameters determined by the user. This gives the equation of the fitted function as,

$$\hat{V}_i = b_0 + b_1 t_i + b_2 T_{\text{IN}}(t_i) + b_3 T_{\text{OUT}}(t_i) + b_4 \text{tiltx}(t_i) + b_5 \text{tilty}(t_i), \quad (6.21)$$

where \hat{V}_i is the fit value of the i^{th} data point. Once \hat{V}_i has been calculated for all i data points the error, e , of the fit can be determined from,

$$e_i = V_i - \hat{V}_i. \quad (6.22)$$

As before, this error is used to calculate the sum of the square of the errors, S , which, once minimised, will give the best fit values of the unknown coefficients in Equation (6.20).

$$S = \sum_{i=1}^n e_i^2 = \sum_{i=1}^n (V_i - \hat{V}_i)^2. \quad (6.23)$$

A full, and detailed, description of this method can be found in [157].

Once the best fit form of \hat{V}_i has been found, it is subtracted from the raw capacitive sensor data. This should remove any drift or noise associated with the noise sources that are taken into account in the fitting process. The top plot in *Figure 6.18* shows the raw time series capacitive sensor data, in blue,

before noise has been removed and the calculated least squares fit based on the function given in Equation (6.20), shown in red. The second plot shows the time series data of the capacitive sensor after the best fit has been subtracted.




Figure 6.18: The top plot shows the raw capacitive sensor data, in blue, before the least squares fit data, in red, is subtracted from it. The bottom plot shows the capacitive sensor data after the fit has been subtracted, leaving the residual noise.

6.8 Sensitivity of the Torsion Balance After Improving the Set up

In order to increase the sensitivity of the instrument further, the sensor and actuator plates were placed a few millimeters closer to the torsion bob². Once all of the mentioned improvements to the set up had been made, and the multiple regression analysis was able to successfully remove known excess low frequency noise, the torsion balance was calibrated again. The new calibration constant was calculated to be $243 \mu\text{V}/\mu\text{rad}$ using the same method described in Section 6.4. This shows a definite improvement from the calibration discussed in Section 6.4. A measurement of the best sensitivity of the torsion balance was made in late August 2011. The noise level for this measurement is shown in *Figure 6.19*. At this noise level the torsion balance is sensitive to torques of $4.43 \times 10^{-15} \text{ Nm}$ at 1 mHz.

This measurement was taken while the torsion balance was in free running mode and without a charged sample in the vacuum chamber. The sensitivity of the instrument would be expected to decrease slightly when taking a measurement of a charged silica disc. This is because the shield around one of the bob arms would have to be removed in order to make the measurement.

The thermal noise limit shown in *Figure 6.19* was calculated using [158, 159],

$$S_{\tau}(f) = \sqrt{\frac{4k_B T \kappa}{2\pi Q f}}. \quad (6.24)$$

where k_B is the Boltzmann constant, T is temperature, κ is the stiffness of the suspending fibre, Q is the quality factor of the torsion balance and f is the frequency. The instrument is within a factor of 10 of the thermal noise at 0.1 mHz which is expected due to the lack of temperature and tilt control. This was a sufficient sensitivity to begin measurements of charging noise.

²The sensor and actuator plates were initially placed as far as possible from the bob so that the torsion balance was easier to control.




Figure 6.19: A plot showing the highest sensitivity reached by the torsion balance. The torsion balance noise is shown in blue, the capacitive sensor noise is shown in black and the thermal noise limit is shown in red.

6.9 Torsion Balance - Mark II

As with all experimental work, it was found that the design of the torsion balance could be improved after the limitations of the instrument were determined. A Mark II version of the instrument was designed that should reach higher sensitivities than the previous version of the instrument. Two designs were considered for the Mark II instrument, one which consisted of a thin plate for the bob, much like the University of Washington Eöt-Wash instrument [160], and one which consisted of a four armed cross with cube test masses on the ends of the arms, much like the instrument used at the University of Trento [161]. The latter design was chosen as the design of the Mark II instrument.

6.9.1 The Torsion Bob

The bob of the Mark II instrument, shown in *Figure 6.20*, has four isolated cube masses attached to the ends of each of its four arms instead of plates. There is also a small copper plate bolted to the center of the bob with a $300\mu\text{m}$ hole drilled in the center of it. The center hole in the copper plate is for a copper tube, with an outer diameter of $250\mu\text{m}$ and inner diameter of $50\mu\text{m}$, to be inserted. The torsion fibre will then be threaded into the tube and the tube will be crimped to hold the fibre in place. This will allow the torsion fibre to be better centered during the set up process rather than using a fibre clamp. The fibre clamp of the Mark I instrument had quite a large margin for error and often the fibre would have to be re-clamped several times before it was reasonably centered. The reason the copper plate was made separate to the rest of the bob is so that if the fibre breaks and needs to be replaced, the plate is easily removed.




Figure 6.20: A photograph of the Mark II torsion bob. The cube masses are isolated from the rest of the bob so that the sensing signal can be applied to the cubes capacitively while keeping the rest of the bob grounded.

In the Mark I version of the instrument the sensing signal was applied at the top fibre clamp and down the fibre to the bob. When the signal is applied down a thin exposed fibre it is sensitive to stray capacitance. This method of applying the sensing signal also leaves the whole bob floating at approximately 1 V. Ideally, it would be preferred if the bob was grounded so it is at the same potential as its surroundings. The Mark II instrument will apply the sensing signal capacitively onto two of the cube masses using an electrode in the top plate of the sensor housing, discussed in Section 6.9.2. Since the cubes are electrically isolated from the rest of the bob, the bob and fibre can remain grounded while the cubes are left floating. An optical position sensing method may also be developed for the Mark II instrument. This would also leave the bob grounded and avoid the use of electrostatics during measurements.

Another benefit of not connecting the sensor signal at the main fibre clamp is that the bob can now be continuously rotated. With the sensor signal wire connected to the fibre clamp, the bob can only be rotated so far until the wire starts to wrap around the eddy current damper. This would be very useful when making measurements of charging noise as the DC force from the surface charge is usually very large. The bob could be rotated through a very large angle until the majority of the DC force is canceled out.

6.9.2 The Sensor Housing

One of the major problems with the Mark I instrument was that the sensor plates were left exposed to dielectric surfaces inside the vacuum tank. This created a significant amount of noise which was reduced by shielding the sensor plates and the bob. This is why, for the Mark II instrument, it was decided that the sensor and actuator plates would be housed in an enclosure. The housing is made of aluminium and has three copper electrodes which can be fixed into the top plate and two of the side plates of the housing. The electrode in the top plate is for applying the drive signal to the cube mass, one of the electrodes

in one of the side plates is for applying servo control and the other electrode in the opposite side plate is for capacitive sensing. The fully constructed sensor housing is shown in *Figure 6.21* and a view of the sensor housing with the top plate raised is shown in *Figure 6.22*. There is also an “exploded” view of the top half of the sensor housing shown in *Figure 6.23*.




Figure 6.21: A photograph of the sensor housing for the Mark II torsion balance.




Figure 6.22: The electrode in the roof of the housing is for applying the drive signal to the cube mass of the torsion balance.

6.10 Conclusion

A torsion balance apparatus has been developed that is of sufficient sensitivity to begin studies of charging noise. A direct measurement of charging noise has not been made yet and thus a measurement will allow us to determine whether the Weiss theory of charging noise is the most appropriate description.

Initial measurements have shown that it is possible to measure both the polarisation charge and surface charge present on a silica disc. With the in-




Figure 6.23: A photograph showing an exploded view of the inside of the Mark II sensor housing. The sensor, actuator and drive electrodes are highlighted in the photograph.

clusion of a Kelvin probe in the experimental set up it should be possible to make accurate measurements of the surface charge on samples of silica and the time constant of the charging noise. This will allow accurate estimates of the signal expected from charging noise.

The limiting noise sources of the instrument have been investigated and the sensitivity of the instrument has been greatly improved to the stage where it should be possible to make a successful measurement of charging noise. This measurement will provide greater understanding of charging noise and allow accurate predictions of charging noise in advanced gravitational wave detectors.

A Mark II torsion balance instrument has also been designed and built. The design of the Mark II torsion balance addressed some of the problems with the design of the Mark I instrument. The improved design should make the instrument easier to work with and should improve the sensitivity. Due to lack of time, the performance of the Mark II instrument has not been included in this thesis.

Chapter 7

Measuring Charging Noise with a Torsion Balance

7.1 Introduction

With the torsion balance at a sufficient sensitivity, measurements of charging noise could now be attempted. This chapter presents a measurement of the torsion balance sensitivity with a charged fused silica disc near the bob. There will be a detailed discussion on how the torsion balance was calibrated for this measurement as it is very different from the method described in the previous chapter. This is because for a measurement of charging noise the bob had to be under servo control, therefore, the closed loop gain of the servo control had to be taken into account. There will also be a detailed explanation of how the torque noise due to moving surface charges was estimated. This involved estimating the average torque acting on the torsion balance and determining the correlation time of the surface charge from a Kelvin probe measurement. The chapter will conclude with a discussion of the results and any disparity between the theoretical predications and the measurement.

7.2 Experimental Set Up

The experimental set up for the measurements presented in this chapter is virtually the same as described in the previous chapter. However, there are a few minor changes.

The initial plan for this experiment was to charge a silica disc fixed to the sample stage and a silica disc attached to the bob. The fixed sample would then be moved into position using the stage motors so that the two charged surfaces would be facing each other and a measurement could be made. It was found that the bob was far too difficult to control using this method, therefore, the silica disc on the bob was substituted for an aluminium plate. This way it is only the silica disc fixed to the sample stage that has to be charged. To keep the bob balanced, the copper plates attached to the bob that were used for capacitive sensing and servo control were also substituted for the aluminium plates. The aluminium plates are slightly larger than the original copper plates so as to increase the area which interacts with the surface charge on the silica sample. A photograph of the bob with the aluminium plates attached is shown in *Figure 7.1*.

To charge the sample, a large rectangular rubber pad was made from a Viton o-ring and then covered with a thin sheet of copper. This was attached to the fixed platform, near the bottom of the vacuum chamber, using an aluminium post and was positioned near the bob. The copper was electrically isolated from the grounded fixed platform and connected to an electrical feedthrough so that voltages could be applied to it. This allowed the pad to be grounded during measurements, preventing unwanted excess noise. The sample could be easily charged in vacuum by moving it away from the bob and pressing it against the pad using the motorised stages. The copper had a voltage of 10 V applied to it when the silica disc was pressed against it. The copper pad can be seen in both *Figure 7.1* and *Figure 7.2*.

In order to make measurements of charging noise, the bob of the torsion




Figure 7.1: A photograph of the torsion bob with new aluminium plates attached to the ends of each of its four arms. The Kelvin probe, silica sample and copper pad, used for charging the sample, can also be seen in the photograph.

balance would have to be under servo control due to the large DC torque exerted on it by the surface charge on the silica sample. The servo voltages were originally applied to the actuator plates using a National Instruments PCI 6229 ADC. This could only supply a maximum of 10 V to the actuator plates which would not be enough to keep the bob under servo control while a charged sample was present in the vacuum chamber. A Trek 603 dual channel power amplifier, that could supply 50 V, was briefly used but it was found to be too noisy for this experiment. One channel of the Trek power amplifier was also improperly calibrated and always had an offset of a few volts. Eventually two high voltage OPA445 op-amps were used to supply the servo plates with 40 V. This provided adequate servo control for the purposes of this experiment. Both of the OPA445 op-amps were set up in a standard non-inverting configuration, as shown in *Figure 7.3*.




Figure 7.2: A photograph of the silica sample positioned near the torsion bob to make a measurement of charging noise. The decaying surface charge on the silica disc, which is monitored with the Kelvin probe, exerts a decaying torque on the torsion balance plate. The disc was charged by pressing it against the copper pad whilst the pad had a voltage applied to it.




Figure 7.3: A schematic of the non-inverting op-amp circuit used to provide a servo control voltage to the actuator plates.

A Kelvin probe, which can be seen in *Figures 7.1* and *7.2*, was added to the torsion balance set up so that the charge on the fused silica disc sample could be monitored. Since the charged face of the silica disc has to be facing

the bob, it was not possible to directly measure this surface charge density. Instead, the Kelvin probe was positioned at the opposite face. This meant that it was probably a polarisation charge induced in the opposite face that was being measured. Since it was a polarisation charge that was being measured, the calibration for the Kelvin probe would be different from that described in Chapter 3. Unfortunately an alternative calibration method could not be developed in time. However, a measurement of the surface charge density on the silica disc is not essential for the charging noise measurement as the average torque acting on the bob, needed for estimating the level of charging noise expected, can be calculated from the rotation of the bob and the stiffness of the torsion fibre. The Kelvin probe was still used to measure the correlation time of the charge.

It should also be noted that the method of computationally removing known noise sources from the torsion balance data was not used for the results presented in this chapter. Since the decay of the surface charge has a very long time constant, it can appear as if there is a large linear drift in the data. The multiple regression fit program would fit to the charging noise signal and remove it. Therefore, the data presented in this chapter is the raw, unedited, torsion balance data that was recorded during the measurements. The drift in the signal due to changes in temperature and tilt are also insignificant compared to the strong signal of the charging noise and should not greatly affect the data anyway.

7.3 Calibrating the Servo Control

The measurements made in Chapter 6 were taken when the torsion bob was in free running mode. This means that the servo control was only used to damp down the oscillations of the torsion bob and then the servo control was turned off for the duration of the measurements while the bob was allowed to naturally oscillate. To make a measurement of charging noise servo control

must be used. This is because the servo control will have to cancel out the large DC torque exerted on the bob by the surface charge on the silica sample.

Calibrating the torsion balance while under servo control is slightly more complex than calibrating it in free running mode. *Figure 7.4* shows a schematic of the gains of the torsion balance set up that will have to be considered. The gain of the PID controller, G , and the capacitance gradients, C , between the bob and each of the servo plates must be taken into account, as well as the response function of the bob, T . Transfer functions T and G can be defined in the Laplace domain as

$$T = \frac{1}{Is^2 + bs + \kappa} \quad (7.1)$$

and

$$G = \alpha + \frac{\beta}{s} + \gamma s \quad (7.2)$$

where I is the moment of inertia of the bob, b is the damping constant, κ is the stiffness of the torsion fibre and α , β and γ are the gains of the P, I and D controls respectively. C is a constant and was measured to be 1.2×10^{-11} F/rad, as stated in Chapter 6. By applying a known torque, Γ_{IN} , on the bob while it is under servo control, it should be possible to calculate the gain required to generate the feedback torque, Γ_F , needed from the servo voltage, V_M , to keep the torsion balance at its set point. The gain which will be calculated is the closed loop gain of the system.

An expression must be found that relates Γ_{IN} to the measured servo voltage, V_M , so that the servo voltages can be converted into a torque. Using control system theory, and with the aid of *Figure 7.4*, it is possible to derive the relation between Γ_{IN} and V_M . First, it is possible to relate Γ_F to the capacitive sensor output, V_{OUT} , using,

$$\Gamma_F = CGV_{OUT}. \quad (7.3)$$




Figure 7.4: A schematic of the gains that have to be taken into account in the torsion balance set up. For a torque acting on the bob, Γ_{IN} , the torsion balance will rotate and V_{OUT} will be measured by the capacitive sensor. V_{OUT} is also used to calculate the servo voltage, V_M . V_M applies a negative feedback torque, Γ_F , to the bob to counteract the bob's motion and keep it at its set point.

It is also possible to relate V_{OUT} to Γ_{IN} and Γ_F using,

$$V_{OUT} = T (\Gamma_{IN} - \Gamma_F). \quad (7.4)$$

The expression above essentially states that if the correct Γ_F is applied to keep the bob at its set point then no change in V_{OUT} should be observed. Rearranging Equation (7.3) for V_{OUT} and substituting into Equation (7.4) gives,

$$\frac{\Gamma_F}{GC} = T (\Gamma_{IN} - \Gamma_F). \quad (7.5)$$

Equation (7.5) can be rearranged for Γ_F to obtain the expression,

$$\Gamma_F = \frac{GTC \Gamma_{IN}}{(1 + GTC)}. \quad (7.6)$$

A new expression for Γ_F can be substituted into Equation (7.6) to relate Γ_{IN} to V_M . From Figure 7.4 it can be seen that Γ_F can be expressed as,

$$\Gamma_F = CV_M, \quad (7.7)$$

therefore, Equation (7.6) becomes,

$$\Gamma_{IN} = \frac{V_M (1 + GTC)}{GT}, \quad (7.8)$$

where the constant term in the equation, $(1 + GTC)/GT$, is the closed loop gain of the system.

With Equation (7.8) derived, it is now possible to calculate the gain of the control loop at different frequencies. To do this, the stepper motor that controls the rotation stage, shown in *Figure 6.3*, was programmed to move in a sine wave at selected frequencies with an amplitude of 100 steps, or $1310\mu\text{rad}$. When the torsion balance oscillates in this manner the servo voltages will also oscillate as they try and apply enough feedback torque to keep the bob at its set point. Γ_{IN} can be calculated from the amplitude of the rotation, and the stiffness of the fibre, and the amplitude of V_M can be determined from fitting a sine wave to the servo voltage data using the multiple regression fitting analysis described in Chapter 6. A plot of the stepper motor motion, the servo effort and the fit to the servo data is shown in *Figure 7.5*.

The gain in Equation (7.8) can be calculated simply by dividing the amplitude of Γ_{IN} by the amplitude of V_M . The gain measurement was carried out for a range of frequencies and the results are shown in *Figure 7.6*. It can be seen from *Figure 7.6* that the measured gain starts to level off at low frequencies. This is because at low frequencies the movement of the rotation stage essentially seems like a very slow drift from the set point. The integrator of the PID controller dominates at this point in order to keep the bob at its set point. This is also the reason why the servo effort and the stepper motor motion are out of phase with each other in *Figure 7.5*.

The error in *Figure 7.6* comes from the uncertainty in the angular displacement in the estimate of the torque being applied to the torsion balance. As discussed in Chapter 6, this error is taken as the square root of the number of stepper motor steps. For this case the uncertainty in angular displacement will




Figure 7.5: A plot showing the change in servo effort, blue, as the torsion balance is moved in a sine wave motion, green. The fit to the servo data is also shown, in red. The servo action is out of phase with the stepper motor motion because the integral action is dominant at low frequencies.

be $\pm 131 \mu\text{rad}$ which gives an uncertainty in the torque of $\pm 8.52 \times 10^{-12} \text{ Nm}$. As stated previously, this is likely to be an overestimate of the uncertainty but as there was not sufficient time to carry out repeated measurements of the calibration this is the best estimate of the uncertainty.

So that the gain could be calculated for any frequency, a best fit line was plotted through the data points in *Figure 7.6*. It was found that the function that fitted the data best [162] was,

$$\text{Gain} = \frac{1}{A_0 - Bf \exp\left(\frac{f}{f_0}\right)}, \quad (7.9)$$

where f is frequency and A_0 , B and f_0 are constants. The best fit of this function was found using fit tools in Matlab and is represented by a red dashed line in *Figure 7.6*. The measured servo voltage could then be converted into a torque by multiplying it by the appropriate gain value calculated using the




Figure 7.6: A plot of the closed loop gain of the torsion balance PID system measured at different frequencies. The expression given in Equation (7.9) was used to fit to the data points so that the gain could be calculated for any frequency.

best fit line to the gain data.

This calibration was initially carried out with no charged sample in the vacuum chamber. This was to compare the torque noise of the torsion balance under servo control with the torque noise of the torsion balance in free running mode, as both of these noise levels should be approximately the same. It should be noted that even though the torsion balance was in free running mode there were fixed servo voltages applied to the servo plates for this measurement. The reason for this was incase the presence of voltage on the servo plates created additional noise through coupling via electrostatic stiffness. For an accurate comparison, the noise would have to be present in both noise measurements. *Figure 7.7* shows plots of the torque sensitivity of the torsion balance measured while it was in free running mode, red, and while it was under servo control, black. It can be seen from *Figure 7.7* that the torque sensitivities of the torsion balance in these two modes of operation compare well. Thus confirming that

the calibration method is accurate and reliable. The increase in sensitivity around 5/6 mHz in the fixed servo data is due to the resonant frequency of the torsion balance in free running mode.




Figure 7.7: A plot of the torque sensitivity of the torsion balance while in free running mode, red, and while under servo control, black.

The gain measurements for the PID set up during the charging noise measurement were slightly trickier. The closed loop gain was measured in the same way as described previously except that the PID values were changed to those used during the charging noise measurement. Once the charging noise measurement had been made, the sample was moved away from the bob and some air was released into the chamber to remove the excess charge from the silica disc. In order to keep the bob in place during the charging noise measurement the proportional gain had to be very high, which can cause the PID control to become unstable. Unfortunately without the surface charge the high proportional gain made the torsion balance uncontrollable. This meant that the control loop gain could not be measured with exactly the same proportional gain as used during the charging noise measurement. It was found that

the highest proportional gain that could be used, without the charged sample present, to keep the torsion balance reliably controlled was half the value used for the charging noise measurement. The amplitude of the servo voltage oscillation was measured at 0.5 mHz for different proportional gain values to see if the amplitude of the servo voltage oscillation for the proportional gain used for the charging noise measurement could be extrapolated from the data. These measurements are shown in *Figure 7.8*. From *Figure 7.8* it appears as if changing the proportional gain does not have a significant effect on the overall control loop gain, therefore, it will be assumed that the closed loop gain for the proportional gain used during the charging noise measurement will be approximately the same as for a proportional gain of half that value. The error in the estimation of the servo amplitude is negligible as the servo voltages can be measured to high accuracy. It was not known how reproducible the calibration process was as there was not enough time to gather sufficient statistics. Therefore, it is unknown whether the outlying data point, for the proportional gain value of 0.25, is truly an outlying data point or within in the standard deviation that one might expect from repeating this measurement.

A plot of the stepper motor motion, the servo effort and the fit to the servo data is shown in *Figure 7.9* and the control loop gain as a function of frequency for the new PID values is shown *Figure 7.10*. For these measurements the servo control appears more in phase than in *Figure 7.5*. This is likely because, with such a high proportional gain, the proportional term is the most dominant in this circumstance.

The gain measurements were taken between 0.25 and 2 mHz. Measurements were not taken below 0.25 mHz because it would have taken too long and measurements were not taken above 2 mHz because there would not have been enough time to remove backlash in the stepper motor. If the motor direction needed to be changed the backlash in the motor had to be removed. The number of steps needed to remove backlash in the motor used here was




Figure 7.8: A plot of the amplitude of the servo voltage oscillation measured for different values of proportional gain. It does not appear as if changing the proportional gain has a significant effect on the overall control loop gain. The proportional gain used for the charging noise measurement was 0.7.

75. During the sine wave motion of the rotation stage the motor had to stop at the peaks and troughs of the wave and remove the backlash. This is why the peaks and troughs of the servo voltage shown in *Figure 7.9* appear slightly flattened. It is unknown whether the gain measurements at 1.75 mHz and 2 mHz are dropping off as the frequency increases or if they fall within the expected standard deviation of the measurement. Therefore, the data that will be analysed will be below 1.25 mHz because the data definitely appear to level off at these frequencies. The average gain of these data points was found to be 1.98×10^{-9} Nm/V. This value will be used to convert the servo voltages recorded during the charging noise measurement into a torque. The error in the estimation of the servo amplitude will be taken as 8%, as this was the standard deviation of the data points in *Figure 7.8*, and it was assumed that they were approximately the same value during the charging noise measure-




Figure 7.9: A plot showing the change in servo voltage, blue, the fit to the servo voltage data, red, and the torsion balance sine wave motion, green, for the PID values used during the charging noise measurement.




Figure 7.10: A plot of the closed loop gain measured at different frequencies for the PID values used during the charging noise measurement.

ment. As mentioned previously, there was not sufficient time to obtain good statistics on the standard deviation of repeated gain measurements. There is also an error in the estimate of the torque exerted on the torsion bob from the sinusoidal rotation of the rotation stage as mentioned previously. This value was estimated to be 10 % and will be added in quadrature to the 8 % error in the servo voltage amplitude to give a combined error of 12.8 % on the torque noise.

7.4 Results

A silica disc was charged under vacuum by pushing it against the copper pad, shown in *Figure 7.2*, held at a potential of 10 V. The disc was pushed against the pad multiple times until the disc had acquired a sufficient amount of charge. The sample was then moved into position using the motorised stages so that the charged surface of the silica sample was facing one of the plates at the ends of the arms of the torsion bob. The silica disc was spaced approximately 18 mm from the plate of the torsion balance. Initially, the bob could not be servo controlled due to the large DC torque exerted on the bob by the charge. The bob had to be rotated through an angle of 0.55 rad, or 31.5 °, so that the restoring torque of the fibre would help cancel out the DC torque from the surface charge. Once the bob had been rotated, servo control was achieved and the torsion balance was left for several days to take a measurement. The servo voltages that were measured are shown in *Figure 7.11*. Typical servo voltages measured without a charged sample in the vacuum chamber are shown in *Figure 7.12*.

When the charged sample is in the vacuum chamber the change in the servo voltages is far greater than when the sample is removed from the set up. This is due to the fact that the surface charge is decaying and therefore the DC torque being exerted on the bob is decaying. The decaying DC torque of the surface charge is significantly larger than any of the other noise sources of




Figure 7.11: *Plots of the applied servo voltages during the charging noise measurement.*

the instrument that were discussed in the previous chapter. This is why the change in servo voltage without the charged sample in the chamber is of the order of tens of millivolts and the servo voltage in the presence of the charged sample changes by several volts.

Figure 7.13 shows the ASD of the time series servo voltage data measured during the charging noise measurement and the ASD of servo voltage data measured without the charged sample in the vacuum tank. When comparing the noise of the instrument with and without the charged sample in the vacuum chamber it can be seen that the noise clearly increases significantly when the charged sample is present.

The servo voltages shown in *Figure 7.11* are also strongly correlated with the Kelvin probe signal, further indicating that the increased noise observed in *Figure 7.13* is most likely due to the decay of the surface charge on the silica sample rather than just the presence of the charge. *Figure 7.14* shows a plot of




Figure 7.12: Plots of the typical servo voltages applied to the servo actuators when a charged sample is not present in the vacuum chamber.




Figure 7.13: A plot of the torque sensitivity of the torsion balance with a charged sample near the bob, red, and without a charged sample near the bob, black.

the linear correlation between one of the servo voltages and the Kelvin probe signal during the charging noise measurement.




Figure 7.14: A plot of the Kelvin probe voltage against the servo voltage being applied to one of the actuator plates during the charging measurement. Clearly there is a strong linear correlation between the two data sets.

7.4.1 Estimating the Expected Charging Noise Level

The level of charging noise expected from this measurement must be calculated to determine whether the noise observed is consistent with the Weiss theory of charging noise. Charging noise is dependent on the average Coulomb force, or torque in this instance, exerted on the torsion bob and the correlation time of the charge. So, to determine the charging noise signal that should be expected from this measurement the average DC torque acting on the bob and the correlation time of the charge must be estimated.

7.4.1.1 Estimating the DC Torque Exerted on the Torsion Bob

The DC Coulomb torque was estimated from the restoring torque acting on the torsion bob. The restoring torque and the Coulomb torque from the surface charge should be equal, but opposite, when the bob is perfectly centered between the sensor and actuator plates. The restoring torque comprises of two components: the restoring torque of the fibre and the torque applied by the servo actuators. The magnitude of the restoring torque of the fibre, Γ_1 , can be calculated from the angle through which the bob was rotated, ϕ , and the stiffness of the fibre, κ ,

$$\Gamma_1 = \kappa\phi. \quad (7.10)$$

From Chapter 6 it was found that κ was 6.5×10^{-8} Nm/rad and it was mentioned previously in this chapter that ϕ was 0.55 rad. This gives Γ_1 a value of 3.58×10^{-8} Nm.

The servo torque that is exerted on the bob by the servo actuators can be calculated using,

$$\Gamma_2 = \frac{1}{2} \frac{dC}{d\theta} (V_1 - V_2)^2, \quad (7.11)$$

where $\frac{dC}{d\theta}$ is the average capacitance gradient between the bob and the two servo plates, V_1 is the average servo voltage applied to servo actuator 1 over the course of the measurement and V_2 is the average servo voltage applied to servo actuator 2 over the course of the measurement. The average capacitance gradient was calculated to be 1.2×10^{-11} F/rad from measurements made in Chapter 6, the average voltage applied to servo actuator 1 was 11.26 V and the average voltage applied to servo actuator 2 was 14.96 V. This gives the value of Γ_2 to be -8.21×10^{-11} Nm, therefore, the combined restoring torque is 3.57×10^{-8} Nm.

7.4.1.2 Estimating the Charging Noise Correlation Time

The surface charge on the silica sample was monitored throughout the torsion balance measurement with a Kelvin probe. As mentioned in Section 7.2, the

probe wasn't directly measuring the charged surface of the silica sample because this surface had to face the bob. The Kelvin probe was located at the surface opposite the one that had charge on it and was measuring a polarisation charge. This should still give an accurate measurement of the correlation time because as the charge on the surface facing the bob decreases so should the induced charge being measured by the Kelvin probe. *Figure 7.15* shows the decaying Kelvin probe signal that was measured. The correlation time, τ_0 , of the charge can be obtained from the inverse of the time coefficient in the exponential term of the exponential fit to the data, also shown in *Figure 7.15*. τ_0 was measured to be 3.01×10^6 s or approximately 35 days.

7.4.1.3 Modeling Torque Noise Due to Charging

Once τ_0 and the average torque, $\langle \Gamma \rangle$, had been determined it was possible to calculate the level of charging noise that should be expected in the experiment using the Weiss theory of charging noise. The torque noise power spectrum was calculated using,

$$\Gamma_{x1}^2(f) = \frac{2\langle \Gamma^2 \rangle}{\pi \tau_0 \left(\frac{1}{\tau_0^2} + (2\pi f)^2 \right)}, \quad (7.12)$$

as stated in Chapter 2, except that force noise has now become torque noise. It was not necessary to use the finite element analysis that was described in Chapter 2 because the average torque was determined from the measurement. Therefore, it was just a simple case of calculating the torque noise for a range of frequencies using Equation (7.12).

It is noted in Appendix A that if the time series of charging noise can be described simply as a decaying exponential then charging noise should follow the fourier transform of a decaying exponential. When taking Fourier transforms the normalisation constant $1/\sqrt{2\pi}$ is often used. This would give charging




Figure 7.15: Plots of the decaying Kelvin probe signal measured during the charging noise measurement (top) and the residuals of the fit to the data (bottom). The correlation time of the charge was calculated to be 3.01×10^6 s from the exponential fit to the data, shown in red (top).

noise the form,

$$\Gamma_{x2}^2(f) = \frac{\langle \Gamma^2 \rangle}{2\pi\tau_0 \left(\frac{1}{\tau_0^2} + (2\pi f)^2 \right)}. \quad (7.13)$$

Both Equations (7.12) and (7.13) are essentially the same in principle, however, it is unknown to the author how the constant of $2/\pi$ appears in Equation (7.12). An estimate of the torque noise created by moving surface charge was calculated using both Equations (7.12) and (7.13). Both of these estimates are shown in Figure 7.16 along with the observed noise spectrum of the torsion balance measured during the charging noise measurement.

The most important part of this result is that the slope of the torsion




Figure 7.16: A plot of the torque sensitivity of the torsion balance with a charged silica disc present in the vacuum chamber (red) and estimates of the charging noise expected using the two different methods described in the text (blue and green). The region of uncertainty on the torsion balance measurement is between the two faint red lines. The error on the theoretical estimates are less than 1 %.

balance measurement matches the slope of the theoretical estimate. This would show that the charging noise has the frequency dependence expected. To show how well the slopes match, the ratio of the torsion balance measurement and one of the theoretical charging noise estimates (the blue line in *Figure 7.16*) are presented in *Figure 7.17*. The plot shows that the ratio of the two lines has a flat response which indicates that the slopes agree well.

From the measurements presented here it would appear that the torsion balance was indeed measuring noise created by the fluctuating forces from charges decaying on the surface of the silica disc. Equation (7.13), which uses a different normalisation constant to Equation (7.12), seems to give a better estimate of the increased noise level observed than Equation (7.12), which overestimates the noise level by nearly a factor of 2.5.

In conclusion this experiment successfully measured excess noise due to




Figure 7.17: A plot of the ratio of the torsion balance measurement and the theoretical estimate of charging noise against frequency. The flat response indicates that the two lines have the same gradient.

fluctuating electric charges on a dielectric surface. It was shown that charging noise (see [Figure 7.18](#)) has the frequency dependence as predicted by the Weiss theory of charging noise. This means that accurate estimates of charging noise can be made for advanced and future gravitational wave detectors which will be extremely useful when planning their construction.

7.5 Discussion

From the torsion balance measurement presented in this chapter it would appear that a direct measurement of charging noise has been made. Depositing surface charge on the surface of a silica disc and positioning it near the torsion balance significantly increased the noise level of the instrument. It is very unlikely that this increase in noise could have been the result of anything else other than the charge decaying on the silica disc due to the good correlation between the servo voltage signal and the Kelvin probe signal.




Figure 7.18: A plot of the total aLIGO noise and charging noise.

Equation (7.12) overestimated the level of noise expected by almost a factor of 2.5 and Equation (7.13) overestimated the level of noise expected by approximately 18 %. It is not entirely clear why the charging noise level observed is not as high as what is predicted by theory. As mentioned in this chapter, the most significant sources of error in the measurement are due to uncertainty in the torque estimate, which is used in both the gain measurements and the estimation of the theoretical charging noise level, and the calculation of the amplitude of the servo voltage oscillation during the calibration process. The uncertainty in the torque estimate should not really be more than stated because after the torsion bob is moved in a sine wave motion, with an amplitude of 100 stepper motor steps, the servo voltages generally settle on the values that they had before the bob was moved. It can be seen in *Figures 7.5 and 7.9* that there is not an obvious drift in the signal, thus indicating that the bob

has not taken too few or too many steps in a particular direction.

It is entirely possible that a larger error exists in the estimation of the average torque used in the theoretical estimates of the charging noise. This is because the bob had to be rotated through quite a large angle, which required almost 42,000 stepper motor steps, and this was done over a period of weeks. It is possible that some stepper motor steps were miscounted. However, it is doubtful that this additional error is on the scale of the deviation observed here.

It could be that the uncertainty in the estimate of the servo amplitude is greater than expected. As mentioned, it was unknown whether two of the gain data points in *Figure 7.10* were in fact dropping off as the frequency increased or if there was just a larger error in those data points. Repeated measurements should be taken of the servo gain at a particular frequency in order to calculate the standard deviation observed for a greater data population.

7.6 Conclusion and Future Work

In this chapter it was shown that moving surface charges do create a fluctuating force noise that can increase the noise level of sensitive instruments. When the Weiss theory of charging noise was used to estimate the level of charging noise expected during the measurement, it was found that it overestimated the noise level observed. Equation (7.12), the charging noise equation as stated in Weiss's technical paper, and Equation (7.13), essentially Equation (7.12) but with a different normalisation constant, were both used to estimate the charging noise expected. It was found that Equation (7.12) overestimated the level of charging noise by almost a factor of 2.5 and Equation (7.13) overestimated the level of charging noise by approximately 18%. It is believed that the errors considered in this experiment are accurate, however, it is entirely possible that there is some source of error that has been unaccounted for, or that one of the sources of error has been slightly underestimated, and this is

why there is some disparity between measurement and theory. Even with the slight constant offset between measurement and theory, the most important part of the measurement was that the frequency dependence of the measured and predicted charging noise agreed very well. Thus confirming that the theoretical prediction of charging noise discussed in this thesis should provide good estimates of charging noise for future gravitational wave detectors.

The torsion balance is a highly sensitive instrument that is extremely useful for measuring very weak forces. As shown in this experiment, it is possible to construct a torsion balance experiment of reasonable sensitivity without having an overly complicated set up. At the end of Chapter 6 there was a discussion on the development of a Mark II version of the torsion balance. With the improved sensitivity of the Mark II instrument over the one that was used for the measurements in this chapter, it should be possible to improve upon the measurement of charging noise presented in this thesis. This will then lead to measurements of the more subtle aspects of charging noise, such as the role played by polarisation charges and the additional forces generated by a charged moving earthquake stop. This new set up could also be used for the study of charging on materials that will be used for third generation detectors.

Chapter 8

Conclusions

Kilometer scale gravitational wave detectors have conducted several collaborative science runs in an attempt to directly detect gravitational waves. Even though there has not been a direct detection yet there are several astronomical sources that provide indirect evidence for the existence of gravitational waves. Within the next decade second generation detectors will be taking part in collaborative science runs that will no doubt detect the first gravitational waves. These detectors aim to have an improved sensitivity of at least a factor of ten greater than their predecessors. Improvements in the detectors sensitivity comes from years of collaborative work from researchers around the globe who have created new and innovative methods of reducing the main limiting noise sources of gravitational wave detectors.

As the sensitivity of gravitational wave detectors increase they could be more prone to noise created by the movement of excess electrical charges on the surface of the detector optics. If electrical charge is deposited on the optics of a gravitational wave detector it can compromise its control and sensitivity. It is thought that enough precautions have been taken for charging noise not to be a significant problem in second generation detectors, however, research of charging issues must continue to ensure that it will not be a problem in third generation detectors. It is likely that materials such as sapphire and silicon will be used in these detectors and it is unknown what level of charging noise

these materials could create if they became charged whilst under vacuum. It is also unknown whether the charge mitigation measures being taken for second generation detectors would be appropriate for third, and future, generation detectors.

For future gravitational wave detectors it is imperative that they have at least one in situ charge mitigation procedure. This procedure should effectively reduce excess surface charge on the detector optics, not cause damage to the reflective coating on the optics and minimise detector down time. It has been shown in this thesis that the use of argon, ionised by an AC glow discharge, as a means of reducing excess charge would be very effective in a gravitational wave detector such as aLIGO. There are also other studies that have shown that reducing excess surface charge can be accomplished by using ionised nitrogen. A study carried out by Weiss also showed that ionised nitrogen did not damage the TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ multi-layer reflective coating that will be used for aLIGO. It is unlikely that the ionised argon will cause significant damage to the aLIGO coating either, as it is used in the deposition process of the reflective coating. However, the main issue that has to be resolved for the process presented in this thesis is eliminating the possibility of UV being exposed to the TiO_2 doped $\text{Ta}_2\text{O}_5/\text{SiO}_2$ coating. As discussed here, this shouldn't be a significant problem. Although it is untested, this thesis presented a device which contained the glow discharge so that ionised argon can flow out of the device while leaving the UV trapped inside. This would also allow nitrogen gas to be used which is easier to pump out the vacuum chamber than argon. If this could be shown to work in practice it would provide a charge mitigation procedure to future detectors that is safe, easily set up and reproducible.

The moment when gravitational wave detector optics are at the highest risk of becoming charged is during the cleaning of the optics. It is still unknown whether mixing carbon nano-tubes (CNT) into the First ContactTM (FC) solution significantly decreases the resistivity of the FC when it dries to

leave a polymer film. However, as shown in this thesis, mixing high concentrations of CNT into FC will not deposit significantly less surface charge on the detector optics than ordinary FC. Therefore, the most promising method of reducing excess surface charge on the detector optics during the cleaning procedure is, again, through the use of ionised gas. Ionising nitrogen using a commercially available anti-static bar is a very easy method of significantly reducing the amount of charge deposited on the detector optics. This method did encounter some problems when trying to mitigate charge from between two closely spaced silica samples, but it is very difficult to recreate the kind of set up one would encounter in a real detector. It was still shown that this method should reduce surface charge to a level that will not limit aLIGO. Some additional work may have to be carried out at LASTI, or at the detectors themselves, in order to optimise this process for mitigating charge between the test mass and the reaction mass. Mitigating charge on insulators through use of ionised gas created by a corona discharge is a process that has been well researched and is used in many different types of industries and should be suitable for use in future gravitational wave detectors.

Possibly the most important measurement presented in this thesis is the direct measurement of charging noise with a servo controlled torsion balance. This result shows that charge deposited on the surface of an insulator does decay exponentially and produces a power spectrum of the form described by Weiss. The theoretical estimate of the level of charging noise expected agreed well with the noise level observed. This indicates that the Weiss theory of charging should give a good estimate of the charging noise one might expect in a gravitational wave detector. The charging noise measurement presented in this thesis was made with the Mark I instrument, however, there was discussion of the development of a Mark II instrument. This instrument should not only be easier to use, but should also provide greater sensitivity than the Mark I. This could provide more accurate measurements of charging noise and provide

additional insight into more subtle aspects of charging noise. The role of polarisation charges created in the silica test masses is an issue that has been becoming more of a concern in recent years. The Mark II version of the torsion balance could be used to see if these polarisation charges create additional noise that perhaps doesn't follow the same Markov assumption of the charging noise discussed here.

This is truly an exciting time in the gravitational wave community as we approach the beginning of a new era of second generation detectors. When these detectors can make routine detections of gravitational waves they will no doubt provide a wealth of new physics and astronomy. The Weiss theory of charging noise should allow accurate predictions of charging noise in gravitational wave detectors. This will be incredibly important in the design of third and future generation detectors. Charge mitigation procedures have been well researched for second generation detectors, however, there has been little discussion of how excess charge on the optics of third generation detectors will be handled. It is entirely possible that the methods discussed for second generation detectors can be directly applied to third and future generation detectors. However, with the possibility of new materials and optical coatings being used for the test masses of these detectors, it is unknown what effect the charge mitigation procedures already developed will have on these materials and coatings. Additional research should be conducted in order to ensure that future gravitational wave detectors do not risk having their low frequency sensitivity limited by charging noise.

Appendix A

Derivation of the Charging Noise Equation for an Exponentially Decaying Surface Charge

If it is assumed that surface charge decays exponentially across the surface of an optic, with a single correlation time τ_0 , it will create a fluctuating force, F_x , which acts on the detector optic,

$$F_x = \langle F_0 \rangle \exp \left(-\frac{t}{\tau_0} \right), \quad (\text{A.1})$$

where $\langle F_0 \rangle$ is the average Coulomb force acting on the optic and t is time. To convert this time series into a frequency spectrum a Fourier transform of Equation (A.1) must be taken and can be expressed as,

$$F_x(\omega) = \frac{\langle F_0 \rangle}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp \left(-\frac{t}{\tau_0} \right) \exp (i\omega t) dt, \quad (\text{A.2})$$

where ω is angular frequency. For this exponential signal it will be assumed that,

$$F_x = \begin{cases} \langle F_0 \rangle \exp \left(-\frac{t}{\tau_0} \right) & \text{for } t > 0 \\ 0 & \text{for } t < 0 \end{cases} \quad (\text{A.3})$$

Therefore, the limits of the integral in Equation (A.2) are changed to give,

$$F_x(\omega) = \frac{\langle F_0 \rangle}{\sqrt{2\pi}} \int_0^\infty \exp\left(-\left(\frac{1}{\tau_0} - i\omega\right)t\right) dt. \quad (\text{A.4})$$

Integrating Equation (A.4) gives,

$$F_x(\omega) = \frac{\langle F_0 \rangle}{\sqrt{2\pi} \left(\frac{1}{\tau_0} - i\omega\right)} \exp\left(-\left(\frac{1}{\tau_0} - i\omega\right)t\right) \Bigg|_0^\infty, \quad (\text{A.5})$$

$$F_x(\omega) = \frac{\langle F_0 \rangle}{\sqrt{2\pi} \left(\frac{1}{\tau_0} - i\omega\right)}. \quad (\text{A.6})$$

To remove the complex term in Equation (A.6) it is multiplied by its complex conjugate,

$$F_x^2(\omega) = \frac{\langle F_0^2 \rangle}{2\pi \left(\frac{1}{\tau_0} - i\omega\right) \left(\frac{1}{\tau_0} + i\omega\right)}, \quad (\text{A.7})$$

which gives,

$$F_x^2(f) = \frac{\langle F_0^2 \rangle}{2\pi \left(\frac{1}{\tau_0^2} + (2\pi f)^2\right)}. \quad (\text{A.8})$$

In order to have the correct units of $\text{Nm}/\sqrt{\text{Hz}}$, Equation (A.8) has to be divided by the bandwidth of the signal being measured. In this instance the bandwidth will be taken as the correlation time, τ_0 , to give the expression,

$$F_x^2(f) = \frac{\langle F_0^2 \rangle}{2\pi\tau_0 \left(\frac{1}{\tau_0^2} + (2\pi f)^2\right)}. \quad (\text{A.9})$$

This is slightly different from Equation (2.3) quoted in Chapter 2. It is possible that the author of the paper from which Equation (2.3) is quoted used a different normalisation constant in the Fourier transform of the exponential function. However, Equation (2.3) and Equation (A.9) are essentially the same in principle.

Appendix B

Additional SEM Images

There are more SEM images of the surface of a silica disc cleaned with FC mixed with multi-walled CNT, discussed in Chapter 5, shown below. The SEM images in *Figures B.1 to B.10* are taken at a magnification of 300 and the SEM images in *Figures B.11 to B.18* are taken at a magnification of 800. From these images it does not appear as if any carbon is deposited on the silica samples from the FC.




Figure B.1: *An SEM image of a silica disc before FC with CNT is applied. This image was taken at the edge of the sample.*




Figure B.2: *An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken at the edge of the sample.*




Figure B.3: An SEM image of a silica disc before FC with CNT is applied. This image was taken approximately 1 mm from the edge of the sample.




Figure B.4: An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken approximately 1 mm from the edge of the sample.




Figure B.5: An SEM image of a silica disc before FC with CNT is applied. This image was taken approximately 2 mm from the edge of the sample.




Figure B.6: An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken approximately 2 mm from the edge of the sample.




Figure B.7: An SEM image of a silica disc before FC with CNT is applied. This image was taken approximately 3 mm from the edge of the sample.




Figure B.8: An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken approximately 3 mm from the edge of the sample.




Figure B.9: An SEM image of a silica disc before FC with CNT is applied. This image was taken approximately 4 mm from the edge of the sample.




Figure B.10: An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken approximately 4 mm from the edge of the sample.




Figure B.11: An SEM image of a silica disc before FC with CNT is applied. This image was taken approximately 1 mm from the edge of the sample.




Figure B.12: An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken approximately 1 mm from the edge of the sample.




Figure B.13: An SEM image of a silica disc before FC with CNT is applied. This image was taken approximately 2 mm from the edge of the sample.




Figure B.14: An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken approximately 2 mm from the edge of the sample.




Figure B.15: An SEM image of a silica disc before FC with CNT is applied. This image was taken approximately 3 mm from the edge of the sample.




Figure B.16: An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken approximately 3 mm from the edge of the sample.




Figure B.17: An SEM image of a silica disc before FC with CNT is applied. This image was taken approximately 4 mm from the edge of the sample.




Figure B.18: An SEM image of a silica disc after FC with CNT has been removed from its surface. This image was taken approximately 4 mm from the edge of the sample.

Bibliography

- [1] A. Einstein. Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik*, **49**:769, 1916.
- [2] B. P. Abbott, R. Abbott, R. Adhikari, P. Ajith, B. Allen, G. Allen, et al. LIGO: The laser interferometer gravitational-wave observatory. *Reports on Progress in Physics*, **72**:076901, 2009.
- [3] T. Accadia et al. Status of the Virgo project. *Classical and Quantum Gravity*, **28**(11):114002, 2011.
- [4] H. Grote et al. The GEO 600 status. *Classical and Quantum Gravity*, **27**:084003, 2010.
- [5] R. Takahashi et al. Status of TAMA300. *Classical and Quantum Gravity*, **21**(5):S403–S408, 2004.
- [6] G. M. Harry et al. Advanced LIGO: the next generation of gravitational wave detectors. *Classical and Quantum Gravity*, **27**:084006, 2010.
- [7] F. Acernese et al. Advanced Virgo Baseline Design. Technical Report VIR-027A-09, 2009.
- [8] B. Willke et al. The GEO-HF project. *Classical and Quantum Gravity*, **23**:S207, 2006.
- [9] I. Newton. *Philosophiæ Naturalis Principia Mathematica*. London, 1st edition, 1687.

- [10] C. W. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation*. W. H. Freeman, *San Francisco*, 1973.
- [11] C. Cutler and K. S. Thorne. An overview of gravitational-wave sources. In *Proceedings of the 16th International Conference on General Relativity and Gravitation*, page 72, 2002.
- [12] B. F. Schutz. Gravitational-wave sources. *Classical and Quantum Gravity*, **13**:A219–A238, 1996.
- [13] B. F. Schutz. *A first course in general relativity*. Cambridge University Press, *Cambridge*, 1985.
- [14] B. S. Sathyaprakash and B. F. Schutz. Physics, astrophysics and cosmology with gravitational waves. *Living Reviews in Relativity*, **12**(2), 2009.
- [15] M. R. Abernathy et al. Einstein gravitational wave telescope conceptual design study. Technical Report ET-0106C-10, June 2011.
- [16] C. Messenger and G. Woan. A fast search strategy for gravitational waves from low-mass X-ray binaries. *Classical and Quantum Gravity*, **24**:S469, 2007.
- [17] K. S. Thorne. Gravitational radiation. In *Annals of the New York Academy of Sciences*, volume **759**, pages 127–152, 1995.
- [18] R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. *The Astrophysical Journal*, **195**:L51–L53, 1975.
- [19] J. H. Taylor, R. A. Hulse, L. A. Fowler, G. E. Gullahron, and J. M. Rankin. Further observations of the binary pulsar PSR 1913 + 16. *The Astrophysical Journal*, **206**:L53–L58, 1976.

- [20] P. R. Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors*. World Scientific Publishing, New Jersey, 1994.
- [21] J. H. Taylor. Binary pulsars and relativistic gravity. *Reviews of Modern Physics*, **66**:711–719, 1994.
- [22] B. P. Abbott, R. Abbott, F. Acernese, R. Adhikari, P. Ajith, B. Allen, et al. An upper limit on the stochastic gravitational-wave background of cosmological origin. *Nature*, **460**:990–994, 2009.
- [23] J. Weber. Observation of the thermal fluctuations of a gravitational-wave detector. *Physical Review Letters*, **17**:1228–1230, 1966.
- [24] J. Weber. Gravitational-wave-detector events. *Physical Review Letters*, **20**:1307–1308, 1968.
- [25] J. Weber. Evidence for discovery of gravitational radiation. *Physical Review Letters*, **22**:1320–1324, 1969.
- [26] J. L. Levine and R. L. Garwin. Absence of gravity-wave signals in a bar at 1695 Hz. *Physical Review Letters*, **31**:173–176, 1973.
- [27] J. A. Tyson. Null search for bursts of gravitational radiation. *Physical Review Letters*, **31**:326–329, 1973.
- [28] R. W. P. Drever, J. Hough, R. Bland, and G. W. Lessnoff. Search for short bursts of gravitational radiation. *Nature*, **246**:340–344, 1973.
- [29] V. B. Braginsky, A. B. Manukin, E. I. Popov, V. N. Rudenko, and A. A. Khorev. Upper limit of density of extraterrestrial gravitational radiation. *Soviet Physics, JETP*, **66**:801–812, 1974.
- [30] H. Grote and D. H. Reitze. First generation interferometric gravitational-wave detectors. In *Proceedings of the 46th Rencontres de Moriond QCD and High Energy Interactions Conference*, 2011.

- [31] M. E. Gertsenshtein and V. I. Pustovoit. On the detection of low frequency gravitational waves. *Soviet Physics, JETP*, **16**:433–435, 1962.
- [32] LIGO timeline of events (<http://www.ligo-la.caltech.edu/LLO/-time.htm>) accessed 22nd Feb 2013.
- [33] S. Braccini, C. Casciano, F. Cordero, F. Corvace, M. De Sanctis, R. Franco, et al. The maraging-steel blades of the Virgo super attenuator. *Measurement Science and Technology*, **11**:467–476, 2000.
- [34] R. Abbott, R. Adhikari, G. Allen, S. Cowley, E. Daw, D. DeBra, et al. Seismic isolation for advanced LIGO. *Classical and Quantum Gravity*, **19**:1591–1597, 2002.
- [35] J. Heefner, R. Bork, and R. Abbott. The LIGO interferometer sensing and control system. In *Proceedings of the 8th International Conference on Accelerator and Large Experimental Physics Control Systems (ICALPCS 2001)*, pages 392–394, 2001.
- [36] T. Accadia et al. The status of Virgo. *Classical and Quantum Gravity*, **23**(8):S63–S69, 2006.
- [37] H. Grote, G. Heinzel, A. Freise, S. Goßler, B. Willke, H. Lück, et al. Alignment control of GEO 600. *Classical and Quantum Gravity*, **21**:S441, 2004.
- [38] R. W. Drever. A gravity-wave detector using optical cavity sensing. In *Proceedings of the 9th International Conference on General Relativity and Gravitation*, page 265, 1983.
- [39] R. Weiss. Electronically coupled broadband gravitational antenna. *Quarterly Progress Report, Research Laboratory of Electronics (MIT)*, **105**:54, 1972.

- [40] B. J. Meers. Recycling in laser-interferometric gravitational-wave detectors. *Physical Review D*, **38**:2317–2326, 1988.
- [41] J. R. Smith et al. The path to the enhanced and advanced LIGO gravitational-wave detectors. *Classical and Quantum Gravity*, **26**:114013, 2009.
- [42] H. Grote et al. The status of GEO 600. *Classical and Quantum Gravity*, **25**:114043, 2008.
- [43] B. Abbott et al. Astrophysically triggered searches for gravitational waves: status and prospects. *Classical and Quantum Gravity*, **25**:114051, 2008.
- [44] B. P. Abbott, R. Abbott, F. Acernese, R. Adhikari, P. Ajith, B. Allen, et al. Searches for gravitational waves from known pulsars with science run 5 LIGO data. *The Astrophysical Journal*, **713**:671–685, 2010.
- [45] J. Abadie et al. Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors. *Classical and Quantum Gravity*, **27**:173001, 2010.
- [46] The LIGO Scientific Collaboration and the Virgo Collaboration. *Commissioning and Observing Scenarios for the Advanced LIGO and Advanced Virgo Gravitational-Wave Observatories*. Internal LIGO Document P1200087-v4, 2012.
- [47] K. Kuroda. Status of LCGT. *Classical and Quantum Gravity*, **27**:084004, 2010.
- [48] J-P. Zendri, L. Baggio, M. Bignotto, M. Bonaldi, M. Cerdonio, L. Conti, et al. Status report and near future prospects for the gravitational wave detector AURIGA. *Classical and Quantum Gravity*, **19**:1925–1933, 2002.

- [49] S. Hild. Beyond the second generation of laser-interferometric gravitational wave observatories. *Classical and Quantum Gravity*, **29**:124006, 2012.
- [50] K. Danzmann, T. A. Prince, P. Binetruy, P. Bender, S. Buchman, J. Centrella, et al. LISA Unveiling a hidden Universe. Technical Report ESA/SRE(2011)3, 2011.
- [51] P. Amaro-Seoane, S. Aoudia, S. Babak, P. Binétruy, E. Berti, A. Bohé, et al. Low-frequency gravitational-wave science with eLISA/NGO. *Classical and Quantum Gravity*, **29**:124016, 2012.
- [52] GWINC wiki page (<https://awiki.ligo-wa.caltech.edu/aLIGO/GWINC>) accessed 22nd Feb 2013.
- [53] S. A. Hughes and K. S. Thorne. Seismic gravity-gradient noise in interferometric gravitational-wave detectors. *Physical Review D*, **58**:122002, 1998.
- [54] P. R. Saulson. Terrestrial gravitational noise on a gravitational wave antenna. *Physical Review D*, **30**:732–736, 1984.
- [55] N. A. Robertson, G. Cagnoli, D. R. M. Crooks, E. Elliffe, J. Faller, P. Fritschel, et al. Quadruple suspension design for advanced LIGO. *Classical and Quantum Gravity*, **19**:4043, 2002.
- [56] M. V. Plissi, C. I. Torrie, M. E. Husman, N. A. Robertson, K. Strain, H. Ward, et al. GEO 600 triple pendulum suspension system: Seismic isolation and control. *Review of Scientific Instruments*, **71**:2539–2545, 2000.
- [57] M. Evans, S. Ballmer, M. Fejer, P. Fritschel, G. Harry, and G. Ogin. Thermo-optic noise in coated mirrors for high-precision optical measurements. *Physical Review D*, **78**:102003, 2008.

- [58] K. Kawabe and the LIGO Collaboration. Status of LIGO. *Journal of Physics Conference Series*, **120**:032003, 2008.
- [59] A. Amico, L. Bosi, L. Carbone, L. Gammaitoni, M. Punturo, F. Travasso, et al. Fused silica suspensions for the Virgo optics: status and perspectives. *Classical and Quantum Gravity*, **19**:1669–1974, 2002.
- [60] G. M. Harry, M. R. Abernathy, A. E. Becerra-Toledo, H. Armandula, E. Black, K. Dooley, et al. Titania-doped tantala/silica coatings for gravitational-wave detection. *Classical and Quantum Gravity*, **24**:405–415, 2006.
- [61] R. Bassiri, K. B. Borisenko, D. J. H. Cockayne, J. Hough, I. MacLaren, and S. Rowan. Probing the atomic structure of amorphous Ta₂O₅ mirror coatings for advanced gravitational wave detectors using transmission electron microscopy. *Journal of Physics Conference Series*, **241**:012070, 2010.
- [62] M. R. Abernathy, S. Reid, E. Chalkley, R. Bassiri, I. W. Martin, K. Evans, et al. Cryogenic mechanical loss measurements of heat-treated hafnium dioxide. *Classical and Quantum Gravity*, **28**:195017, 2011.
- [63] M. Pitkin, S. Reid, J. Hough, and S. Rowan. Gravitational wave detection by interferometry (ground and space). *Living Reviews in Relativity*, **14**(5), 2011.
- [64] C. M. Caves. Quantum-mechanical noise in an interferometer. *Physical Review D*, **23**:1693–1703, 1981.
- [65] C. M. Caves. Quantum-mechanical radiation-pressure fluctuations in an interferometer. *Physical Review Letters*, **45**:75–79, 1980.

- [66] V. Henning, S. Chelkowski, B. Hage, A. Franzen, K. Danzmann, and R. Schnabel. Coherent control of vacuum squeezing in the gravitational-wave detection band. *Physical Review Letters*, **97**:011101, 2006.
- [67] M. Hewitson, K. Danzmann, H. Grote, S. Hild, J. Hough, H. Lück, et al. Charge measurement and mitigation for the main test masses of the GEO 600 gravitational wave observatory. *Classical and Quantum Gravity*, **24**:6379, 2007.
- [68] S. Rowan, S. Twyford, R. Hutchins, and J. Hough. Investigations into the effects of electrostatic charge on the Q factor of a prototype fused silica suspension for use in gravitational wave detectors. *Classical and Quantum Gravity*, **14**:1537–1541, 1997.
- [69] M. J. Mortonson, C. C. Vassiliou, D. J. Ottaway, D. H. Shoemaker, and G. M. Harry. Effects of electrical charging on the mechanical Q of a fused silica disk. *Review of Scientific Instruments*, **74**:4840–4845, 2003.
- [70] D. Ugolini, R. Amin, G. M. Harry, J. Hough, I. Martin, V. Mitrofanov, et al. Charging issues in LIGO. In *Proceedings of the 30th International Cosmic Ray Conference*, volume **3**, pages 1283–1286, 2008.
- [71] D. Ugolini, M. Girard, G. M. Harry, and V. P. Mitrofanov. Discharging fused silica test masses with ultraviolet light. *Physics Letters A*, **372**(36):5741–5744, 2008.
- [72] S. Buchman, R. L. Byer, D. Gill, N. A. Robertson, and K-X Sun. Charge neutralization in vacuum for non-conducting and isolated objects using directed low-energy electron and ion beams. *Classical and Quantum Gravity*, **25**:035004, 2008.
- [73] S. Reid, I. W. Martin, A. Cumming, L. Cunningham, M. M. Fejer, G. Hammond, et al. The mechanical loss of tin (II) oxide thin-film

- coatings for charge mitigation in future gravitational wave detectors. *Classical and Quantum Gravity*, **29**:035002, 2012.
- [74] D. Ugolini, Q. Funk, and T. Amen. Note: Discharging fused silica test masses with ionised nitrogen. *Review of Scientific Instruments*, **82**:046108, 2011.
- [75] R. Weiss. *Surface charge control of the advanced LIGO mirrors using externally introduced ions*. Internal LIGO Technical Note T1100332-v1, 2010.
- [76] P. Campsie, L. Cunningham, M. Hendry, J. Hough, S. Reid, S. Rowan, and G. D. Hammond. Charge mitigation techniques using glow and corona discharges for advanced gravitational wave detectors. *Classical and Quantum Gravity*, **28**:215016, 2011.
- [77] P. F. O’Grady. *Thales of Miletus: The Beginnings of Western Science and Philosophy*. Ashgate, Surrey, U.K., 2002.
- [78] J. Lowell and A. C. Rose-Innes. Contact electrification. *Advances in Physics*, **29**:947–1023, 1980.
- [79] W. R. Harper. *Contact and Frictional Electrification*. Oxford University Press, Oxford, U.K., 1967.
- [80] J. Lowell and W. S. Truscott. Triboelectrification of identical insulators: II. Theory and further experiments. *Journal of Physics D: Applied Physics*, **19**:1281–1298, 1986.
- [81] P. E. Shaw. Experiments on tribo-electricity. I. The tribo-electric series. In *Proceedings of the Royal Society of London Series A*, volume **94**, pages 16–33, 1917.

- [82] P. E. Shaw and C. S. Jex. Tribo-electricity and friction. III. Solid elements and textiles. In *Proceedings of the Royal Society of London Series A*, volume **118**, pages 108–113, 1928.
- [83] L. S. McCarty and G. M. Whitesides. Electrostatic charging due to separation of ions at interfaces: Contact electrification of ionic electrets. *Angewandte Chemie International Edition*, **47**:2188–2207, 2008.
- [84] R. Weiss. *Note on electrostatics in the LIGO suspensions*. Internal LIGO Technical Note T960137-00, 1995.
- [85] L. G. Prokhorov, P. E. Khramchenkov, and V. P. Mitrofanov. Measurement of relaxation of electrical charge distribution on fused silica sample. *Physics Letters A*, **366**:145–149, 2007.
- [86] R. Adhikari, G. Mueller, N. Robertson, D. McClelland, Y. Chen, S. Waldman, et al. *LSC Instrument Science White Paper*. Internal LIGO Technical Note T1100309-v5, 2011.
- [87] B. Lantz. *Simple calculations for charge noise for advanced LIGO*. Internal LIGO Technical Note T080214-v1, 2008.
- [88] J. Kissel. *Advanced LIGO Active Seismic Isolation*. Internal LIGO Document G1100431-01, 2011.
- [89] R. Amin. *Brief of the May 2006 LIGO Charging Event*. Internal LIGO Technical Note G070572-x0, 2007.
- [90] D. Ugolini and C. Walker. *Charge Deposition from Earthquake Stop Materials*. Internal LIGO Technical Note G0900132-v1, 2009.
- [91] E. Gustafson, G. Harry, and N. Robertson. *Charging Work - An Update*. Internal LIGO Technical Note G1000460-v1, 2010.

- [92] Photonic Cleaning Technologies homepage (<https://www.photoniccleaning.com>) accessed 22nd Feb 2013.
- [93] G. M. Harry. *Estimate of noise from charging in initial, enhanced and advanced LIGO*. Internal LIGO Technical Note T080019-01, 2008.
- [94] V. B. Braginsky, O. G. Ryazhskaya, and S. P. Vyatchanin. Limitations in quantum measurements resolution created by cosmic rays. *Physics Letters A*, **359**:86–89, 2006.
- [95] V. P. Mitrofanov, L. G. Prokhorov, and K. V. Tokmakov. Variation of electric charge on prototype of fused silica test mass of gravitational wave antenna. *Physics Letters A*, **300**:370–374, 2002.
- [96] V. P. Mitrofanov, L. G. Prokhorov, K. V. Tokmakov, and P. Willems. Investigation of effects associated with variation of electric charge on a fused silica test mass. *Classical and Quantum Gravity*, **21**:S1083–S1089, 2004.
- [97] T. Sumner, H. Araújo, D. Davidge, A. Howards, C. Lee, G. Rochester, et al. Description of charging/discharging processes of the LISA sensors. *Classical and Quantum Gravity*, **21**:S597, 2004.
- [98] K-X. Sun, B. Allard, S. Buchman, S. Williams, and R. Byer. LED deep UV source for charge management of gravitational reference sensors. *Classical and Quantum Gravity*, **23**:S597, 2004.
- [99] H. Armandula, S. Buchman, R. Byer, M. Fejer, G. M. Harry, N. Lein-decker, et al. *Effects of Ultraviolet Irradiation on LIGO Mirror Coating*. Internal LIGO Technical Note G080150-x0, 2008.
- [100] K. A. Strain, K. Danzmann, J. Mizuno, P. G. Nelson, A. Rüdiger, R. Schilling, and W. Winkler. Thermal lensing in recycling interfero-

- metric gravitational wave detectors. *Physics Letters A*, **194**:124–132, 1994.
- [101] S. Buchman and K-X. Sun. *Charge Management with UV-LED Space Experiment*. Internal LIGO Technical Note G1100380-v1, 2011.
- [102] K. Besocke and S. Berger. Piezoelectric driven Kelvin probe for contact potential difference studies. *Review of Scientific Instruments*, **47**:840–842, 1967.
- [103] N. A. Robertson, J. R. Blackwood, S. Buchman, R. L. Byer, J. Camp, and D. Gill. Kelvin probe measurements: investigations of the patch effect with applications to ST-7 and LISA. *Classical and Quantum Gravity*, **23**:2665, 2006.
- [104] D. K. Davies. The examination of the electrical properties of insulators by surface charge measurement. *Journal of Scientific Instruments*, **44**:521–524, 1967.
- [105] T. R. Foord. Measurement of the distribution of surface electric charge by use of a capacitive probe. *Journal of Physics E: Scientific Instruments*, **2**:411–413, 1969.
- [106] W. Thomson. Contact electricity of metals. *Philosophical Magazine*, **46**:82–120, 1898.
- [107] W. A. Zisman. A new method of measuring contact potential differences in metals. *Review of Scientific Instruments*, **3**:367–370, 1932.
- [108] M. Nonnenmacher, M. P. O’Boyle, and H. K. Wickramasinghe. Kelvin probe force microscopy. *Applied Physics Letters*, **58**:2921–2923, 1991.
- [109] J. M. R. Weaver and D. W. Abraham. High-resolution atomic force microscopy potentiometry. *Applied Physics Letters*, **58**:2921–2923, 1991.

- [110] Besocke Delta Phi homepage (<http://www.besocke-delta-phi.de/-index.htm>) accessed 22nd Feb 2013.
- [111] G. J. Sloggett, N. G. Barton, and S. J. Spencer. Fringing fields in disc capacitors. *Journal of Physics A: Mathematical and General*, **19**:2725–2736, 1986.
- [112] Matweb - Material Property Data (<http://www.matweb.com/search/DataSheet.aspx?MatGUID=a028dc09086e4e6f9885e90ae49ecc32>) accessed 22nd Feb 2013.
- [113] Matweb - Material Property Data (<http://www.matweb.com/search/DataSheet.aspx?MatGUID=dc842aba6fa24c48b7e0155ee448d6f0>) accessed 22nd Feb 2013.
- [114] R. J. D’Arcy and N. A. Surplice. The effects of stray capacitance on the Kelvin method for measuring contact potential difference. *Journal of Physics D: Applied Physics*, **3**:482–488, 1970.
- [115] N. A. Surplice and R. J. D’Arcy. A critique of the Kelvin method of measuring work functions. *Journal of Physics E: Scientific Instruments*, **3**:477–482, 1970.
- [116] I. D. Baikie, E. Venderbosch, J. A. Meyer, and P. J. Z. Estrup. Analysis of stray capacitance in the Kelvin method. *Review of Scientific Instruments*, **62**:725–735, 1991.
- [117] T. Sugimoto and Y. Higashiyama. Surface potential of insulator plate equipped with a holed-grounded electrode charged by ion spray. In *Conference Record of the 2005 IEEE Industrial Applications Conference*, volume **2**, pages 1423–1428, 2005.
- [118] J. R. Roth. *Industrial Plasma Engineering Volume 1: Principles*. IOP Publishing, Bristol, U.K., 1995.

- [119] T. Farouk, B. Farouk, A. Gutsol, and A. Fridman. Atmospheric pressure methane-hydrogen DC micro-glow discharge for thin film deposition. *Journal of Physics D: Applied Physics*, **41**:175202, 2008.
- [120] A. Ohsawa. Precisely balanced ionizer using atmospheric pressure glow discharge in air. *Journal of Electrostatics*, **63**:45–57, 2005.
- [121] F. Paschen. Ueber die zum Funkenübergang in Luft, Wasserstoff und Kohlensäure bei verschiedenen Drucken erforderliche Potentialdifferenz. *Annalen der Physik*, **273**:69–96, 1889.
- [122] K. T. A. L. Burm. Calculation of the Townsend discharge coefficients and the Paschen curve coefficients. *Contributions to Plasma Physics*, **47**:177–182, 2007.
- [123] J. D. Cobine. *Gaseous Conductors*. Dover Publications, New York, U.S.A., 1958.
- [124] G. Harry, T. P. Bodiya, and R. DeSalvo, editors. *Optical Coatings and Thermal Noise in Precision Measurement*. Cambridge University Press, New York, U.S.A., 2012.
- [125] K. L. Brown and G. W. Tautfest. Faraday cup monitors for high energy electron beams. *Review of Scientific Instruments*, **27**:693–702, 1956.
- [126] A. Visioli. *Practical PID Control*. Springer, London, U.K., 2006.
- [127] E. Kreyszig. *Introductory Mathematical Statistics: Principles and Methods*. Wiley, New York, U.S.A., 1970.
- [128] F. J. Massey Jr. The Kolmogorov-Smirnov test for goodness of fit. *Journal of the American Statistical Association*, **46**:7504–78, 2007.
- [129] Varian Inc. Vacuum Technologies. *The Turbo-V 3K-G System*. Varian Inc. Vacuum Technologies Brochure, 2012.

- [130] A. Alexandrovski, M. Fejer, A. Markosyan, and R. Route. Photothermal common-path interferometry (PCI): new developments. In *Solid State Lasers XIII: Technology and Devices Proceedings of the Society of Photographic Instrumentation Engineers (SPIE)*, volume **7193**, 2009.
- [131] A. Markosyan. *Study of the UV exposure effect on the optical absorption of Ad-LIGO optimized mirrors with the PCI technique*. Internal LIGO Technical Note G1000249-v1, 2010.
- [132] R. Dannenberg. *Running Log of COC Cleaning Development*. Internal LIGO Document T0900351-v8, 2009.
- [133] W. R. Thruber, J. R. Lowney, R. D. Larrabee, and J. R. Ehrstein. AC impedance method for high-resistivity measurements of silicon. *Journal of The Electrochemical Society*, **138**:3081–3085, 1991.
- [134] P. G. Muzykov, Y. I. Khlebnikov, S. V. Regula, Y. Gao, and T. S. Sudarshan. High resistivity measurement of SiC wafers using different techniques. *Journal of Electronic Materials*, **32**:505–510, 2003.
- [135] A. R. Blythe. Electrical resistivity measurements of polymer materials. *Polymer Testing*, **4**:195–209, 1984.
- [136] J. Choi, J-H. Jung, J. T. Lee, K. H. Lee, B. H. Yim, C-Y. Choi, et al. Investigation on reduction technique of anomalous leakage current generation in CMOS image sensor adopting n-type substrate. *Sensors and Actuators A: Physical*, **167**:14–18, 2011.
- [137] [Static Clean International homepage \(http://www.staticclean.com\)](http://www.staticclean.com) accessed 22nd Feb 2013.
- [138] M. Phelps. *First ContactTM application and removal procedure*. Internal LIGO Document E1000079-v3, 2010.

- [139] R. S. Amin. *Brief of the May 2006 LIGO Charging Event*. Internal LIGO Document G070572-00, 2007.
- [140] H. Cavendish. Experiments to determine the density of the Earth. *Philosophical Transactions of the Royal Society of London*, **88**:469–526, 1798.
- [141] G. G. Luther and W. R. Towler. Redetermination of the Newtonian gravitational constant G . *Physical Review Letters*, **48**:121–123, 1982.
- [142] J. H. Gundlach and S. M. Merkowitz. Measurement of Newton’s constant using a torsion balance with angular acceleration feedback. *Physical Review Letters*, **85**:2869–2872, 2000.
- [143] G. T. Gillies and R. C. Ritter. Torsion balances, torsion pendulums, and related devices. *Review of Scientific Instruments*, **64**:283–309, 1993.
- [144] I. Cruceru, G. Nicolescu, and O. G. Dului. A dedicated torsion balance to detect neutrinos by coherent scattering on high Debye temperature monocrystals. In *Journal of Physics: Conference Series*, volume **136**, page 042089, 2008.
- [145] L. Carbone, A. Cavalleri, G. Ciani, R. Dolesi, M. Hueller, D. Tombalato, et al. Torsion pendulum facility for direct force measurements of LISA GRS related disturbances. In *AIP Conference Proceedings*, volume **873**, pages 561–565, 2006.
- [146] G. D. Hammond, A. Pulido-Paton, C. C. Speake, and C. Trenkel. Novel torsion balance based on a spherical superconducting suspension. *Review of Scientific Instruments*, **75**:955–961, 2004.
- [147] Matweb - Material Property Data (<http://www.matweb.com/search/DataSheet.aspx?MatGUID=37fde7152eb94e3dbcbde8e2edc0a497>) accessed 22nd Feb 2013.

- [148] Y. Su, B. R. Heckel, E. G. Adelberger, J. H. Gundlach, M. Harris, G. L. Smith, et al. New tests of the universality of free fall. *Physical Review D*, **50**:3614–3636, 1994.
- [149] A. Savitzky and M. J. E. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, **36**:1627–1639, 1964.
- [150] G. D. Hammond, C. C. Speake, A. J. Matthews, E. Rocco, and F. Peña-Arellano. Development of a second generation torsion balance based on a spherical superconducting suspension. *Review of Scientific Instruments*, **79**:025103, 2008.
- [151] SUPA Advanced Data Analysis Course, Prof. Martin Hendry, January 2009 (http://www.astro.gla.ac.uk/users/martin/supa-da/-supa_ada2011_sec2.pdf) accessed 22nd Feb 2013.
- [152] S. Schlamminger, C. A. Hagedorn, M. G. Famulare, S. E. Pollack, and J. H. Gundlach. High sensitivity torsion balance tests for LISA proof mass modeling. In *AIP Conference Proceedings*, volume **873**, pages 151–157, 2006.
- [153] K. Kuroda. Anelasticity in g experiments. *Measurement Science and Technology*, **10**:435–438, 1999.
- [154] Y. T. Chen, A. H. Cook, and A. J. F. Metherell. An experimental test of the inverse square law of gravitation at range of 0.1m. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, **349**:47–68, 1984.
- [155] S. Schlamminger, K.-Y. Choi, T. A. Wagner, J. H. Gundlach, and E. G. Adelberger. Test of the equivalence principle using a rotating torsion balance. *Physical Review Letters*, **100**:041101, 2008.

- [156] J. K. Hoskins, R. D. Newman, R. Spero, and J. Schultz. Experimental tests of the gravitational inverse-square law for mass separations from 2 to 105cm. *Physical Review D*, **32**:3084–3095, 1985.
- [157] G. Linfield and J. Penny. *Numerical Methods using Matlab*. Prentice Hall, New Jersey, 2nd edition, 1999.
- [158] E. G. Adelberger, J. H. Gundlach, B. R. Heckel, S. Hoedl, and S. Schlamminger. Torsion balance experiments: A low-energy frontier of particle physics. *Progress in Particle and Nuclear Physics*, **62**:102–134, 2009.
- [159] P. R. Saulson. Thermal noise in mechanical experiments. *Physical Review D*, **42**:2437–2445, 1990.
- [160] S. E. Pollack, M. D. Turner, S. Schlamminger, C. A. Hagedorn, and J. H. Gundlach. Charge management for gravitational-wave observatories using UV LEDs. *Physical Review D*, **81**:021101, 2010.
- [161] A. Cavalleri, G. Ciani, R. Dolesi, M. Hueller, D. Nicolodi, D. Tombolato, et al. Direct force measurements for testing the LISA pathfinder gravitational reference sensor. *Classical and Quantum Gravity*, **26**:094012, 2009.
- [162] J. B. Wachtman, W. E. Tefft, D. G. Lam, and C. S. Apstein. Exponential temperature dependence of Young’s modulus for several oxides. *Physical Review*, **122**:1754–1759, 1961.