



University
of Glasgow

Evans, Keith Lee (2012) *Computational modelling of amorphous mirror coatings for use in advanced gravitational wave detectors.*

PhD thesis

<http://theses.gla.ac.uk/4457/>

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

Computational Modelling of Amorphous mirror
coatings for use in Advanced Gravitational wave
detectors

Keith Lee Evans

School of Physics and Astronomy

College of Science and Engineering

University of Glasgow

Presented as a thesis for the degree of Ph.D

Supervised by Dr Ian Maclaren and Professor Sheila Rowan

November 2012

Abstract

Albert Einstein in 1916 predicted the existence of a Gravitational Wave in his General Theory of Relativity. These waves, which propagate at the speed of light transmit gravitational information through the Universe. Since its prediction by Einstein, astronomers and physicists have searched for them and developed method to detect them. Though so far unsuccessful, the search of Gravitational waves goes on and great efforts are being made to develop the most sensitive detectors yet in the hope of that first detection.

Currently ground based detectors are limited by coating Brownian thermal noise due to excitation of the reflective coatings applied to the test masses. Through measurement of mechanical loss of a material the magnitude of the Brownian thermal noise can be determined. It is necessary to determine the root cause of mechanical loss in current coatings (Ta_2O_5 and SiO_2). Work towards this goal is taking place on multi paths, directly, through characterisation of mechanical loss and indirectly through microscopy studies to determine the structural cause.

In this thesis, the effect of TiO_2 doping and heat treatment of Ta_2O_5 has been investigated. It has been previously shown that a TiO_2 doping of Ta_2O_5 reduces the mechanical loss and that, that reduction is at a maximum at 30% TiO_2 . It has been determined through Electron Diffraction experiments that the structure of TiO_2 doped Ta_2O_5 becomes more homogenous up to 30% doping. Through computation modelling of these structures using Density Functional Theory it has also been determined that the abundance of TiTaO_2 ring formations also maximises at 30% doping. Further modelling has also determined that the TiTaO_2 rings are more flexible than their counter parts of Ta_2O_2 and Ti_2O_2 . From this it has been hypothesised that the overall flexibility of a structure is a strong component of the mechanical response of the structure. Hence by increasing the flexibility through TiO_2 doping the mechanical loss (as Thermal Noise) is decreased similarly and this response would also be expected using similarly flexibility improving dopants.

Table of Contents

Table of Contents	i
Tables of Figures	v
Acknowledgements.....	xvi
Preface	xviii
Chapter 1: Introduction to gravitational waves, advanced interferometry and Thermal noise	1
1.0 Chapter Introduction	1
1.1 Gravitational Waves	1
1.1.1 What are Gravitational Waves?.....	1
1.1.2 Why are they important?	2
1.1.3 How do we detect Gravitational Waves?	4
1.2 Interferometric Gravitational Wave (IGW) detectors	5
1.3 Sensitivity limits.....	8
1.3.1 Seismic Noise	9
1.3.2 Gravitational Gradient Noise	10
1.3.3 Quantum Noise	11
1.3.4 Thermal Noise and Mechanical Loss	12
1.4 Mechanical Loss.....	13
1.4.1 Measuring Mechanical Loss.....	13
1.4.2 Temperature Dependent Mechanical Loss.....	14
1.4.3 The Importance of Tantala and Titania Doped Tantala.....	14
1.5 Coatings and modelling	17
1.5.1 Crystalline Materials	18
1.5.2 Amorphous Materials	19
1.6 Chapter Summary.....	22

Chapter 2: Methods of Experimental and Analytical Transmission Electron Microscopy and Density Functional Theorem Constrained Modelling.	23
2.0 Chapter Introduction	23
2.1 Transmission Electron Microscopy and Microscopes	24
2.1.1 Electron Sources	26
2.1.2 Electron Lenses	28
2.1.3 Physical Apertures	30
2.1.4 Typical Detectors used in electron microscopy.....	31
2.1.5 Cross Section Encapsulation Sample Preparation	32
2.1.6 Electron Diffraction.....	34
2.1.7 Electron Energy loss spectroscopy	37
2.2 Reduced Density Functions	40
2.2.1 Practical Application of RDF method	44
2.3 Reverse Monte Carlo	47
2.4 CASTEP and Density Functional Theorem	50
2.4.1 Density Functional Theorem.....	50
2.4.2 The Kohn-Sham Equations.....	51
2.4.3 Exchange Correlation Functionals	52
2.4.4 Pseudopotentials	52
2.4.5 Self Consistent Field Theory Iterations.....	53
2.4.6 Molecular Dynamics	54
2.4.7 Geometry optimisation.....	56
2.5 Cohesive Energy	57
2.6 X-ray Reflectometry.....	58
2.7 Chapter Summary.....	60
Chapter 3: Developement of a Semi-empirical modelling process	62
3.0 Chapter Introduction	62

3.1 CBED improvements.....	62
3.1.1 SIS Camera Adjustments.....	63
3.1.2 High Dynamic Range Imaging	65
3.1.3 Qmax increase spliced images.....	75
3.2 X-ray Diffraction	78
3.3 Improving the Quality of Reduced Density Functions.....	81
3.3.1 High Dynamic Range Reduced Density Functions	83
3.3.2 Reduced Density Function Analytical Techniques.....	87
3.3.3 Area Mapping	88
3.4 Coating Density Measurements from X-ray Reflectometry.....	95
3.5 Improving Density Functional Theorem Simulations.....	100
3.5.1 Decreased Simulation time.....	104
3.5.2 Molecular Dynamics simulations vs Geometry Optimisations.....	108
3.5.3 Molecular Dynamics with Geometry optimisation	110
3.5.4 Iterative Modelling Method	111
3.6 Chapter Summary.....	116
Chapter 4: The effect of Titania doping on Tantalum.....	118
4.0 Chapter Introduction	118
4.1 Experimental Analysis.....	119
4.2 Modelling Analysis.....	126
4.2.1 Analysis of models upon completion of the DFT stage of modelling.....	126
4.2.2 The role of Planar Fragments in Titania doped Tantalum	134
4.2.3 Analysis of the final models	139
4.2.4 Analysis of Density Functional Theory Models Performed Using Optimised Parameters.....	145
4.3 Chapter Summary.....	150
Chapter 5: Modelling of Annealed Titanium Dioxide doped Tantalum Pentoxide	152

5.0 Chapter Introduction	152
5.1 Experimental Analysis.....	152
5.1.1 Image Analysis	152
5.2 Modelling Analysis.....	164
5.3 Chapter Summary	166
Chapter 6: General Conclusions	169
Conclusions on the improvement of the techniques	169
Conclusions on the structure of Titania doped Tantalum	171
Possible Future Work	173

Tables of Figures

Figure 1.1.3.1: Timeline diagram of the space surrounding a bar-type GW detector as a GW propagates through it. The space (and the bar) expands in one direction, returns to its equilibrium position and then expands in the orthogonal direction before returning to its equilibrium position. The magnitude of a gravitational wave 'h' is defined as $h = 2\Delta L/L$ where L is the initial distance between two points in space and ΔL is the difference in distance induced by the passage of gravitational wave of magnitude h.4

Figure 1.2.1. Diagram of a basic Michelson-type laser interferometer showing the laser (a), beam splitter (b) the reflecting mirrors (c1 and c2) the arms (d1 and d2) and the detector (e)6

Figure 1.3.1.1. Diagram (not to scale) of the quad suspension similar to the design of the aLIGO test masses, from the front (left) and the side (right). The final suspension layer is quasi-monolithic i.e. that the mass and the wires are one continuous piece of fused silica.10

Figure 1.4.3.1. Mechanical loss of amorphous Tantalum heat treated at 300°C, 400°C, 600°C and 800°C as a function of temperature.....15

Figure 1.4.3.2: Convergent beam electron diffraction pattern of amorphous Tantalum heat treated 800°C using a 860mm camera length using a FEI T20 TEM. The diffraction pattern shows that the illuminated area is mainly a single crystal.16

Figure 1.4.3.3. Room temperature mechanical loss of 6 Titania doped Tantalum coatings ranging from 0 - 53.7 cation % against cation %.....17

Figure 1.5.1.1. CBED image of a single crystal Silicon sample along the <110> axis using a 860mm camera length using a FEI T20 TEM and recorded on a SIS Megaview III CCD camera. Each bright spot corresponds to diffraction from particular inter-planer spacing.18

Figure 1.5.2.1. Amorphous Structure of Ta₂O₅ showing the same planar fragment of Ta₂O₂ as is in the crystalline phase.....19

Figure 1.5.2.2. TEM dark field image (left) and CBED pattern (right) of a Hafnia coating, heat-treated to 100°C showing partial crystallisation despite low heat-treatment temperature.....20

Figure 1.5.2.3. CBED pattern from a Tantalum coating heat-treated to 600°C showing that the coating has remained amorphous in the illuminated area; no areas of crystallisation were found in this sample.....21

Figure 1.5.2.4. CBED pattern from a Silicon coating heat-treated to 450°C showing that the coating has remained amorphous in the illuminated area; no areas of crystallisation were found in this sample.....21

Figure 2.1.1: Basic design of a thermionic electron gun, current is passed through the filament causing significant heating releasing the electrons. The electrons are then columnated and accelerated down the column.....28

Figure 2.1.2. Design of a basic magnetic lens, a magnetic field is induced by passing current through a series of coils. This magnetic field results in a force which acts on the electrons as they pass through the field adjusting their trajectory. Thus by varying the current in the coils the magnetic field strength can be changed and used to control the path of the electrons.....29

Figure 2.1.5. Schematic diagram of the cross section encapsulation method of sample preparation. Beginning from a 1" disc (a), it is halved (b), a sliver is sliced off (c) leaving a small cuboidal sliver of sample (d). This is then cut in two and the pieces glue together (e), the substrate is then removed (f) and then the whole sample is slotted into a molybdenum rod and fit into a brass tube (g). The tube is then cut into a series of discs (h) each of which is mechanically dimpled (i) finally the sample is thinned to electron transparency using an ion milling system (j).....33

Figure 2.1.6.2 Selected Area Diffraction (SAED) pattern of a silicon crystal sample (left) shown in false colour; notice the oversaturation of the central spots and the distortion in the lefthand side of the image due to the prism. Convergent Beam Electron Diffraction (CBED) pattern of a silicon crystal sample (Right). These two images were taken at the same camera length and on the same crystalline orientation37

Figure 2.1.7.1. Typical EELS spectra showing the zero loss peak at $\sim 0\text{eV}$ and the Plasmon excitations at 20 and 40eV.....38

Figure 2.1.7.2. Typical EELS spectra showing the core loss region, notice the strong edge at $\sim 530\text{eV}$ and the weaker edges at 250eV, 280eV and 400eV.....39

Figure 2.2.1. Typical CBED pattern of Tantalum in false “temperature” colour taken using the SIS camera at a camera length of 860mm.....41

Figure 2.2.2 Intensity Pattern (left) from radial average of a masked CBED image in figure 2.1.1.1. The saturation of the camera occurs at ~ 3600 counts, indicating this image is safely below saturation. The q range has been cut from 1.2-16.5 which represents the edge of the central beam and the edge of the camera.41

Figure 2.2.3 Reduced intensity functions (left) and Reduced density functions (right) for the intensity profile shown in figure 2.2.2. Notice that the undamped Φq shows considerable noise at high Q and that damping factor of 0.4 can significantly reduce this. The periodic oscillations seen in the undamped $G(r)$ (blue) are also significantly reduced by the introduction of the damped factor.45

Figure 2.3.1 200 atom randomly arranged model of Silica (SiO_2) this model would be the starting point for an RMC refinement.47

Figure 2.3.2 200 atom RMC model of Silica (SiO_2) this model is now in experimental alignment with a silica RDF. This model shows considerably more bonding than the randomly arranged model and would serve as the starting point for a molecular dynamics

simulated anneal. The “free” oxygen seen in the model is an artefact from the plotting program and does not imply free oxygen within the structure.48

Figure 3.1.1.1 High contrast false colour diffraction image (left) showing the distortion due to the prism and beam blank. False colour diffraction pattern (right) with masking over the distortion and beam blank both images were taken at a camera length of 860mm.64

Figure 3.1.1.2 False colour CBED image taken at a camera length 600mm, notice the lack of features surrounding the central beam. Blooming has obscured the low spatial frequency information at the centre of the pattern, making it unusable.65

Figure 3.1.2.1. CBED image of Amorphous Silicon exposed for 160s in false colour scheme. This image was processed manually using 160s Gain reference and Black Reference images with the resultant image being multiplied by the maximum intensity of the gain reference (to scale the intensity correctly).....66

Figure 3.1.2.2. CBED image of Amorphous Silicon exposed for 10s in false “temperature” colour scheme. This image was processed automatically by the iTEM image acquisition software. Despite the automatic processing this image is still acceptable as processing induced noise is less prominent at low exposures67

Figure 3.1.2.3. Scatter plot for the image in figure 3.1.2.1 showing the areas considered saturated (green) and the area used calculate the scaling factor (blue).....68

Figure 3.1.2.4. HDR processed CBED image of Amorphous Silicon in false “temperature” colour scheme. This image was created using 7 images with exposure times ranging from 2-160s with the higher exposure time images being processes manually to remove any artefacts.69

Figure 3.1.2.5. Reduced Intensity Functions of amorphous Silicon annealed at 450°C for 4 hours using the Low and high Dynamic Range processes. The LDR reduced intensity function shows considerably worse SNR at the extreme of the q range.70

Figure 3.1.2.6. CBED image of Amorphous Silicon exposed for 160s in false “temperature” colour scheme. This image was processed automatically by the ITem acquisition program. Notice the large amount of noise and artefacts caused by the sub-optimal processing this effect is most noticeable at higher exposures71

Figure 3.1.2.8 Intensity patterns for typical images from the Orius camera (blue) and the SIS (green). The maximum q for the sis at 860mm camera length is $\sim 17\text{\AA}^{-1}$ for the Orius at 8cm camera length the maximum is $\sim 39\text{\AA}^{-1}$. The maximum useable q range for the Orius is $\sim 25\text{\AA}^{-1}$ which is a significant increase over the SIS. The SNR at maximum for the Orius is also significantly better than the SIS.....74

Figure 3.1.3.1: Position of the Tantalum oxygen (left) and tantalum-tantalum RDF peaks as a function of q max. Both plots show convergence at high q max underlining the need for high q data sets when calculating G(r).....75

Figure 3.1.3.2. Series of RDFs from the same image produced by varying q max; the RDFs appear to converge as q max is increased. The shorter distance information is affected more by the decreasing q max, this is because q is a reciprocal distance and as such high spatial frequency data is cut off first.76

Figure 3.1.3.3. Off centre CBED image taken at 860mm camera length, there is a distortion on the left hand side of the image which was later masked to remove its affects.....77

Figure 3.2.1 A Comparison between electron and X-ray diffraction data, the q range shown represents the maximum range attainable using an energy of 8KeV on beamline b16 at Diamond Lightsource. Note the splitting of peak at $4\text{-}5\text{\AA}^{-1}$ in the X-ray data and the enhanced peak at $6\text{-}8\text{\AA}^{-1}$, sadly the 2\AA^{-1} peak is almost complete obscured by noise.80

Figure 3.3.1. Example of RDF calculated from low quality data, this represents the best fit which could be achieved to the intensity which results in a negative noise peak below 1Å81

Figure 3.3.1.1 RDFs from an HDR image (blue) and an LDR image (green) showing the significantly improved sub angstrom noise peak for the HDR image.....84

Figure 3.3.1.2. A series of RDFs from HDR and LDR images of amorphous Ti:Ta₂O₅ showing the spread due to noise. Notice that the peak at approximately 1.1Å is resolved in the HDR images considerably more than in the LDR images.86

Figure 3.3.3.1 Area map of diffraction study sites for the heat treated (left) and as deposited (right) samples. The red markers indicate the areas in which diffraction data was taken. The Substrate for both samples is crystalline Silicon and marked as SUB in the images.90

Figure 3.3.2.2. Heights (top), FWHM (middle) and positions (bottom) of the 1st and 2nd RDF peaks for amorphous Tantalum heat treated to 600°C as a function of depth of the nanovolume. Both the heights and positions show clear trends with depth indicating that the peaks drop in height and move to short distances toward the interface. The trend in the peak width is less clear but it seems that the 1st peak becomes narrower and the 2nd peak broader with depth.93

Figure 3.4.1 Density measurements for the series Ti:Ta₂O₅ coatings discussed in chapter 4. The density for the 6.5% doped coatings was measurements using XRR. Values for the 0 and 20% doped coatings were measured using alternative methods. The remaining values were linearly interpolated from the measured values as was a value for pure amorphous Titania.....97

Figure 3.4.2. Raw XRR data for coatings 1, 2, 4 and 5 showing the critical angles of silica and Ti:Ta₂O₅ layers at ~0.2° and 0.32° respectively.98

Figure 3.4.3. XRR data for coatings 1, 2, 4 and 5 rescaled such that the critical angle for silica is fixed.....	99
Figure 3.5.1. Total energy of a 100 atom, Tantalum structure as a function of the kinetic energy cut off.	100
Figure 3.5.2. A comparison of RDFs from experimental (green), un-optimised simulation (blue) and optimised simulation (red) sources. The optimised DFT RDF shows a closer agreement to the experimental RDF especially in the 2 nd peak at ~ 3.5Å.....	102
Figure 3.5.3. Energy evolution for two models with the same starting configuration. The optimised model (blue) used a cut off energy of 600eV and 8 k-points, the un-optimised model used a cut off energy of 200eV and 1 k-point.	103
Figure 3.5.1.1 Total energy evolution of two separate simulations which differ only the amount of time the system is allowed to find achieve equilibrium at each temperature. The 'slow' simulation (green) was allowed 5ps (2500 iterations) at each stage, the 'fast' simulation (blue) was allowed 2ps (1000 iterations) at each stage.	106
Figure 3.5.1.2 Total energy evolution of two simulations with a different amount of stages in the SN process. The fast simulation (blue) from figure 3.5.1.1 uses the 8 stage SN explained early where as the other simulation (green) uses a 2 stage SN.....	107
Figure 3.5.2.1 Energy evolution of a geometry optimisation (green) and molecular dynamics simulation (blue). Both simulations used the same input parameters and input structure of amorphous Tantalum.	109
Figure 3.5.3.1 Energy evolution of 100 atom amorphous Tantalum structure where the final MD structure has been optimised. The time axis for the geometry optimisation is arbitrary, since the each successive iterative is an atomic step and not an time step as is the case with the MD.	110

Figure 3.5.4.1 RDFs of 25.7 Ti cat % doped Ta₂O₅ coating from the experiment (blue) the RMC generated structure (green) and the structure after the DFT MD simulation.....112

Figure 3.5.4.2 The effect of an increased number of RMC steps on the fit between the theoretical and experimental RDFs.....113

Figure 3.5.4.2 The result of a geometry optimisation performed on a structure refit to the experimental RDF after a simulated anneal simulation.115

Figure 4.1.1: Diffraction patterns for Ti:Ta₂O₅ ranging from 0% (top left) to 53.7% (bottom right) doping. These images were taken using a standard set of microscope settings, 860mm camera length, $\alpha = 1.19 \pm 0.05$ mrad and spot size 5.....120

Figure 4.1.2 RDFs from the 6 Ti:Ta₂O₅ coatings, the Y axes are dimensionless and have an arbitrary scale and thus are not shown, however each plot has same Y axis.122

Figure 4.1.3a. Peak Positions for the 6 LMA Titania Doped Tantalum coatings against the Ti cat. %.....123

Figure 4.1.3b. Peak Heights for the 6 LMA Titania Doped Tantalum coatings against the Ti cat. %.....123

Figure 4.1.3c. Peak Widths for the 6 LMA Titania Doped Tantalum coatings against the Ti cat. %.....124

Figure 4.1.4. Room Temperature mechanical loss for the 6 LMA coatings as a function of Ti cat. %. Also plotted is the FWHM for the 1st peak ($\sim 2\text{\AA}$) as a function of the Ti Cat. %. The absolute values of these quantities are not related and as such they are shown on arbitrary Y-axes.125

Figure 4.2.1.1 Calculated cohesive energy of the LMA coatings against the full width at half maximum of the 1 st peak in their respective RDFs. The values for the Y-axes are arbitrary and unrelated and as such have been removed.	128
Figure 4.2.1.2 Experimental (blue) and theoretical (green) RDFs for the LMA coatings, The theoretical RDFs are calculated for the structure upon completion of the DFT stage.....	130
Figure 4.2.1.4. Ratio of abundances of Ta-O to Ti-O nearest neighbours against Ti Cat % as measured as the heights of the peaks in the figure 4.2.1.3.	132
Figure 4.2.1.5. Nearest neighbour distances of the Ta-O and Ti-O nearest neighbours against Ti Cat. %	133
Figure 4.2.2.1. 3 variants of the 4 atom planar fragments found in Tantalum and Titania doped Tantalum models after DFT stage.	135
Figure 4.2.2.2. Ratio of Ta...Ta to Ti...Ta distances against Ti Cat. % calculated from Gaussian parameters. Also shown is the room temperature mechanical loss on an appropriate scale.	136
Figure 4.2.3.1 Partial RDFs for the LMA series of coatings after the final refinement stage of model, each model here had ~ 2700 atoms. All 6 plots share the same Y axis scale however the absolute values are arbitrary and so have been omitted. The X-axes have the limits 0 – 5 Å and have a spacing of 1Å.....	140
Figure 4.2.3.2. The ratio of Ta...Ta distances summed with Ti...Ti distance against Ti...Ta Distances. This is an indication of the ratio between the pure and doped planar fragment variants. Shown are four functions, each has a different scaling of the TiTi contributions, using A = 1, 1.9, 1.9 ² and 1.9 ³ . Shown alongside for comparison is the room temperature mechanical loss of Titania Doped Tantalum.	144

Figure 4.2.4.1 RDFs from generated from the experimental data (red), the models which used Unoptimised parameters (green) and optimised parameters (blue). Notice that for the most part, the RDFs from optimised models more closely match the experimental than the unoptimised models. All 5 plots share the same Y axis scale however the absolute values are arbitrary and so have been omitted. The X-axes have the limits 0 – 5 Å and have a spacing of 1Å.147

Figure 4.2.4.2 Partial RDFs for the completed “optimised” models for the LMA coating series. The model for the 8.26% model is missing as it was corrupt and unusable. All 5 plots share the same Y axis scale however the absolute values are arbitrary and so have been omitted. The X-axes have the limits 0 – 5 Å and have a spacing of 1Å.149

Figure 5.1.1. Diffraction patterns for CSIRO 25% Titania doped Tantalum in increasing annealing temperature left to right154

Figure 5.1.2 Diffraction patterns for CSIRO 55% Titania doped Tantalum in increasing annealing temperature left to right155

Figure 5.1.1.3 Experimentally determined RDFs for the four CSIRO coatings doped at 25%, as can be seen, only slight changes can be seen between the different heat treatment temperatures.156

Figure 5.1.1.4 Experimentally determined RDFs for the four CSIRO coatings doped at 55%, as can be seen, again, only slight changes can be seen between the different heat treatment temperatures.157

Figure 5.1.3. Heights of the first and second RDF peaks for the 25 (bottom) and 55% (top) doping CSIRO coatings as a function of heat treatment temperature159

Figure 5.1.4. Positions of the first and second RDF peaks for the 25 and 55% doping CSIRO coatings as a function of heat treatment temperature162

Figure 5.1.5. Heights of the first and second RDF peaks for the 25 and 55% doping CSIRO coatings as a function of heat treatment temperature163

Figure 5.2.1 The Initial energies and energy changes due to due to the simulated anneal for the 25% doped models.....165

Figure 5.2.2. Average cohesive energies for the 4 25% doped CSIRO coatings as a function of heat treatment temperature.166

Acknowledgements

First and foremost I would like to convey my sincerely gratitude to my supervisors Dr Ian Maclaren and Prof. Sheila Rowan. Their guidance, support and teachings have been of tremendous help during this investigation. Thanks also go to Prof. Jim Hough for his input and support during my Ph.D. My thanks go to Dr Riccardo Bassiri for his continued help during my Ph.D, for his help acquiring the experimental data this Thesis is based on and for our many spirited conversations regarding the direction and focus of our research. A special thank you to Dr Iain Martin for his continuous help, input, guidance aswell as may hours of stimulating discussion of the subject. I would also like to particularly thank Drs Damian McGrouther, Stephen McVitie and Giles Hammond for the training I received initially and throughout my Ph.D, I found it invaluable. I would also like to thank Dr Richard Beanlean of the University of Warwick for this input into this project and his patience whilst working with me.

Thanks go to Dr Sam McFadzean and Colin How for their monumental patience and understand when aiding and instructing me in use of the experimental equipment aswell as their maintenance of said. A special thank you to Brian Miller for all the samples he tirelessly prepared for me, may he rest in peace. Thanks go to Dr Konstantin Borisenko, Yixen Chen and the late Prof David Cockayne of Oxford University for the expertise and guidance they lent to my Ph.D. Thanks go to Dr Stuart Reid and Dr Ronny Nawrodt for their instruction and my education in thermal noise and mechanical loss.

Heartfelt thanks to my office mates Dr Matthew Abernathy, Dr Karen Haughian, Dr Alan Cumming, Dr Peter Murray, Dr Nicola Beverage, Chris Bell, Kieran Craig, Paul Campsie and Rahul Kumar for countless productive discussions and for making my office a happy place to work and my Ph.D enjoyable. Thanks to Jean Grieg and Jamie Scott for their administrative and IT support during my time in Glasgow.

From the University of Florida I would like to start by thanking Prof Hai-Ping Cheng for her hospitality accepting me into her group for the short time I worked with her. To all

the Quantum Theory Project in Florida I would like to say thank you for immediate acceptance of me and willingness to include me in your research while I visited UFL. Furthermore I'd like to thank Yuning Wu of UFL for all his help and patience educating me on the VASP software package and in DFT methods.

Finally I would like to thank all the member of the Institute for Gravitational Research, Materials and Condensed Matter Physics Group and Quantum Theory Project for all that I have learnt from them during my Ph.D.

One final thank you to my parents Catherine Glynis and David Alan Evans, my sister Elizabeth Leigh Hornsey and my niece Libby Rae Hornsey for their unwavering support throughout my education and research, this thesis is dedicate to them. A Thank you also, to all my friends and to my love Naziyah Noor Mahmood for their love, support and for listening to my endless complaints.

Preface

This thesis is an account and record of research carried out at the Institute for Gravitational Research and Materials and Condensed Matter Physics at the University of Glasgow as well as the Quantum Theory Project at the University of Florida in Gainesville, Florida. The research took place between October 2009 and November 2012 and is based on computational modelling of amorphous mirror coatings for use in advanced gravitational wave detectors.

Chapter 1 is summary and concise description of gravitational waves, gravitational wave detectors, their limitations and the work underway to understand and surpass them.

Chapter 2 is a brief but detailed description of the methods employed during this Ph.D including (but not limited to) Transmission Electron Microscopy, Density Function Theory and Reduced Density Functions.

In chapter 3 various adaptations to the method of semi-empirical modelling outlined and developed by Prof David Cockayne, Dr Konstantin Borisenko and Dr Yixin Chen of Oxford University are discussed and their relevance to this thesis. All electron diffraction data is courtesy of Dr Riccardo Bassiri except for that contained under the head High Dynamic Range (HDR) data. The preliminary HDR data was taken by the author and is explained in section 3.1.2, HDR data from the CSIRO coatings (also shown in chapter 3) was acquired by Mr Martin Hart and Miss Kirsty Annand as part of their Master Degree Research Project. Further HDR data was acquired using the JEOL 2100 microscopy at the University of Warwick on an Gatan Orius CCD Camera with aid of Dr Richard Beanland. X-ray diffraction experiments were carried out with Dr Ian Maclaren and Dr Riccardo Bassiri at the Diamond Light Source facility as part of the Rutherford Appleton Laboratory in Didcot, Oxfordshire. X-ray Reflectometry experimental were carried out in the Materials department at the Leeds University in collaboration with Dr Timothy Comyn and under the supervision of Dr Ian Maclaren. Analysis was performed by the author,

under the supervision Dr Ian Maclaren with continuous input from Dr Riccardo Bassiri, Dr Iain Martin and Prof Sheila Rowan. Reverse Monte Carlo software was provided and developed by Dr Konstantin Borisenko as were the initial input for the DFT models.

Mechanical loss measurements for the LMA coatings discussed in chapter 4 were taken and analysed by Prof. Gregg Harry and Dr Matthew Abernathy. The electron diffraction studies for these were performed by Dr Riccardo Bassiri and the analysis by the Author. The RDF statistical analysis for the LMA coatings was inspired by Dr Iain Martin and performed by the author. The modelling of these coatings was performed jointly between the author responsible for DFT calculations and Dr Riccardo Bassiri responsible for the Reverse Monte Carlo simulations. The electron microscopy samples were prepared by the late Brian Miller.

In chapter 5 the mechanical loss (undocumented) was measured, analysed and characterised by Dr Matthew Abernathy and Mr Kieran Craig. The electron diffraction studies were performed by Mr Martin Hart, Miss Kirsty Annand and Dr Ian Maclaren and the analysis by the author. The modelling at both the Reverse monte carlo and Density function theory stage were performed by the author as was the analysis.

Chapter 1: Introduction to gravitational waves, advanced interferometry and Thermal noise

1.0 Chapter Introduction

In this chapter we will define what Gravitational Waves (GWs) are; this will be done without venturing deeply into General Relativity (GR) or using significant derivations. Instead GWs will be defined and described in more qualitative terms and emphasis will be placed on their uses. Since the form and magnitude of GW signals are highly dependent on the type of signal source, these will also be discussed to provide adequate context for the use of GWs. GW detectors are likely to be a vital component of future studies using multi-messenger astronomy. This Chapter will contain a concise outline of the need to develop effective methods for detecting GWs and hence the need to overcome limiting noise in current and future GW detectors. The methods for GW detection both past and present will be outlined and discussed both in the context of the physics and process of detector development. Detector development will be discussed in terms of the noise sources which limit current and future detectors, this will progress onto a clear definition of the aims of this thesis. Finally the progress towards this aim, made prior to this thesis, will be presented which will clarify the motivation for this work and that of future work.

1.1 Gravitational Waves

1.1.1 What are Gravitational Waves?

In nature there are four fundamental forces in nature, electromagnetism (EM) [1], the strong and [2], weak nuclear forces [3] and gravity [4]. The former three can already be defined in terms of waves i.e. EM waves for electrostatics etc. They also have an associated force carrying particle, i.e. photons for EM, *gluon* particles for the strong nuclear force [5, 6] and (amongst others) the *Z* particles for the weak force [7]. These

have all been predicted theoretically and confirmed experimentally, however direct detection of gravitation waves and their force carrying particles remains elusive. It is assumed that the same state exists for gravity as for the other three forces and that it too can be defined with a wave [8] and a particle [9]. Albert Einstein showed that a wave solution to his General Theory of Relativity [8] exists and that it must propagate at the speed of light. This wave is of course known as a Gravitational Wave and by solving the associated wave equation it can be shown that this wave has two quadrupolar polarisations, 'plus' and 'cross' [8].

A GW can be described conceptually in geometric terms [8] i.e. in terms of the effect it has on space itself. GR tell us that space and time are not constant [8] and that it is instead, the space-time interval which remains constant [10]. So as a GW propagates it stretches and compresses space in order to essentially "update" the gravitational gradient of that area of the universe this can be seen diagrammatically in Figure 1.1.3.1. In essence whenever mass moves non-axisymmetrically [8], it changes the gravitational gradient of the Universe. This information is transmitted by means of a propagating wave, altering the gravitational gradient of universe. To clarify, only non-axisymmetric changes in mass distribution create GWs, because axially-symmetric masses do not have a time dependent spatial distribution of their own mass. For example a spinning, axisymmetric neutron star has the same mass distribution one second as it does the next. However a neutron star with a small mountain at the equator has a time-varying spatial distribution of its mass. As such the symmetric neutron star does not emit GWs, whereas the non-symmetric one does.

1.1.2 Why are they important?

In the field of astronomical observation it is a rare occurrence when less data is preferable over more data. When a major astronomical event occurs, such as a supernova, astronomers seek to collect as much data as possible from a variety of sources. This is referred to as 'multi-messenger' astronomy and as one would expect, it

provides greater information than astronomy using a single observational channel. The concept is simple; if for example an optical telescope witnesses a supernova in its initial stages then astronomers will move to “point” as many other telescopes at that area of space. This of course means that instead of simply recording the visible part of the supernova, the X-ray emission/absorption spectra and the radio wave emission spectra etc, can be recorded. Combining the data from all these sources will then provide a much clearer picture of the supernova leading to a greater understanding of it.

Eventually, once GW detectors and the associated analytical methods are able to detect and categorise GWs regularly and efficiently, then GW detection will become an integral part of multi-messenger astronomy. The advantages of GW detection for multi-messenger astronomy is that GWs are not absorbed by dust, gas or other obstacles between the source and the detector, a problem that does occur with many EM waves. In addition, observed EM radiation typically has been emitted from the exterior of an object, leaving observers to infer the internal dynamics of systems. In contrast GWs are produced purely by the acceleration of mass, thus they contain information resulting directly from the mass-dynamics of systems, including those which are electromagnetically dark. Their observation should thus enable us to obtain information about sources which is unobtainable by any other means.

So what are we likely to “see” with GW detectors? To answer this question we need to define the types of sources of GWs. GW sources are grouped together by their duration and form and are classified either *burst*, *continuous* or *stochastic* accordingly. *Burst* sources such as supernovae, are onetime events creating (relatively) short, non-repeating signals through which information such as initial mass can be determined. Continuous events create repeating signals with predictable periods, these sources include aspherical neutron stars; from these signals the total mass of the system and degree of asphericity can be determined. Stochastic sources are repeating signals with an inherently unpredictable period, these signals include the stochastic background, a

remnant of the big bang which could lead to the greater understanding of the moments proceeding the big bang.

1.1.3 How do we detect Gravitational Waves?

So we know what GWs are, where they come from and what we can learn from them but how do we detect them? To answer this question we have to think about how GWs change space-time as they propagate. Since it is known that space and time are not constant we can use this property to detect them. As GWs propagate they change the length scale of space and they do so at a finite velocity. The first detectors [11] took the form of large, heavy aluminium bars suspended by thin (low mass) wires. The theory behind this is that as a GW propagates through the bar, it would briefly expand/compress as the GW stretched and compressed space (figure 1.1.3.1) [12].




Figure 1.1.3.1: Timeline diagram of the space surrounding a bar-type GW detector as a GW propagates through it. The space (and the bar) expands in one direction, returns to its equilibrium position and then expands in the orthogonal direction before returning to its equilibrium position. The magnitude of a gravitational wave 'h' is defined as $h = \frac{2\Delta L}{L}$ where L is the initial distance between two points in space and ΔL is the difference in distance induced by the passage of gravitational wave of magnitude h. *With the Weber Bar having an approximate sensitivity of 10^{-16} .*

A famous example of this was the ‘Weber’ bar detector built in the 1960s [11], which had several detections [13, 14], which were later found to be unverifiable [15-18].

Bar-type detectors have the disadvantage of having a narrow frequency range to which they are sensitive. Interferometric Gravitational Wave Detectors (IGWDs) [19] (discussed in more detail in section 1.2) have a much broader sensitive frequency range [20], limited by a range of different noise sources (which are outlined in section 1.3).

1.2 Interferometric Gravitational Wave (IGW) detectors

Although the first GW detectors were bar-type, today most active GW detectors are IGWDs based on Michelson-type interferometers [21]. A Michelson-type laser interferometer [22] (Figure 1.2.1) uses a beamsplitter to split a laser into separate beams; these are then sent down different paths, typically two “arms” in orthogonal directions. The laser beams are then re-combined, causing interference due to the phase difference in the separate beams. The phase shift between two beams is due to the path length difference between the two arms.

IGWDs are based on Michelson interferometers but include highly advanced additions to decrease noise and increase sensitivity. The theory of IGWDs is that as a GW propagates through the detector, it will cause extension and contraction of the lengths of the arms in each orthogonal direction in turn (Figure 1.2.2). Before detection the interferometer is setup such that the arm lengths have a π phase difference resulting in an interference pattern at the detector centred on a dark fringe, making it sensitive to increases in intensity. A GW will result in a change in path length this will in turn alter the observed intensity of interference pattern.




Figure 1.2.1. Diagram of a basic Michelson-type laser interferometer showing the laser (a), beam splitter (b) the reflecting mirrors (c1 and c2) the arms (d1 and d2) and the detector (e)




Figure 1.2.2. A diagram similar to that of figure 1.1.3.1. showing that expansion and contraction of space containing a IGWD to show that the two types of detection work on the space principle.

Each source of GWs mentioned above (and further types not included here for brevity) will have a distinct signal pattern indicative of its method of emission. So by

appropriate modelling of the signal from known astrophysics it is possible to determine the nature of the object(s) which caused the GW signal. This allows for quite detailed analysis to be performed on the signal and from that a great deal can be learnt about the source, which is one of the key motivations for detection of GWs.

One of the most advanced IGW detectors in the world [23] is known as the Laser Interferometry Gravitational wave Observatory (LIGO) [24]. Other detectors include GEO 600 in Germany [25], Virgo (currently under upgrade to Advanced Virgo) in Italy [26] and Tama in Japan [27]. LIGO is a system of two detectors based in the US, one in Hanford, WA and the second in Livingstone, LA [28]. Each detector is made up of two perpendicular 4km long arms using 10.7kg silica test masses [24]. LIGO is known as a 1st generation detector, it is currently undergoing an upgrade to Advanced LIGO (aLIGO), a 2nd generation detector designed to improve its sensitivity [29]. Both LIGO and aLIGO are designed to have their highest sensitivity in the range from several hundred Hz to several kHz [30]. This allows aLIGO to witness in spirals from black hole and neutron star binaries however missing at the coalescence of said due to the high frequency nature of said. aLIGO is also incapable of detecting low frequency events such as big bang remnants. aLIGO incorporates some highly advanced optical equipment designed to increase efficiency and reduced noise in the detector [30]. The test mirrors are coated in multilayer coatings of alternating high (Ti:Ta₂O₅) and low (SiO₂) refractive index materials [31]. This maximises the reflectivity of the mirror in the desired laser frequency range as described by the Fresnel Conditions eq.1.2.1 [32].

$$R_s = \left(\frac{n_1 \cos \theta_i - n_t \cos \theta_t}{n_i \cos \theta_i + n_t \cos \theta_t} \right)^2$$

[eq.1.2.1]

Where n_1 , n_2 , θ_i and θ_t are the reflective indices at the frequency of the laser for the two layers and the angle of the incident and transmitted beams normal to the surface respectively.

Each layer is $\frac{1}{4}$ wavelengths in thickness [30] so as to maximise the reflected wave. The test mirrors are also 40kg compared to the 10.7kg of LIGO [30], the extra mass helps to reduce radiation pressure noise (see section 1.3.3.). These test mirrors will be hung on pendulum suspensions having four stages ‘quadruple pendulums’ to improve seismic isolation and thus reduce the effects of seismic noise on detector sensitivity (section 1.3.1). Fabry-Perot cavities [33], are implemented in each arm help to increase the effective laser power which reduces the effect of photon shot noise [34] (section 1.3.3). These improvements along with many others [35] are designed to increase the sensitivity of LIGO tenfold [35]. The improved sensitivity increase the effective range of the detector by 10 (by being able to detect weaker signals) and thus increases the volume of space within the detectable range of aLIGO by ~ 1000 [30]. This will significantly increase the GW detection rate of LIGO since more signals will be above the limiting noise level and this will hopefully lead to detection. aLIGO will of course have its own sensitivity limiting noise sources which will need to be overcome in the future [35].

1.3 Sensitivity limits

Since gravitational wave signal strength is so small (detectable strain of approximately $h \sim 10^{-21}$ [36]) detectors must be incredibly sensitive in order to detect them. IGWDs are fundamentally susceptible to noise sources that would not normally be considered in most areas of physics. For example, US IGWDs are subjected to a repeating noise source due to waves crashing into the western shore of the United States hundreds of miles away [37]. This is an example of seismic noise [38], another source is gravitational gradient noise [38] caused by the change in local gravity which occurs when massive objects move, such as clouds floating over the detectors or tectonic shifts. There is also quantum noise, caused by the laser itself in the form of photon shot [39] and

radiation pressure noise [10]. The most important noise source for this thesis is thermal noise, specifically coating thermal noise [40].

1.3.1 Seismic Noise

When detection of a signal depends on the distance between 2 points it seems reasonable to accept that any event which is capable of moving those points would cause unwanted signal. This is the case with seismic noise in IGW detectors; seismic noise is simply any event that causes reverberations in the detector resulting in spurious signal. Like all sources of noise relating to GW detection the problem is compounded by the fact that GWs cause very small length changes in aLIGO. This means that IGWs must be incredibly sensitive and as such are incredibly sensitive to very weak noise sources. When one thinks of 'seismic' one might be mistakenly think about earthquakes, however IGWs like aLIGO are sensitive to footsteps, cars passing miles away and tides hitting the coast 100s of miles away. Because of this, IGWs have had to push the field of seismic isolation to new heights to achieve the required sensitivity. In aLIGO each 'test mirror' is suspended at the end of a quad suspension chain (figure 1.3.1.1) [41], i.e. a pendulum on a pendulum on a pendulum on a pendulum. It has been proposed that future gravitational wave detectors could be built underground further reducing seismic noise [42].




Figure 1.3.1.1. Diagram (not to scale) of the quad suspension similar to the design of the aLIGO test masses, from the front (left) and the side (right). The final suspension layer is quasi-monolithic i.e. that the mass and the wires are one continuous piece of fused silica.

1.3.2 Gravitational Gradient Noise

The detection of GWs using IGW detectors relies on sensing the relative displacements of the test mass mirrors induced by GWs. So any force which can pull a mirror can cause spurious signal at the detector. This is the principle of gravitational gradient noise; motion of local objects can change local gravity and as such change the direction of the gravitational attraction on the mirror, One example of the causes of gravitational gradient (GG) noise at low frequency is ground movement. The gravitational pull on the test masses induced by seismic motion (as distinct from direct seismic noise) constantly changes the local gravitational gradient. To combat gravitational gradient noise the IGWs can be placed underground which protects them from some of the seismic shifts in local density responsible for GG noise. GG noise is not a significant problem for aLIGO but has played a major part in the design of 3rd generation detectors are they will be sensitive enough to be limited by GG noise [38].

1.3.3 Quantum Noise

Quantum noise is a combination of different phenomena regarding the behaviour of photons, these are photo-electron shot noise and radiation pressure. Shot noise is due to effects related to photon detection at the photodiode at the interferometer output. It is assumed that photons obey Poisson statistics and as such the average magnitude of these statistical fluctuations in detected photon number is \sqrt{N} where N is the mean number of photons [43]. The limit to measurable differential displacement, Δx of an IGWD mirror due to photon shot noise is described by eq.1.3.3.1 [44]. Where P and λ are the laser's power and wavelength, Δf the bandwidth, \hbar and c are the reduced Planck constant and speed of light and Φ is the phase difference between the two arms.

$$\Delta x = \left(\frac{\hbar c \lambda \Delta f}{4\pi P \cos^2(\Phi/2)} \right)^{\frac{1}{2}}$$

[eq.1.3.3.1]

Thus it can be seen that the effects of photoelectron shot noise can be reduced, either by increasing P or forcing $\cos^2(\Phi/2)$ to 1 i.e. Φ to zero which is physically represented by tuning the IGWD to operate on a dark fringe. The second source of quantum noise is radiation pressure noise. Despite being massless, photons still have momentum. This comes from a rearrangement of mass-energy equivalence equation (eq.1.3.3.2) which states that the momentum of a photon is equal to its energy divided by the speed of light [45].

$$\frac{E}{c} = mc$$

[eq.1.3.3.2]

Since the photon has momentum, it also means that it can transfer this momentum to the suspended optic in the interferometer, this being described as radiation pressure (RP). RP noise is thus due to photons transferring momentum to the test mirrors which causes them to move slightly; this of course changes the arm length causing spurious signal. Differential displacement due to RP noise is described by eq.1.3.3.3 where m is the mass of the mirror, ω is the detection angular frequency and all other variables are as defined for eq.1.3.3.1. It can be seen that, eq.1.3.3.1 and eq.1.3.3.3 scale in opposite directions with changing laser power.

$$\Delta x = \left(\frac{16\pi\hbar P \Delta f}{\lambda m^2 \omega^4 c} \right)^{\frac{1}{2}}$$

[eq.1.3.3.3]

Since both Radiation Pressure and Photon Shot noise scale oppositely to the same variable (laser power) then there is a fundamental limit to how much both can be reduced. This is called the standard quantum limit and is unavoidable, although it can be circumvented [46-48], how this is done will not be discussed in this thesis.

1.3.4 Thermal Noise and Mechanical Loss

The final noise source to be discussed is thermal noise, the reduction of which is the underlying objective of this thesis. Thermal noise has two components relevant to aLIGO and future detectors [49] these are Brownian and thermoelastic noise. Brownian noise [50] is due to the finite temperature of the atoms being manifested in their random thermal motion. [51]. Thermoelastic noise is due to inhomogeneous expansion of the coating caused by statistical fluctuations in local temperature [52]. It should be mentioned at this point that for aLIGO, thermal noise is a particularly limiting source of noise from the mirror coatings [49] and it is these coatings which are investigated here.

Thermal noise can in principle be reduced by lowering the temperature of the mirrors, however care must be taken as for some materials, such as fused silica, the mechanical loss of the system actually increases with decreasing temperature, with a broad dissipation peak centred around 40K. aLIGO will be operated at room temperature however cooling remains an option for future detectors/upgrades [29, 53] Lowering the temperature poses certain problems; firstly it is expensive to pump liquid nitrogen or helium around the mirrors to cool them. Furthermore the laser will heat mirror, or at least the area it is in contact with, which will cause a temperature gradient across the mirror causing noise. Another option is to use a coating with lower intrinsic thermal noise, the search for low noise coatings is currently an important area of research in the GW community [54].

1.4 Mechanical Loss

1.4.1 Measuring Mechanical Loss

Measuring thermal noise directly at the frequencies of interest for gravitational wave detection is difficult; however it is related to mechanical loss (eg. of a coated sample) by eq.1.4.1 which is easier to measure. Coating mechanical loss can be measured using a process of exciting vibrational resonances in coated cantilevers [55], disks [56] and bulk rods [56]. The resonant modes are excited electrostatically using an exciter plate [56] driven by an AC supply, the frequency of which is used to excited modes. By modelling the rate of decay of the amplitude of the resonance the mechanical loss of the coated sample can be determined [57]. The mechanical loss of the coating can then be extracted through comparison with an uncoated “witness” sample. Whilst this allows for quantification of mechanical loss for a specific coating, it doesn’t directly determine the mechanisms causing the loss. It is thus necessary to determine the fundamental causes of mechanical loss so that more advanced, lower loss coatings can be developed.

$$S_x(f) = \frac{2k_B T d}{\pi^2 f Y w_0^2} \left(\frac{Y'}{Y} \Phi_{\parallel} + \frac{Y}{Y'} \Phi_{\perp} \right)$$

[eq.1.4.1]

Where $S_x(f)$ is the thermal noise spectral density, k_B is the Boltzmann constant, T and d are the temperature and thickness of the coating respectively. Y and Y' are the Young's moduli of the coating and substrate respectively and Φ_{\parallel} and Φ_{\perp} are the components of the mechanical loss parallel and perpendicular to the coated surface. Finally f is the detection frequency in the IGWD and w_0 is the width of the laser beam.

1.4.2 Temperature Dependent Mechanical Loss

Mechanical loss is not typically temperature dependent; work at the University of Glasgow has shown that coating mechanical loss can change considerably with temperature [58, 59]. Room temperature mechanical loss measurements are essential to aid understanding of the anticipated sensitivities of current detectors and 2nd generation detectors. In addition, the form of the temperature dependence of coating loss can be instructive in identifying the mechanisms responsible for the magnitude of the loss both at room and low temperatures. It has been proposed upgrades to the advanced LIGO detectors and the 3rd generation detector known as Einstein Telescope (ET) will be operated at cryogenic temperatures to reduce thermal noise [42, 60]. For this reason current and proposed future coatings are tested over a range of temperatures down to ~10K [61].

1.4.3 The Importance of Tantalum and Titania Doped Tantalum

As mentioned the reflectivity of the test mirrors is maximised by coating them with a multilayer coating of alternating low and high refractive index coatings. aLIGO is designed to use amorphous Silica (SiO_2) and Tantalum (Ta_2O_5) for the low and high index

materials respectively [31]. Tantalum has a refractive index of 2.27 at 1064nm, [55] the wavelength of laser which aLIGO is designed to use [30]. Previous work on Tantalum by Martin *et al* [58] has shown that the mechanical loss of Tantalum is temperature dependent (figure 1.4.3.1). At low temperature $< 20\text{K}$ a peak in the loss is seen in Tantalum which has been heat treated to 300°C for 24 hours figure 1.4.3.1 [58]. It has also been determined that this low temperature peak is not constant; it changes with heat treatment as seen in figure 1.4.3.1.




Figure 1.4.3.1. Mechanical loss of amorphous Tantalum heat treated at 300°C , 400°C , 600°C and 800°C as a function of temperature.

Heat-treating Tantalum to 400°C , 600°C and 800°C increases the low temperature mechanical loss quite considerably compared to coatings heat treated to 300°C . It should be mentioned at this point that heat-treating Tantalum to 800°C leads to crystallisation as can be seen in figure 1.4.3.2. This increase in low temperature loss could of course be avoided by not heat-treating; however it has also been shown that heat treatment improves

the optical properties [62]. Since these coatings are intended to be used in an optical system; it is of utmost importance to produce coatings of the highest optical quality and as such heat treatment is a necessity. This leads to the need to lower the mechanical loss of Tantalum by means other than avoiding heat treatment.




Figure 1.4.3.2: Convergent beam electron diffraction pattern of amorphous Tantalum heat treated 800°C using a 860mm camera length using a FEI T20 TEM. The diffraction pattern shows that the illuminated area is mainly a single crystal.




Figure 1.4.3.3. Room temperature mechanical loss of 6 Titania doped Tantalum coatings ranging from 0 - 53.7 cation % against cation %.

Work by Harry *et al* [63] has shown that the room temperature mechanical loss of Tantalum can be reduced by doping with Titania. A series of 6 different dopant levels were tested and it was found that the lowest loss occurs at a dopant level of 30% Titania [63] figure (1.4.3.3). This series of coatings will be discussed at length in this thesis as it represents the largest single dataset. How Titania doping affects mechanical loss in Tantalum at low temperature is an area of current research in the Institute for Gravitational Research but is as yet unpublished and will not be discussed.

1.5 Coatings and modelling

Since all the materials investigated in this thesis are amorphous, it is necessary to define what an amorphous material is. Before this can be done it is however, necessary to discuss the main characteristics of crystalline structures. These will be discussed in terms

of their electron diffraction patterns and the characteristics of their atomic structure. Modelling will not be discussed in detail, except for the need for modelling; detailed methods will be discussed in Chapter 2.

1.5.1 Crystalline Materials

Crystalline materials have well defined short and long-range atomic order and are describable by the periodic tessellation of some arrangement of atoms called a unit cell. These unit cells will typically be describable by some simple shape, such as a cube or hexagon with a near constant set of dimensions called lattice parameters. Since the structure is a simple tessellation of the unit cell, planes exist which intersect with identical atomic positions in neighbouring cells. With crystalline structures thicker than a single layer, repetitions of planes occur and thus set up Bragg conditions as defined by $n\lambda = d \sin \theta$. This creates a discrete series of diffraction angles explored by incident electrons and hence a crystalline diffraction pattern appears as a series of spots (figure 1.5.1.1) each corresponding to a unique inter-planar spacing.




Figure 1.5.1.1. CBED image of a single crystal Silicon sample along the $\langle 110 \rangle$ axis using a 860mm camera length using a FEI T20 TEM and recorded on a SIS Megaview III CCD camera. Each bright spot corresponds to diffraction from particular inter-planer spacing.

1.5.2 Amorphous Materials

To define an amorphous material we should start by defining the word amorphous which means “without form”. This is perhaps not an accurate description of an amorphous material which has form (structure) which simply does not extend beyond a few interatomic distances. Were amorphous structures devoid of organised structure (like low pressure gases) they would not exhibit any properties of interest and thereby would not be of interest [64]. In reality an amorphous structure is highly ordered at the short range and shares some of the structural features seen in their crystalline counter parts such as the planar fragment seen in figures 1.5.2.1. These fragments and their role in the structure will be discussed at length in chapter 4.




Figure 1.5.2.1. Amorphous Structure of Ta_2O_5 showing the same planar fragment of Ta_2O_2 as is in the crystalline phase

As a result of sharing many structural features they also share similar nearest neighbour (NN) distances although not necessarily all of them. The reason that these structures only share some of the NNs is because of the long-range order of the

crystalline form. An atom in a crystal structure will see its own NNs as well as each periodic NN, this creates a series of discrete planes of atoms within structure. The amorphous structure does not have the periodicity present in the crystal and as such only has the short range NNs. This is not entirely correct, the longer range NNs do exist in amorphous structures however much less often and as such are weaker (and generally get weaker with increasing distance). Moreover, although as said amorphous and crystalline structures share NNDs, the spread or standard deviation of each of these distances is much different. The periodicity of crystalline structures leads to well-defined NN distances which are represented by a sharp peaks. In amorphous structures, the lack of periodicity leads to much less defined NN distances and much broader distributions for each nearest neighbour distance. For this reason in amorphous systems NNDs are discussed in terms of their modal NND and the width of the distribution.




Figure 1.5.2.2. TEM dark field image (left) and CBED pattern (right) of a Hafnia coating, heat-treated to 100°C showing partial crystallisation despite low heat-treatment temperature.

In material studies it is often difficult to study amorphous materials due to their tendency to crystallise. Crystallisation often leads to a more densely packed, ordered, lower energy arrangement and as such is usually thermodynamically favoured by the material. Crystallisation from an amorphous state can be triggered by amongst other

factors, heat-treatment. Some materials crystallise at relatively low temperatures like Hafnia (HfO_2), which was shown in figure 1.5.2.2 to have crystallised at only 100°C . Other materials can withstand much greater temperatures before crystallising as seen in the figures 1.5.2.3 and 1.5.2.4 for Tantalum and a-Si which remain amorphous at 600°C and 450°C respectively.




Figure 1.5.2.3. CBED pattern from a Tantalum coating heat-treated to 600°C showing that the coating has remained amorphous in the illuminated area; no areas of crystallisation were found in this sample.




Figure 1.5.2.4. CBED pattern from a Silicon coating heat-treated to 450°C showing that the coating has remained amorphous in the illuminated area; no areas of crystallisation were found in this sample.

As discussed above in section 1.4.3 heat-treatment increases the mechanical loss at low temperature for Tantalum [58] whilst decreasing it at room temperature and improves the optical properties [62]. Currently it is unknown why heat treatment affects the mechanical loss; it is assumed that the excess thermal energy changes the atomic structure leading to a different response to mechanical excitation. Studying the link between processing and the atomic structure in coatings is thus of key importance and is a major theme of this thesis.

1.6 Chapter Summary

This chapter has defined what a GW is, a ripple in space-time which carries gravitational information. GWs are important because they will lead to a greater understanding of the universe through providing information about astrophysical events, of a form complementary to that obtained through electromagnetic astronomy. Detection of GWs relies on the quadrupolar nature of GWs, which expand and contract orthogonal lengths in turn, and thus allow for their detection using bar-type and interferometer type detectors. Numerous factors limit the sensitivity of current and future interferometer detectors, the most important of which for this thesis being Brownian thermal noise in the mirror coatings. Although thermal noise is difficult to measure directly, it is related to mechanical loss, which is measurable with suitable equipment, such as in the IGR at the University of Glasgow. Room temperature and cryogenic mechanical loss levels have been shown to depend on the coating material, heat treatment temperature and dopant percentage. These factors are postulated to change the atomic structure of the coating, and in order to understand mechanical loss it is a necessity quantify these structural changes. Most current and proposed coatings of IGWD test masses are amorphous, which adds an additional difficulty since it means the structure has little long-range order. This leads to the need to perform structural modelling to determine the average structure of the amorphous material and this will be discussed at length in Chapter 2.

Chapter 2: Methods of Experimental and Analytical Transmission Electron Microscopy and Density Functional Theorem Constrained Modelling.

2.0 Chapter Introduction

The field of material science has available to it, a host of different methods of study and data collection. Electrons [65], Neutrons [66, 67], photons [68] and many other base particles can be used to probe the structures of materials and unlock their properties and characteristics. Branching out from these base particles are dozens of methods, X-ray diffraction [69], Neutron scattering [70], optical absorption, Electron Energy Loss Spectroscopy [71], all these and much more can be used to provide insights into the properties of a material.

In this thesis, the method of using electrons to generate accurate structural models of material will be discussed. The chosen method of this work is Transmission electron microscopy a method which collects the electron transmitted through a sample from an incident electron beam. The method by which these electrons are collected, or how the Transmission Electron Microscope (TEM) [72] is configured determines the type of signal available. Bright Field, Dark Field, High Angle Annular Dark Field (HAADF) [73], Convergent Beam Electron Diffraction (CBED) [74], Electron Energy Loss Spectroscopy (EELS) [74] are just a few of the possible data sources available in a TEM.

Here it will be detailed how a CBED pattern can be manipulated to form the basis of structural modelling. By stripping away, inelastic atomic scattering the structural information in the diffraction pattern can be isolated. Then using the method of reverse Monte Carlo a structure can be made to emulate this structural information. Following this, through use of Density Functional Theorem Molecular Dynamics the structure can be relaxed into an energetic minimum as would be expected in nature. This results in a structure which is both energetically stable and matches experimentally determined constraints. This structure can then be used to calculate energetic, material, optical properties as necessary.

2.1 Transmission Electron Microscopy and Microscopes

A transmission electron microscope (TEM) allows for the production and focusing of electrons for the purpose of imaging nanoscale structures. In the simplest terms, a TEM consists of an electron source, a series of magnetic lenses and physical apertures, a sample stage and then a collection device, be it a Charge Coupled Device (CCD) or photographic film or a spectrometer [74]. This is shown diagrammatically in figure 2.1.1 and 2.1.2 where the upper and lower sections of a basic TEM are detailed. Although a modern TEM is more complicated than which is alluded to in figures 2.1.1 and 2.1.2 [74].




Figure 2.1.1. Diagram of the illumination system of a basic TEM showing the condenser lenses C1 and C2 and the upper objective lens.

The main lenses of the illumination and projection systems are shown and the electron source would be positioned above the C1 lens. The specimen is placed with the objective lens itself with an upper and lower objective lens above and below respectively. Through adjustment of the projection system the TEM can be altered from diffraction to imaging mode (pictures figure 2.1.2). Further to the system shown in figures 2.1.1 and 2.1.2, further lenses are present beneath the screen to control the electrons onto various detectors. The work detailed here has for the most part been done using a Technai T20 microscope, which uses a thermionic LaB₆ electron source. The T20 has two detectors of interest, a SIS Megaview III CCD camera (SIS) [75] and a Gatan Imaging Filter 2001 (GIF) [76].

The SIS is used for diffraction and bright field work, where the GIF is used for EELS. Two separate sample rods have been used during the course of this work, a FEI single tilt and a FEI double tilt; the former is used for amorphous materials and the latter used for crystalline materials to allow for alignment onto crystalline axes. Some preliminary HDR imaging was performed on a JEOL 2100 TEM, using a Gatan ORIUS CCD with the help of colleagues at the University of Warwick. Finally the experimental results presented in chapter 5 were obtained on a Technai F20 Microscope, employing a thermally assisted FEG and Gatan UltraScan CCD [76].




Figure 2.1.2. Typical state of the projection system of a TEM in “imaging mode” here a 2D image of the sample would be projected and in focus on the screen at the bottom.

2.1.1 Electron Sources

First and for most in a TEMs construction and operation is the electron source or “gun”, the source of the electrons used to image the sample. Many types of gun are available, each with their advantages and disadvantages. Thermionic guns which use heat energy to excite electrons are cheap and easy to maintain and require relatively low vacuum conditions (10^{-7} Torr) [74]. However, typically thermionic sources are less

coherent and have a lesser brightness compared to other guns. Field emission guns (FEGs) use a powerful electric field across sharp, physical tips to produce electrons [74], these are more expensive and require higher vacuum than thermionic source but create much more brighter beams. Thermally assisted FEGs use combination of heated filaments and electric fields to excite electrons, these are very expensive and require ultra-high vacuum (UHV) but create coherent high emission beams and require weaker electric field than normal FEGs. Since the T20 and the JEOL 2100 use thermionic LaB₆ only this type of source will be discussed further here after.

A thermionic electron gun uses heat to excite electrons into their free state, where they can be accelerated down (or sometimes up) the TEM column by strong electric fields. Thermionic current density described by eq.2.1.1 [74], shows us that maximum current density it is achieved by maximising source to a high temperature (T) or by minimising the work function (Φ). High temperature sources use tungsten filaments which benefit from a very high melting point and as such can be heat extensively. The T20 and JEOL 2100 use LaB₆ filaments because LaB₆ has an exceedingly low work function. The work function of LaB₆ is ~2.5eV compared to Tungsten's 4.32-5.22eV. However, typical operating temperatures of Tungsten and LaB₆ filaments are 2600-3000K and 1400-2000K respectively. It is necessary to run a LaB₆ filament at a higher vacuum than a tungsten filament, 10⁻⁷ compared to 10⁻⁵ Torr. However, the typically, the lifetime and brightness of the LaB₆ filament is superior to that of the Tungsten one.

$$J = AT^2 \exp\left(-\frac{\Phi}{kT}\right)$$

[eq.2.1.1]

Where J, A, T, k and Φ are the Thermionic emission current density, a constant, the temperature of the source, the Boltzmann constant and the work function of the source respectively.

Figure 2.1.1 shows the typical setup of a thermionic gun, a high tension potential difference is applied across to the filament to induce heating and release the electrons. The electrons are then accelerated towards the cathode (Wehnelt cup) and enter the main TEM column thereafter [74].




Figure 2.1.1: Basic design of a thermionic electron gun, current is passed through the filament causing significant heating releasing the electrons. The electrons are then columnated and accelerated down the column.

2.1.2 Electron Lenses

So now that we have electrons in our TEM, how do we control them? In photon optics the incoming waves are focused and magnified using lenses, typically made of glass. Similarly in electron optics magnetic lenses are used to focus images and direct the beam [74]. Magnetic lenses work on the simple principle that a charged particle passing through a magnetic field will experience a force perpendicular to both the direction of motion and the magnetic field (The Lorentz Force). This principle is described by $\mathbf{F} = q(\mathbf{v} \times \mathbf{B})$, where \mathbf{F} , \mathbf{v} , and \mathbf{B} are the vector forms of Force on, velocity of and magnetic field passed though by a particle with charge q . Obviously due to the nature of

the cross product in the previous equation, the induced force only alters the path of the electrons and not the speed of it. Thus by generating a magnetic field in the path of the electrons, their path can be controlled similarly to how glass lenses control photons.




Figure 2.1.2. Design of a basic magnetic lens, a magnetic field is induced by passing current through a series of coils. This magnetic field results in a force which acts on the electrons as they pass through the field adjusting their trajectory. Thus by varying the current in the coils the magnetic field strength can be changed and used to control the path of the electrons.

Magnetic lenses are simple iron cores with coils wrapped around them (figure 2.1.2), passing an alternating current through this coil then induces a magnetic field. By adjusting the magnitude of the current passing through the coil the resultant magnetic field can be controlled and with it the path of the electrons. As would be expected, as with photon optics, electron optics suffer from aberrations due to interactions with the lenses. Various species of aberrations exist, chromatic aberrations are caused by a varying electron wavelengths, spherical aberrations are caused by a breakdown of the thin lens approximation. Aberrations cause the focal point to become undefined, instead resulting, in a disk of closest approach. The diameter of that disk d is dependent of the coefficient

of aberration C and the aperture angle α . For chromatic aberration $d_c \propto C_c \alpha$ where as for spherical aberration $d_s \propto C_s \alpha^3$ as such, spherical aberration is typically the controlling factor of resolution.

Chromatic aberration is due to the fact that refraction due to a lens is frequency dependent. Since a typical electron source will accelerate electrons by a varying amount, there will be a spread in the electron wavelengths. This spread will translate into a spread in focal lengths at each lens and result in blurred edges and ultimately reduced resolution. Chromatic aberrations can be eliminated through use of a monochromatic source, with the same De Broglie wavelength all electrons will be deflected by the same amount. A truly monochromatic source is not achievable currently however; it is possible to reduce the spread in the electron wavelengths. Furthermore C_c can be reduced further though use of a monochromator in the column, beneath the gun.

Spherical aberration is caused by the fact that lenses be they optic glass lenses or magnetic electron lenses are not infinitely thin. In the thin lens approximation all rays (electrons) will be focused to the same point as the lens will act equally on all rays. However, in reality this is not true; typically a lens is thinner at its edges than in its centre. This leads to a greater refraction in the beam at the extremes of the lens compared to that at the centre. Again this leads to a focal disk rather than a focal point resulting in blurred images and a decrease in resolution. Spherical aberration can be tackled by imposing anti-aberrations, essentially cancelling out the aberration caused by the working lens. Though the exact methods employed to reduce spherical aberration will not be discussed, a detailed discussion of these methods is available here [74].

2.1.3 Physical Apertures

Apart from magnetic lenses, which are used to control the path of the electrons, TEMs also utilise physical apertures [74]. These apertures are simply, small holes placed in the beam path at specific points in the TEM. The aperture thus blocks/absorbs/reflects any electron whose path (convergence angle) is greater than the diameter of the

aperture. Apertures can improve contrast, decrease illumination and limit the convergence or divergence of the beam to suit the needs of the operator.

2.1.4 Typical Detectors used in electron microscopy

Data collection on TEMs is typically done using a photographic film, a spectrometer or a CCD camera. CCD cameras in electron microscopes do not detect the electrons directly (though direct electron detectors do exist [76]) instead they detect photons produced by the incident electrons. Cameras convert incident electrons into photons by means of scintillator placed above the CCD chip. Photomultiplier tubes are used to increase the photon signal strength for easier detection. These photons then cause excitations in the CCD chip material (typically silicon [77]) and the resultant charge can then be “read off” and an image constructed.

CCD technology was first proposed by Boyle and Smith [78] working at bell labs in the late 60s for which they later were awarded a Nobel Prize. The basic principal is that incident photons excite electrons in the conduction band in to the valence band causing an electron-hole pair. That charge is then passed through several layers of electronics and read off into a computer for interpretation. A full explanation of CCD technology is available here [78] however some of the advantages and disadvantages of using CCDs for data collection will be summarised here.

The main advantage of CCD cameras over photographic film is the speed of read off, CCDs can produce images in less than second, where film requires development, a slow process. CCDs are susceptible to noise for a variety of sources, stray photons can interact with the CCD directly. Stray electrons can interact with the scintillator and create spurious signal aswell. CCD cameras are also susceptible to damage as well, due to an overexcess of electron current. In a situation where a CCD camera is exposed to a high current electron beam it can be damaged in various ways. The scintillator which is responsible for converting incident electrons into photons, while this is a reversible process there is a certain relaxation time associated with the process. In high current

conditions the scintillator material is not able to relax sufficient quickly to accommodate the number of incoming electrons and instead the structure of the scintillator is irreversibly altered. This results in an effective drop in the responsiveness of the scintillator at the point of damage. The CCD itself can also take damage due to a high current beam; this however, typically results in “dead” pixels, which are simply completely unresponsive ones.

For these reasons it is important to protect a CCD camera from high current beams, or high current areas of the beam. For this reason, typically during diffraction work the central beam (where the current is highest) is blocked by a beam blank. Alternatively the current can be reduced by lowering the brightness, or using physical apertures. More advance, robust cameras such as the Gatan Orius [76] are capable of resisting beam damage even from the presence of the central beam. Although the Orius is far from indestructible it can resist significant current and thus a beam blank is unnecessary.

2.1.5 Cross Section Encapsulation Sample Preparation

Samples must be prepared before TEM study; this involves creating a specimen of the correct dimensions and electron transparent. A typical glass sample starts its life as a 1 inch diameter, ¼ inch thick Silica (SiO_2) or Silicon discs which are then coated using dual ion beam assisted sputtering [79]. Coating thickness's range from ¼ wavelength (120-170nm) to 2.5µm and is dependent on the coating vendor or initial purpose for the sample. The nature of samples lends itself well to an encapsulation method [80] in which a small sliver of sample material is inserted into a molybdenum rod which is then “encapsulated” in either a titanium or brass tube. This is then thinned using mechanical and ionic processes resulting in a standard 3mm sized sample with a small electron transparent area in the centre.




Figure 2.1.5. Schematic diagram of the cross section encapsulation method of sample preparation. Beginning from a 1" disc (a), it is halved (b), a sliver is sliced off (c) leaving a small cuboidal sliver of sample (d). This is then cut in two and the pieces glue together (e), the substrate is then removed (f) and then the whole sample is slotted into a molybdenum rod and fit into a brass tube (g). The tube is then cut into a series of discs (h) each of which is mechanically dimpled (i) finally the sample is thinned to electron transparency using an ion milling system (j).

Sample preparation starts by halving the 1inch disc sample (figure 2.1.5 a) and b)), this creates a flat surface from which the TEM samples can be cut. To protect the coating surface during the cutting process, the samples are secured to standard glass microscope slide using wax, coating side down, and the cut is made from the uncoated side. A 2nd cut is then made approximately 2mm from the flat edge of one of the new semicircular disc fragments (figure 2.1.5 c)). This results in a 2*25*6mm sliver (figure 2.1.5 d)), the vast majority of the substrate is then removed (figure 2.1.5 e)). This is then cut in half creating two slivers, each approximately 12.5mm long. The coated sides are then glued together (figure 2.1.5 f)), and inserted into a molybdenum rod, and fitted to a brass tube (figure 2.1.5 g)). The rod is then cut into disc's each one can be made into a separate TEM sample and a typical rod can be made into 5-6 samples if necessary (figure 2.1.5 h)),. The discs are then vacuum baked to set the glue completely and remove any contaminants which maybe present on the surface, the baking also help to remove any free carbon from the glue. The discs are then mechanically dimpled on both sides (figure 2.1.5 i)) and finally ion milled at a relatively low energy until a hole appears at the centre of the disc, creating thin (electron transparent) areas on the sample (figure 2.1.5 j)).

2.1.6 Electron Diffraction

This is an electron diffraction based investigation, it used at every stage in the experimental and computational phases in this work. As such it is important to discuss what diffraction is both in general and specifically in a TEM and how it is used here. Before continuing it is important to define what diffraction is and what scattering is (why will become apparent later). Diffraction is defined as “a deviation in the direction of a wave at the edge of an obstacle in its path” and scattering is defined as “the process in which particles, atoms etc, are deflected as a results of collision”. These definitions are from the Collins’ Dictionary and basically state the diffraction and scattering are the same thing in fact diffraction is a just a special form of elastic scattering [74]. Since electrons can be treated as both waves and particles, their nature allows for both scattering and diffraction.

Diffraction patterns are formed through interference of recombining beams. Let's examine the situation of electron diffraction studies of a crystalline material. A simple crystalline material consists of some primitive cell, tessellated in 3 dimensions to create a complete structure. Since all these primitive cells will have the same structure and orientation there will exist planes of atoms. These planes will intersect with identical atoms in multiple primitive cells throughout the structure. Diffraction between these planes will cause a phase difference (and path difference) and thus lead to interference during recombination. These recombined beams lead to a series of defined "spots" in the diffraction pattern. The spot like pattern is due to the highly ordered nature which essentially discretises the angle of reflection. As would be expected due to interfering reflection and recombination's certain reflections will not be allowed. In certain cases, beams will recombine with a phase difference of π , thus destructively interfering leading to the absence of a diffraction spot.

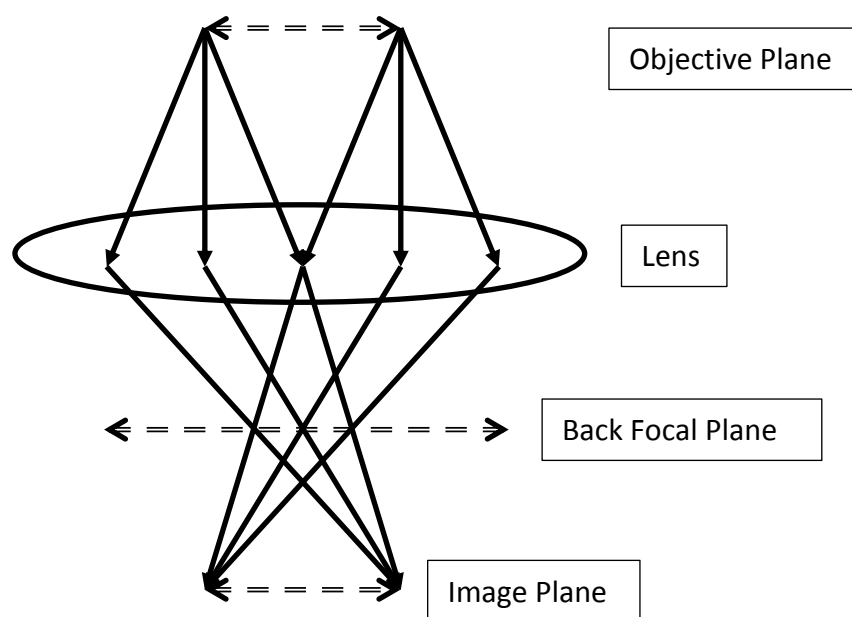


Figure 2.1.6.1 Diagram of the back focal plane of a TEM where a diffraction pattern forms. This plane is projected onto the imaging plane (a screen or camera) during "Diffraction Mode" for diffraction studies.

The diffraction pattern is formed in the back focal plane of the microscope, slightly below the objective lens. During “diffraction mode” this plane is projected on the imaging plane through realignment of the projector lenses (seen in figure 2.1.6.1.).

So electrons can be scattered and diffracted by samples, and exactly how the electrons are scattered and diffracted is of interest to us. Electron scattering can be affected by a great many factors in a TEM, sample thickness, incident electron energy and atomic weight of the target atoms. In diffraction studies, single scattering events are what are most preferable, as the scattering angle is then only dependent on a single interaction. Multiple scattering events can also allow forbidden reflections to occur since the beams will no longer have a phase difference of π . This can obviously lead to false analysis and conclusions as reflections will appear allowed when, in reality they are not. To ensure (or increase the likelihood of) single scattering events samples must be thin <20nm, although electron transparency is possible at much greater thicknesses. Why are single scattering events preferable? While single atomic scattering has a degree of randomness, scattering from structures does not. The basis of diffraction studies is that regular structural features scatter (diffract) electrons in a predictable way and thus by analysing a diffraction pattern structural information can be deduced. However for the most part this process assumes the incident electrons are normal to the surface. In the case of multiple scattering, this assumption breaks down, especially in cases where there has been significantly more than one scattering event. In a multiple scattering events it is impossible to ascertain the true scattering angle and since the angle is indicative of the feature (interatomic distance) that led to it this information is lost. This is why it is absolutely imperative to ensure samples are thin and that single scattering is the dominant process.




Figure 2.1.6.2 Selected Area Diffraction (SAED) pattern of a silicon crystal sample (left) shown in false colour; notice the oversaturation of the central spots and the distortion in the lefthand side of the image due to the prism. Convergent Beam Electron Diffraction (CBED) pattern of a silicon crystal sample (Right). These two images were taken at the same camera length and on the same crystalline orientation

There are many different methods of obtaining diffraction data, historically selected area diffraction has been the standard. CBED has advantages over SAED most notably, information is collected from a much smaller area using CBED which of great use to those working with thin films. In SAED parallel beam conditions are in affect which leads to a greater Q resolution than CBED shown in figure 2.1.6.2. Although this is a disadvantage of CBED over SAED, the advantages of high special resolution of CBED outweigh this.

2.1.7 Electron Energy loss spectroscopy

A useful tool in microscopy is electron energy loss spectroscopy [81] which involves studying the energy loss statistics due to inelastic scattering. The principle of EELS is that as electrons pass through a sample, they will scatter inelastically and such lose energy. The amount of energy lost during a scattering event will be discretised by the specific excitation it caused. Since the energy loss will is determined by the nature of the scattering event, then it is possible to predict which atomic electron or other energy loss process was involved. In order to view EELS spectra it is necessary to first separate out the

electrons as a function of their energy and since an electron source is approximately monochromatic electron energy is easily relatable to energy loss. This energy separation is done using an omega spectrometer, which works by applying a perpendicular magnetic field to the incoming electrons, causing Lorentz deflection, which separates the electrons as a function of their energy.

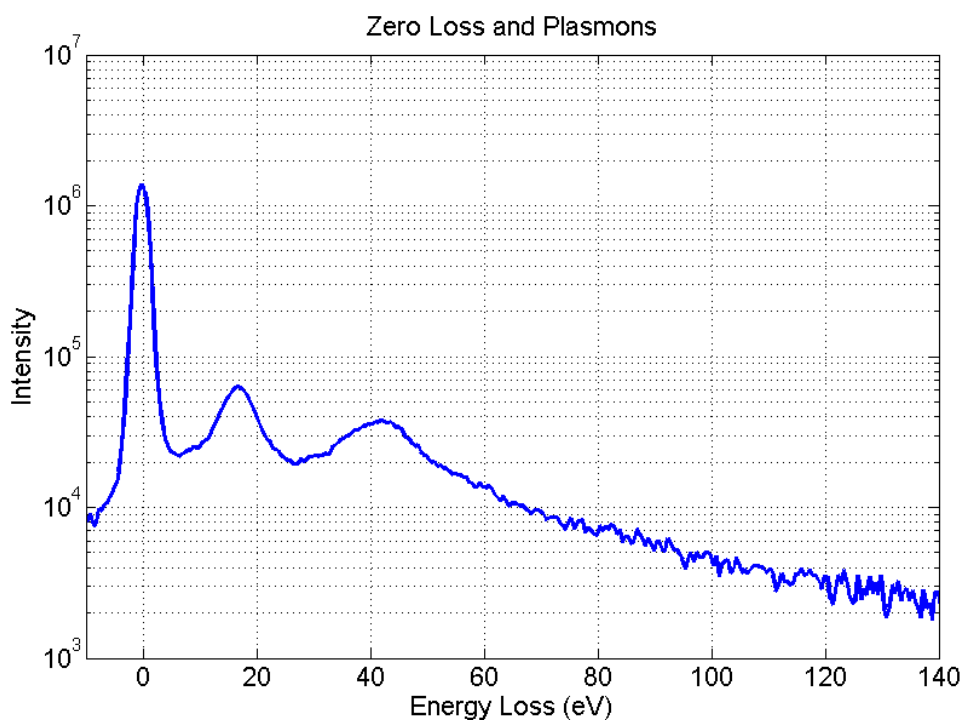


Figure 2.1.7.1. Typical EELS spectra showing the zero loss peak at ~ 0 eV and the Plasmon excitations at 20 and 40eV.

This then produces a typical EELS spectrum which is a combination of signals from a variety of different energy loss processes. Firstly there is a Zero energy loss peak from electrons which haven't undergone inelastic scattering. The next feature is caused by Plasmon resonances in the sample, these are low energy interactions and as such appear low on the energy scale (see figure 2.1.3.1). The next most prominent features are core losses (figure 2.1.3.2) which are caused by interactions with specific atomic electrons and are as such characteristic to the element and each element can have a number of edges at different energies corresponding to various occupied orbitals. All these features are

superimposed on to a background intensity caused by random inelastic scattering which is more probable to occur at low energies and as such the background intensity tails off exponentially from the zero loss.

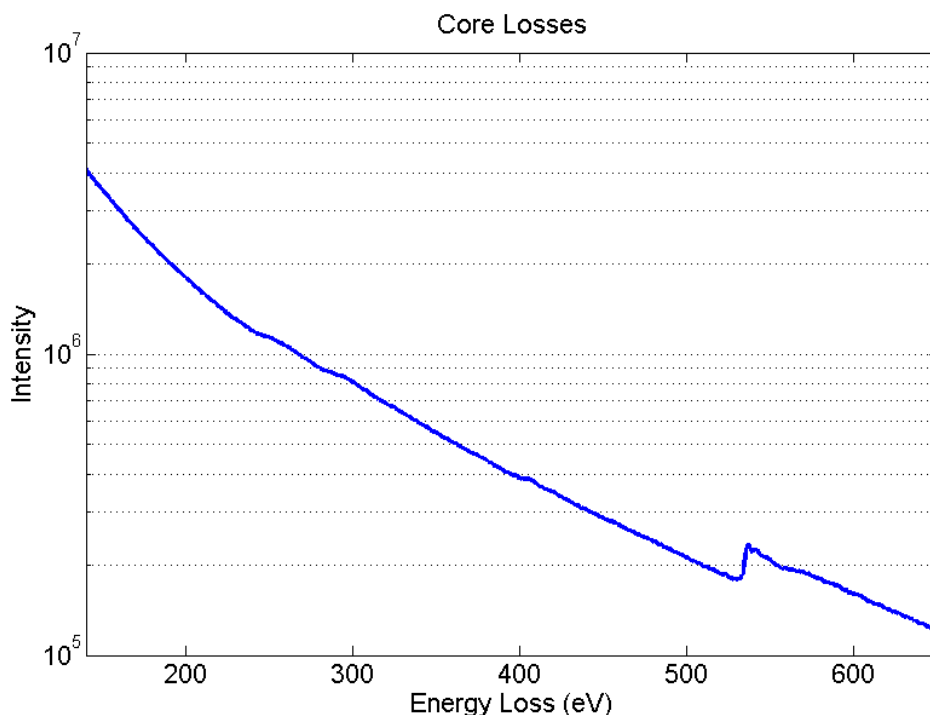


Figure 2.1.7.2. Typical EELS spectra showing the core loss region, notice the strong edge at $\sim 530\text{eV}$ and the weaker edges at 250eV , 280eV and 400eV .

Due to the presence of the background intensity, it is necessary when analysing EELS spectra to model this background in order to determine the true shape of each feature in the spectra. This is done by defining multiple energy bands within spectra away from characteristic edges. Assuming that this intensity is only due to the background intensity from mostly bremsstrahlung radiation which will be decaying exponentially and fitting as such. With the background modelled in the core loss range, it can be subtracted leaving the intensity from the core loss interactions. Through integration under these core loss peaks, the total intensity for a particular "edge" can be determined. Typically a range of energies will be chosen which contains 1 edge from each constituent element, however their relative abundances are not simply a ratio of the total intensity of their

respective peaks. Since each element and each unique edge will have its own reaction cross section, the intensity of an edge will be scaled to reflect this. In order to determine this cross section we must use a calibration sample which contains this edge. For tantalum based coatings we use a crystalline LiTaO_3 [82] sample to calibrate for the Ta-N edge, which is the edge used to determine the stoichiometry of the coatings. For crystalline LiTaO_3 the ratio of Ta:O is known (1:3) and consistent, thus by measuring the Ta:O edge intensity ratio and scaling it relative to the well-known O cross section it is possible to determine the Ta cross section. With the cross sections for the relevant edges for Ti, O and Ta (for a $\text{Ti:Ta}_2\text{O}_5$ coating) known, the intensities can be corrected for them. Finally since, in the case of a $\text{Ti:Ta}_2\text{O}_5$ coating it is assumed that the EELS spectra contains only information from Ti, Ta and O (no contaminations) the relative abundances are then normalised to 1 i.e. $\% \text{Ti} + \% \text{Ta} + \% \text{O} = 1$.

2.2 Reduced Density Functions

A key tool available to the amorphous material scientist is the Reduced Density Function (RDF) A mathematical manipulation of diffraction data [64]. The RDF is a statistical representation of atomic positions within an amorphous structure with respect to a central, arbitrary atom. RDFs can be generated from all types of diffraction data, x-ray [83], or electron [64] all that is required is a series of coefficients which describe the intensity of the scatter beam (be it x-ray, electron etc.) for each constituent element in the structure [64].

The calculation of an electron RDF begins with a typical CBED pattern from an amorphous material such as figure 2.2.1. By determining the coordinates of centre of the pattern a radial average can be performed reducing random noise, amongst other things. This results in the function $I(q)$ which describes the average intensity of the pattern as a function of the distance q from the centre q_0 a typical example of this is shown in figure 2.2.2. Where $q = 4\pi \sin \vartheta / \lambda$, 2ϑ is the scattering angle and λ is the wavelength of the scattered radiation.




Figure 2.2.1. Typical CBED pattern of Tantalum in false “temperature” colour taken using the SIS camera at a camera length of 860mm.




Figure 2.2.2 Intensity Pattern (left) from radial average of a masked CBED image in figure 2.1.1.1. The saturation of the camera occurs at ~ 3600 counts, indicating this image is safely below saturation. The q range has been cut from 1.2-16.5 which represents the edge of the central beam and the edge of the camera.

The function $I(q)$ is described by the Debye formula (eq. 2.2.1) where $r_{ab} = |r_a - r_b|$ the distance between two atoms at points r_a and r_b respectively. The atomic scattering intensity is described by the function f^2 defined by eq. 2.2.2 and N is a fitting parameter related to the relative stoichiometry of the system and $Nf^2(q)$ in this context is a term describing the mean intensity.

$$I(q) = Nf^2(q) + \sum_{a \neq b} \sum_b f^2(q) \sin(qr_{ab}) / (qr_{ab})$$

[eq. 2.2.1]

$$f^2(q) = \left(\sum_{i=1}^{i_{max}} A_i \sum_{n=1}^3 (a_n / ((q^2 / 4\pi^2) + b_n) + c_n e^{-d_n(q^2 / 4\pi^2)}) \right)^2$$

[eq. 2.2.2]

Where A_i is the relative abundance of the i -th atomic species in the material and i_{max} is the total number of unique elements in the material e.g. $i_{max} = 2$ for CO_2 . The constants a_n , b_n , c_n and d_n are the so-called scattering parameters which determine the shape of the Lorentzian (1st term of eq. 2.2.2) and Gaussian (2nd term of eq. 2.2.2) distributions which define the shape of the total scattered intensity and example of which is shown in figure 2.2.2. This process of using a sum of Gaussian and Lorentzian distributions was developed by EJ Kirkland [84] as a modification of a previous system of using Gaussian distributions to describe the scattered intensity [85]. It was determined that whilst Gaussians described low q scattered intensity well, they were inadequate at high q ; Lorentzians, however, describe the intensity extremely well at high q [85] but falter at low q . Hence the combination of the two provides good representation of the scattered intensity across the typical q range.

By integrating eq. 2.2.1 over the sample volume we can obtain eq.2.2.3 show below.

$$I(q) = Nf^2(q) + f^2(q) \sum_b \int g(r_{ab}) \sin(qr_{ab})/(qr_{ab})dV$$

[eq.2.2.3]

Then by defining the reduced density function $G(r) = 4\pi r(g(r) - \rho_0)$ we can write eq.2.2.3 as

$$I(q) = Nf^2(q) + f^2(q) \sum_b \int 4\pi r^2(g(r_{ab}) - \rho_0) \sin(qr_{ab})/(qr_{ab})dr_{ab} \\ + f^2(q) \sum_b \int 4\pi r_{ab}^2 \rho_0 \sin(qr_{ab})/(qr_{ab})dr_{ab}$$

[eq.2.2.4]

The third term in eq.2.2.4 is only contributions significantly at low q so by omitting this range the term can be ignored. For a homogenous, isotropic system we can take $r_{ab} = r$ and define the reduced scattering function $\varphi(q)$ as eq.2.2.5 we can calculate $G(r)$ simply as the Fourier transform of $\varphi(q)$ (eq.2.2.6).

$$\varphi(q) = \left(\frac{I(q) - Nf^2(q)}{Nf^2(q)} \right) q$$

[eq.2.2.5]

$$G(r) = 4 \int_0^\infty \varphi(q) \sin(qr) dq$$

[eq.2.2.6]

2.2.1 Practical Application of RDF method

In practice the application of the RDF method requires slight modifications to account for experimental conditions and compound structures. For compound structures such as Ta₂O₅ eq. 2.2.5 is no longer valid and is replaced by eq.2.2.7 which takes account of each constituent elements relative abundance and individual Kirkland scattering parameters.

$$\varphi(q) = \left(\frac{I(q) - N\langle f^2(q) \rangle}{\langle Nf(q) \rangle^2} \right) q$$

[eq.2.2.7]

Where

$$\langle f \rangle^2 = \frac{(\sum_i N_i f_i)^2}{N^2}$$

$$\langle f^2 \rangle = \frac{\sum_i N_i f_i^2}{N}$$

And

$$N = \sum_i N_i$$

Where N_i and f_i are the relative abundance (in quantity) and associated atomic scattering of the i -th element respectively.




Figure 2.2.3 Reduced intensity functions (left) and Reduced density functions (right) for the intensity profile shown in figure 2.2.2. Notice that the undamped $\Phi(q)$ shows considerable noise at high Q and that damping factor of 0.4 can significantly reduce this. The periodic oscillations seen in the undamped $G(r)$ (blue) are also significantly reduced by the introduction of the damped factor.

Aside from the necessary modification to the base RDF method for use on compound structures, there are also adjustments required to take account of experimental inaccuracies and estimations. The first of these adjustments is necessary to

correct the Fourier transform in eq.2.2.6 for in $\varphi(q)$ of finite q . In order to evaluate eq.2.2.6 it is required that q range from 0 to ∞ , which is impossible due to experimental constraints. To take of account of finite q a truncation factor must then be applied to the Fourier transform [86] which limits it to the range $0 \leq q \leq q_{max}$ where q_{max} is the maximum value of q . Side note* Since this is a Fourier transform, there is an imposed resolution limit of $\Delta r = 2\pi/q$, which lies in the range 0.48-0.37Å for the data collected here with a q_{max} range of 13-17Å⁻¹.

The truncation process inherently introduces oscillations in $G(r)$ with a frequency of $q_{max}/2\pi$ these can be seen in figure 2.2.3. These have a detrimental effect on the interpretation of $G(r)$ if not properly dealt with. Warren [87] detailed a method to reduce the oscillations in an $G(r)$ and the effect of noise in the intensity pattern by introducing a damping factor $exp(-bq^2)$. Many modifications functions are available such as that *Lorch* [88] though will not be described here. Where b is a constant used to define the degree of damped, i.e. noisy data would require heavier damping. This function tends to 0 at high q which is usually the source of oscillations in $G(r)$ due to low SNR. Values of b in the range 0.4 - 0.6 shown in figure 2.2.3 are usually sufficient to eliminate any noticeable oscillation in $G(r)$ without unduly affecting $G(r)$ itself [86].

Finally, it is not often possible to fit $N\langle f^2(q) \rangle$ to $I(q)$ as it requires that $I(q)$ tend to zero at high q which is rarely the case. It is therefore necessary to introduce the value C to the fitting process such that $I(q)$ and $N\langle f^2(q) \rangle$ tend to C at maximum q i.e.

$$C = N\langle f^2(q_{max}) \rangle - I(q_{max})$$

As such eq.2.2.7 becomes.

$$\varphi(q) = \left(\frac{I(q) - N\langle f^2(q) \rangle + C}{\langle Nf(q) \rangle^2} \right) q$$

[eq.2.2.8]

It is then possible to obtain a good fit $I(q)$ by adjusting the values of N and C .

2.3 Reverse Monte Carlo

Reverse Monte Carlo (RMC) [89] structural refinement is a process used to generate atomic structures which reproduce experimental data [89-91]; in the case here, the experimental data is of course an RDF. The process of RMC starts by generating a box of atoms, either by a model of X atoms such as figure 2.3.1 and then sequentially moving a chosen atom in a random direction by a random amount and then accepting that movement if it results in a better fit to the experimental data (figure 2.3.2). There is of course an acceptance criterion in the RMC which allows a detrimental atomic movement to be accepted if the probability of movement leading to a better fit at some point is large enough. This allows the model to escape local minima in favour of a global one by traversing over the barriers between them.




Figure 2.3.1 200 atom randomly arranged model of Silica (SiO_2) this model would be the starting point for an RMC refinement.

In the case here further experimental constraints are placed on the RMC process through EELS and X-ray Reflectometry (XRR) to constrain the stoichiometry and Density

respectively. The only constraints are a maximum atomic step size (0.1\AA unless otherwise stated) and a minimum interatomic distance for each species equal to the sum of the covalent radii taken from the Cambridge Structural Database [92]. Reverse Monte Carlo structural refinement has been used extensively in the literature especially in the X-ray diffraction community [64]. RMC has not however, been used extensively to track the change of a structure with changing composition [64]. The advantage of RMC is that relatively large structures can be generated quickly using a wide range of experimental input [64]. A disadvantage of RMC is that it takes no account of the energy of the structure, although constraints can be put in place to stop exceedingly close approaches between atoms. There are hybridised codes which perform RMC iterations and Self Consistent Field (SCF) style calculations simultaneously [93], however these are computationally very expensive as it the RMC process typically requires many thousands of iterations to achieve convergence [93].




Figure 2.3.2 200 atom RMC model of Silica (SiO_2) this model is now in experimental alignment with a silica RDF. This model shows considerably more bonding than the randomly arranged model and would serve as the starting point for a molecular dynamics simulated anneal. The “free” oxygen seen in the model is an artefact from the plotting program and does not imply free oxygen within the structure.

One of the advantages of the hybrid modelling process discussed here is that by combining molecular dynamics with RMC sequentially a structure can be made to fit experimental data and then be energetic relaxed. This avoids the need to perform an SCF convergence at every RMC iterative massively reducing the computational expense. The energetic relaxation performed using CASTEP [94] allows the structure escape high energy potential wells and undo certain “unnatural” structural configurations which the RMC may have created. However energetic relaxation often leads to model “drifting” away from alignment with the initial experimental constraints imposed during the RMC process.

The drift from experimental alignment is an indication of high energy structural features being created by the RMC process. These are then relaxed by the MD process, however these high energy features are often integral to the fit with experimental constraints and hence the drift. The degree of drift will be a combination of affects from the pseudopotentials being employed, the various constraints of the DFT (energy tolerances, simulation duration) and the accuracy of the original RDF used. To clarify, if the original RDF is not a true representation of structure then degree of drift of the model is likely to be different from that a model based a “ideal” RDF which perfectly describes the system. If the drift is severe enough then it is necessary to refine the post-MD model back into experimental alignment. This can be done once again using RMC, however this 2nd refinement differs in that the maximum movement of any atom is severely restrained compared to the 1st refinement. It is possible that a 2nd refinement process will essentially undo any relaxation the MD process achieved and simply return the structure to its original state. By restraining the maximum movement of the atoms we hoped maintain as much of the energetic relaxation obtained during the MD process although it is accepted that some will be lost. At this point the model is also increased in size to allow for much longer order to be calculated and greater statistical accuracy in the final analysis process. This is done by making copies of the post-MD structure (cell) and arranging these copies into a 3x3x3 “supercell”, the orientation of each cell is randomised to one of the 6

possibilities (for cubic cells) to avoid artificially creating longer range order. This results in a model with ~2700 atoms, which can then be refined back into experimental alignment using the previously mentioned restrained movement. In the future, given greater computing resources it could be possible to setup an iterative process between RMC and MD which would lead to a structure which both experimental aligned and fully energetically relaxed.

2.4 CASTEP and Density Functional Theorem

2.4.1 Density Functional Theorem

As mentioned, the modelling process used here includes a simulated anneal (SN) DFT MD simulation to relax the energy of the models. This part of the process is done using software called CASTEP [94] which is an ab-initio total energy program capable of performing a number of different DFT calculations. Before we can go any further it is necessary to explain what DFT is. In quantum mechanics we seek to solve the Time-dependent (eq.2.4.1.1.) or Time-independent Schrödinger's equations (eq. 2.4.1.2) for the wavefunction ψ to learn about the system and its states. Density functional theory states that the groundstate properties of a system can be described by considering the groundstate charge density

$$i\hbar \frac{\partial}{\partial t} \psi_i(r) = \hat{H} \psi_i(r)$$

[eq.2.4.1.1]

$$\hat{H} \psi_i(r) = \epsilon_i \psi_i(r)$$

[eq.2.4.1.2]

Where the time dependent and independent Hamiltonians [95] are defined by

$$-\frac{\hbar^2}{2m}\nabla^2 + v(r, t)$$

And

$$-\frac{\hbar^2}{2m}\nabla^2 + v(r)$$

2.4.2 The Kohn-Sham Equations

The Hohenberg-Kohn theorem shows us that it is possible to rewrite the Schrödinger equation in terms of the charge density and that this can be used to calculate the system's properties [95]. A method of determining the groundstate charge density has been devised by the Kohn and Sham in the form of the Kohn-Sham equations (eq.2.4.1.3 and 2.4.1.4) [96].

$$\left[-\frac{\hbar^2}{2m}\nabla^2 + v_{eff}(r) \right] \psi_i(r) = \epsilon_i \psi_i(r)$$

[eq.2.4.1.3]

Where v_{eff} , m , and ∇^2 are the effective potential, mass of the particle and the laplacian operator, and ϵ_i is the energy of the Kohn-Sham orbital described by the wavefunction $\psi_i(r)$. The effective potential v_{eff} is described by eq. 2.4.1.4 where $v_{xc}(r)$ is the potential due to the exchange interaction.

$$v_{eff}(r) = v(r) + \int \frac{\rho(r')}{|r - r'|} dr' + v_{xc}(r)$$

[eq.2.4.1.4]

To solve eq.2.4.1.3 it is necessary to determine v_{eff} and hence necessary to determine the exchange correlation potential $v_{xc}(r)$. In DFT we use an exchange correlation

functional to describe $v_{xc}(r)$, creating accurate XC functional has and is the focus of a great deal of research [97].

2.4.3 Exchange Correlation Functionals

One of the more standard functionals is based on the Local Density Approximations (LDAs) which are that the exchange correlation is based purely on electron density at each point in space. Functionals based on the LDAs are extremely successful and have lead onto the more sophisticated approximations used in Generalised Gradient Approximations (GGAs) and hybrid functionals. The GGAs add a further term into the LDAs from the gradient of the density which allows for systems in which the density varies rapidly. The research presented in this thesis was entirely done using the Perdew-Burke-Ernzerhof exchange correlation functional [98]. PBE uses the GGAs and is the de facto standard functional [99] for use with solid state systems as it describes system properties well.

2.4.4 Pseudopotentials

When dealing with periodic systems there are distinct advantages to using a plane wave basis set [98] however to describe the region closest to the atomic core it is necessary to use a pseudopotential. The use of a plane wave basis set creates two issues when describe the atomic core [100]:

- The electrostatic potential energy is asymptotical in the atomic core region
- The valence electron wavefunctions must be orthogonal to the core electron wavefunctions

These two issues lead to the need for a large amount of plane waves being needed to describe the system. Through the use of pseudopotentials this issue can be circumvented whilst maintaining accuracy. This is achieved by generating a new potential which softens the atomic core potential in the core region (defined by some cut-off

radius) whilst describing the all electron potential outside the region. This is justified since the majority of physical properties are controlled by the valence electrons. The valence electron then must see an effective potential which is due the nucleus and core electrons which is, overall, weaker than the true core potential [98]. This creates a “softer” potential within the core region [101], lowering the number of plane waves required to describe the valence electron. Furthermore the core electron wavefunctions no longer need to be calculated as they are handled by the pseudopotential. Pseudopotentials are typically generated using one of the two methods described by Kleinman-Bylander [102] or Vanderbilt [103] and will not be described further here.

2.4.5 Self Consistent Field Theory Iterations

Solving the K-S equations is computationally intensive but is necessary to compute the Eigenstates of the system. That which takes the most time to compute is the K-S wavefunctions and the charge density. To do this, it is necessary to diagonalise the Hamiltonian which is unfeasible given the computational expensive it entails. However since the objective is to calculate the groundstate energy, full diagonalisation can be avoided since it results in Eigenstates then the lowest ones require. It is therefore possible to locate the groundstate through an iterative procedure meant to just compute the lowest eigenstates.

CASTEP employs different methods to achieve this, Density Mixing (DM) and Ensemble Density Functional Theory (EDFT) [94]. DM begins with an estimation of the groundstate charge density, in EDFT the density is recalculated during each iteration and past onto the next. As DM is the method used here, only this one will be discussed further however, a full description of EDFT is available here [104]. With an initial estimation of the charge density the K-S Hamiltonian can be constructed. The Hamiltonian can then be diagonalised and the energy bands calculated as well as the energy. The energy is then compared to the previous calculation, if they equate, within some energy tolerance then the cycle exits. If they do not, then the charge density is calculated and “mixed” to estimate the true groundstate density [34], this is then used in the subsequent cycle. This

process is called the Self Consistence Field (SCF) cycle [94] and is used in both the DM and EDFT methods though the exact calculations differ between them.

2.4.6 Molecular Dynamics

Molecular Dynamics (MD) is a method of simulation in which the equations of motion are solved for the simulated system [100]. Typically MD simulations are performed at some finite temperature giving the system thermal energy, allowing the atoms to move. During a MD simulation iteration, a SCF calculation is performed to determine the energy of, and forces on the system. Simulation time then steps forward some interval, and a subsequent iteration is performed. Between subsequent iterations, the system will move and reconfigure due to the presence of thermal energy and internal forces (stresses). Over a long enough period of time, the system will evolve towards some point of minimal stress. Though this is not a directed evolution as with Geometry Optimisation it is a more natural, more unpredictable process.

A typical MD simulation will last perhaps several picoseconds, during which time the energy of the system will vary seemingly (at times) randomly, which is due to thermal vibrations. However assuming the system is free to explore a large enough potential energy landscape (as opposed to becoming trapped in a well) the energy of the system will (on average) fall over time as the system approaches a more energetically favourable configuration. The purpose of MD is not to achieve an energetic groundstate but to simply relax the structure. This relaxation can be internal i.e. only the atomic positions change; lattice relaxations are also possible, in which the lattice parameters are “unlocked” and allowed to vary according to stress and such. Here the “lattice” is define the as dimension of the simulation box, and should not be confused with the lattice of a crystalline material. To clarify, “lattice relaxation” in this context is a process allowing the box dimensions (predefined cubically) to vary in the simulation. The end result of an MD simulation is a snapshot of an energetically relaxed structure; “snapshot” refers to the fact that the end point of an MD simulation is arbitrary and were the simulation allowed to continue, it would not remain at the same energy (mainly due to thermal motion).

Our use of MD involves running a so-called “simulated anneal” (SN) simulation, which rapidly melts the structure and then allows it to cool slowly to attain an energetically relaxed state [105]. Here, a SN starts by heating the structure to 3000K, which encourages significant shifting of the atomic positions from their start points. At this point it is necessary to reiterate the fact that the starting structure for an SN simulation is an RMC refined structure. The result of the melt phase of the SN is to cause the structure to shift quite significantly from the positions imposed by the RMC. This is to allow the structure escape any potential well it maybe be trapped in by the RMC. The structure is then allowed to cool slowly, by lowering the temperature in steps and allowing equilibrium to be achieved at each temperature. This process continues until room temperature is reached (300K) at which point the simulation is stopped and the final structure taken as the snapshot.

Obviously questions arise about whether this method is appropriate for this work, and whether in fact the simulation has any effect. So does the SN method have any advantages over a prolonged simple MD simulation ran at 300K i.e. is the melt necessary? The advantage of the SN method, is the melt phase which provides sufficient energy to the initial structure to escape any potential well. The initial structure of the MD is a “good guess” of the model i.e. that it fits the experimental data. However since this structure was generated with no thought to the energy then it requires a push to get it away from the high energy states which it will be either trapped in, or situated closed to. Running MD at 300K for a prolonged period, will never move the structure far enough away from the initial structure and as such never find as low as energy state as a SN state will, and this can be seen in a comparison of the energy timelines of bother methods.

As stated, the SN method pushes the structure away from the RMC refined structure, this obviously leads to a significant drift in the experimental alignment. This is firstly an indication on the scale of the energy difference between RMC refined and MD processed structures. It also proves that the structures created by the RMC are not representative of the actual structures. This begs the question, is it possible to have both experimental alignment and energetic relaxation? The answer to this question is

obviously yes since the structures exist, and nature always seeks to achieve the lowest energy state. Achieving this in modelling terms is more challenging, one possibility however to use an iterative process. It is possible to pass the output of a SN back into RMC for further refinement and then pass that refined structure back into a SN. Given sufficient iterations this may bring about convergence between the two methods resulting a structure which is both energetically relaxed and in perfect alignment with experimental data.

So far this iterative method has not been explored with any great enthusiasm mainly due to the computation expense it represents, however the current method of modelling follows a similar tack. As explained earlier the current method is to multiply snapshot structure from the SN into a supercell and refine back to experimental alignment, which essentially represents 1.5 iterations, given sufficient resources, the process could be continued.

2.4.7 Geometry optimisation

GEOMOPT is a method used to determine the groundstate energy and structural configuration of system. This is done by first performing a singlepoint energy calculation [100], which determines the energy of the system and the forces on the structure [106]. Then by moving the atoms towards a point of minimal force (over a series of iterative steps) it is possible to achieve the groundstate configuration/energy. This process is of course highly dependent on the initial structural input. If the initial is far from the groundstate in terms of its configuration, then GEOMOPT can result in a local minimum as opposed to a global one. It is possible to improve the GEOMOPT process by first running an MD simulation to relax the structure from the initial “guess” thereby increasing the likelihood of finding the global minimum (groundstate). For this reason using GEOMOPT to complement our MD run SN simulation would further decrease the energy of the system. However due to the fact that the end of the SN is not the end of the modelling process (i.e final refinement) combined with the computational expense it was determined that a GEOMOPT was not necessary. GEOMOPT is, however, a necessity

when calculating the cohesive energy of the system since it is imperative that the system be in its groundstate.

2.5 Cohesive Energy

“The cohesive energy of a solid is the energy required to break the atoms of the solid into isolated atomic species” [107]. For the simple two body problem, the cohesive energy is relatively straight forward to calculate as the energy of one of the atoms only depends on the contribution from the other. However for large atomic structures the calculation of the cohesive energy becomes nigh-on impossible unless we assume that the cohesive energy of an individual constituent atomic is same or does not differ too much from the average cohesive energy. Using this assumption the calculation becomes relatively straight forward as it simple involves calculation of the total energy of the structure and energies and abundances of each constituent atom.

$$E_{cohesive} = E_{solid} - \sum_A E_A^{isolated}$$

[eq.2.5.1]

The calculation of the cohesive energy is described by eq.2.5.1, where $E_{cohesive}$, E_{solid} and $E_A^{isolated}$ are the total cohesive energy of the structure, total energy of the solid and energy of A-th atom. The term $\sum_A E_A^{isolated}$ for an n element compound structure expands into eq.2.5.2, where N_n and E_n are the quantity and energy of the nth constituent element. To evaluate eq.2.5.2 we must first obtain the values of E_n which involves single atom single point energy calculations [100] for each of the n elemental species. Care must be taken to determine the correct size of the simulation cell for each single atom calculation as small boxes will lead to error due to periodic boundary conditions.

$$\sum_A E_A^{isolated} = N_1 E_1 + N_2 E_2 \dots \dots N_n E_n$$

[eq.2.5.2]

Cohesive energy is experimentally measurable, and this would be an important step for future work however even lacking experimental data cohesive energy calculations are still a useful tool. Although it is impossible to know whether a cohesive energy calculation is an under or overestimate without experimental proof it is still possible to use the trends of the cohesive energy. This is because typically a particular pseudopotential will consistently over or underestimate the cohesive energy [108-112], however it can still accurately reproduce trends since the pseudopotential is essentially introducing systematic error. So by ignoring the absolute value of the cohesive energy, and instead focusing on the trends in the data we can still learn much from the calculations. It is not certain to what extent PBE over or underestimates the cohesive energy and this will have to be determined at some later date.

2.6 X-ray Reflectometry

A major constraint in any amorphous material modelling process is the density of the coating. Some allow the model to determine its own density by allowing cell parameters to vary with during the simulation. This is not an unwise approach for most modelling, however it can lead to error if the lattice is allowed to relax too much or (in some cases) too early. Primarily the work presented here is semi-empirical and as such we find it necessary to constraint the density experimentally. An obstacle exists for this work when it comes to density measurements which is that the materials studied herein are thin films and it is assumed that their density and properties are different from that of their bulk counterparts. This led to the proposed use of X-ray Reflectometry (XRR), which can be used to calculate the density of a thin film.

XRR is a non-destructive X-ray technique which allows for the determination of density, thickness and surface roughness [113-116]. The process involves using a monochromatic x-ray source at grazing incidence to the film and then monitoring the intensity of the beam reflected beam. The grazing angle ω is then varied and the x-ray

intensity will vary accordingly. At some angle θ_c the critical angle, the x-rays will begin to penetrate the thin film a down to the next surface. After the critical angle the intensity includes reflections from the top and bottom surface which interfere causing fringes, which can be used to determine the sample thickness. The angular position of the critical angle is determined largely by the density of the coating and this can be determined through modelling.

Determination of the density, roughness and thickness is usually obtained through modelling of the intensity pattern. This of course assumes these properties are constant, whether this is a valid assumption is questionable. If the assumption is invalid then of course, it would introduce error into any properties determined from the XRR modelling. The magnitude of the error would be controlled by the degree of variation in the thickness, roughness and density. It is also possible with few approximations to calculate the density without modelling. Starting from eq.2.6.1 which describes the complex x-ray refractive index \bar{n} in terms of δ and β , the dispersion and absorption respectively.

$$\bar{n} = 1 - \delta + i\beta$$

[eq.2.6.1]

For a single layer thin film on a substrate it is acceptable to assume the coating is absorption free. However it should be noted that this is not a valid assumption for modelling of the XRR intensity. By considering the reflections at the interface between air ($n_{air} = 1$) at low angles $\theta < \theta_c$ eq.2.6.1 can be expressed as $\bar{n} = 1 - \delta$. Hence using Snells law and the small angle approximations at the critical angle we can express θ_c in terms of the density ρ .

$$\begin{aligned} 1 - \delta &= \cos \theta_c \\ &\cong 1 - \frac{\theta_c^2}{2} \\ \theta_c &\cong \sqrt{2\delta} \end{aligned}$$

The critical angle which is measurable quantity is now expressed in terms of the dispersion parameter which is related to the density as per eq.2.6.2

$$\delta = \frac{r_0 \lambda^2}{2\pi} n_e$$

[eq.2.6.2]

Where

$$n_e = Z n_{Atom}$$

And

$$n_{Atom} = \frac{N_A}{A} \rho$$

Where $r_0, \lambda, n_e, Z, n_{Atom}, N_A, A$ and ρ are the Bohr radius, X-ray wavelength, electron density, Atomic number, Atomic density, Avogadro's number, Atomic weight and density of the coating respectively. This leads to eq.2.6.3, however Z is usually replaced with a complex atom form factor $\bar{f} = f_0 + f' + if'' = Z + f' + if''$.

$$\theta_c = \sqrt{\frac{r_0 \lambda^2}{\pi} N_A \frac{(Z + f')}{A} \rho}$$

[eq.2.6.3]

2.7 Chapter Summary

There are a host of different tools available to the electron microscopist both from an experimental point of view and the analytical one. Electron diffraction, specifically convergent beam electron diffraction is the most important of those tools for this project. CBED forms the basis of all analysis and modelling undertaken during this research and as

such data must be taken with the utmost care. The method of CBED implementation has been discussed here as has the subsequent analytical treatment. The method of EELS, used to determine the stoichiometry of a sample has also been discussed as this is also a valuable source experimental constraints.

Diffraction data is a very important and versatile source of experimental information and widely used in the X-ray, Neutron and Electron communities. Here, the electron diffraction data is manipulated into the reduced density function of the coating. The method of RDF calculation has been discussed here as has the approximations necessary for practical implementation. The RDF forms both the primary experimental constraint of subsequent modelling and the source of preliminary experimental results. Methods of RDF analysis will be discussed and results will be presented in the following chapters.

The multistage semi-empirical modelling process has been described in detail and the individual components in concise terms without significant derivation. The initial models are generated using experimentally determined densities and stoichiometries using XRR and EELS respectively. RMC is then used to reconfigure the structure to fit with the experimentally measured RDF. A simulated anneal using DFT MD is used to relax the structure and reduce the energy into a minimum. Following the simulated annealing, a supercell is generated from the result structure and a final RMC refinement is performed to return the structure to alignment with the experimental RDF. Following the modelling procedure various energetic and structure properties are then calculated such as the cohesive energy.

Chapter 3: Development of a Semi-empirical modelling process

3.0 Chapter Introduction

Previous work in this area developed the modelling process used in this thesis and some of the analytical tools necessary to make conclusions [64, 86, 90, 91, 105, 117, 118]. However as part of this thesis, new and improved experimental, analytical and modelling techniques will be outlined to demonstrate the evolution of the process. Improved experimental techniques have been employed to increase the quality of the diffraction data leading to more accurate modelling. Furthermore the well-established experimental techniques of X-ray diffraction and X-ray Reflectometry have been explored to add new constraints to the modelling process. New analytical techniques such as RDF analysis have provided new insight into structural evolution as a function of heat treatment and doping. Techniques such as area mapping have also shown great promise in determining the structural variability within amorphous thin films. Accurate modelling is always the goal of the quantum modeller; however accuracy must be tempered by speed. With greater accuracy comes greater computational expense and can lead to unfeasible simulation times. Here the effect of optimising the cut off energy and k-point grid on the energy of the model and its simulation cost will be discussed. Various possible future improvements to the modelling process have also been investigated and will also be discussed.

3.1 CBED improvements

The experimental method involved in RDF calculation has been well defined in the literature [64]. In short, CBED data is taken, the “background” intensity is modelled and subtracted and the resulting function is Fourier transformed into the final RDF. This is method in the ideal case and does not entirely represent the practical method employed due to various corrections and limitations in the experimental equipment. In the following pages, these limitations will be outlined as will the methods used to correct for them. Unless otherwise stated the T20 microscope is setup using a standard configuration which is a 30 μm C2 aperture, resulting in a semiconvergence angle of 1.19 ± 0.05 mrad.

Furthermore spot size 5 is used for data taking which leads to a probe full width at half maximum of 80nm in microprobe mode.

3.1.1 SIS Camera Adjustments

As previously mentioned the primary experimental apparatus used for this work was a Tecnai T20 Transmission electron microscope. In the case of diffraction studies on the T20, a SIS Megaview III camera, side mounted above the phosphor screen is used to record data. The SIS camera uses a YAG crystal scintillator and 1376*1032 pixel CCD [75]. Like most TEM cameras the SIS uses a solid substrate scintillator to convert the scattered electrons in photons [75]. This causes the first of 2 problems that our work experiences when using the SIS camera. It is the job of the scintillator to convert the scattered electrons into photons which are detected by the CCD via a photomultiplier. This means that the scintillator bears the brunt of the incident electron beam. In bright field or dark field mode this is not an issue as the beam is typically spread, reducing current density.

In diffraction mode and especially for amorphous materials most of the current arrives at one spot since most of the electrons pass through the sample unscattered. The intensity of the unscattered beam can cause serious and irreparable damage to a scintillator and the SIS camera is not an exception. This means that the unscattered beam must be blocked during diffraction imaging to prevent damage. This is done with a pointer (or beam blank) shown in figure (3.1.1.1), the use of which also causes issues, the least of which is the loss of data. Due to the side mounted position of the SIS camera, the photons must be directed on to the CCD using a simple glass prism. When the pointer is used in conjunction with the SIS, it causes a defect in left half of the camera shown in figure (3.1.1.1). It is caused by multiple internal scattering and surface scattering on the prism. This defect and the beam blank are as such deleted before further analysis, which leads to a significant loss in data, figure 3.1.1.1.




Figure 3.1.1.1 High contrast false colour diffraction image (left) showing the distortion due to the prism and beam blank. False colour diffraction pattern (right) with masking over the distortion and beam blank both images were taken at a camera length of 860mm.

The second issue faced with using SIS camera is “blooming”, which is charge leakage from saturating pixels into neighbouring ones. Blooming simply occurs when the CCD emits secondary electrons which are detected by nearby pixels. This causes an increase in intensity around high intensity areas (such as the inner rings of EDPs). This artificial increase in intensity leads to a bad fit to the intensity pattern when calculating $G(r)$ as $N\langle f^2(q) \rangle + C$ is no longer sufficient to represent the “background” intensity. Furthermore the increase in intensity at the EDP centre means that the exposure must be lowered to avoid overexposure, lowering the SNR at the extremes of the camera. The blooming can also be limited by lowering the exposure time such that the centre of the EDP does not approach saturation, however this lowers the SNR at the extremes even further.

The standard diffraction imaging conditions on the T20 for this work uses a camera length of 860mm. This provides a good q resolution but limits the maximum q achievable when the unscattered beam is centred with respect to the camera. Lower camera lengths have been explored however problems occur due to high current density.

At the 860mm camera length, blooming is an acceptable limitation; however at lower camera lengths e.g. 600mm, blooming totally dominates the central area due to high current density, making the EDPs unusable. This can be seen in figure (3.1.1.2) where blooming has obscured the first diffraction ring; this effect becomes even more prominent at lower camera lengths. Recent experiments with High Dynamic Range (HDR) imaging led to a possible solution to the blooming issue since the central area (where the blooming is most damaging) can be acquired at lower exposure times avoiding excessive blooming. HDR imaging can also significantly increase the SNR at high q which makes it an attractive option for future work.




Figure 3.1.1.2 False colour CBED image taken at a camera length 600mm, notice the lack of features surrounding the central beam. Blooming has obscured the low spatial frequency information at the centre of the pattern, making it unusable.

3.1.2 High Dynamic Range Imaging

High Dynamic Range Imaging (HDRI) is a method of image processing used to more accurately represent the intensity differences within a scene. This is done by enhancing the dynamic range of an image to allow both the lightest and darkest components of images to be seen in detail. Two main methods of HDR are computer rendering which is explained at length in the literature [119-123] and the merging of multiple low dynamic range images [124]. The latter has been employed to improve the quality of the diffraction images obtained using the SIS camera on the T20. In a practical

sense the method involves recording images of diffraction patterns at a range of exposure levels and stitching them together to create one complete HDR image. This increases the SNR of the faintest areas of the image (those at high Q) whilst maintaining accurate representation of the centre of the diffraction pattern.

High exposure images are taken such that the minimum intensity on the image is over 1000 counts (approximately 1/3 of the saturation level). This ensures a good SNR at max q whilst avoiding issues concerning non-linear response of the CCD at high exposure. These high exposure images pictured in figure 3.1.2.1 are largely over saturated at the centre due high current density at low Q. These are combined with low exposure images with no central saturation as pictured in figure 3.1.2.2.




Figure 3.1.2.1. CBED image of Amorphous Silicon exposed for 160s in false colour scheme.

This image was processed manually using 160s Gain reference and Black Reference images with the resultant image being multiplied by the maximum intensity of the gain reference (to scale the intensity correctly).




Figure 3.1.2.2. CBED image of Amorphous Silicon exposed for 10s in false “temperature” colour scheme. This image was processed automatically by the iTEM image acquisition software. Despite the automatic processing this image is still acceptable as processing induced noise is less prominent at low exposures

The Final HDR image (figure 3.1.2.4) is constructed by replacing saturated pixels (figure 3.1.2.3) on high exposure images with their non-saturated counterparts in the low exposure images. To address the scaling issue between the high and low exposure images, an identical mask is generated on each image (figure 3.1.2.3), outside of the saturation area. The pixels within this mask are used to generate the scaling factor used to calculate the intensity as described in equation 3.1.2.1.

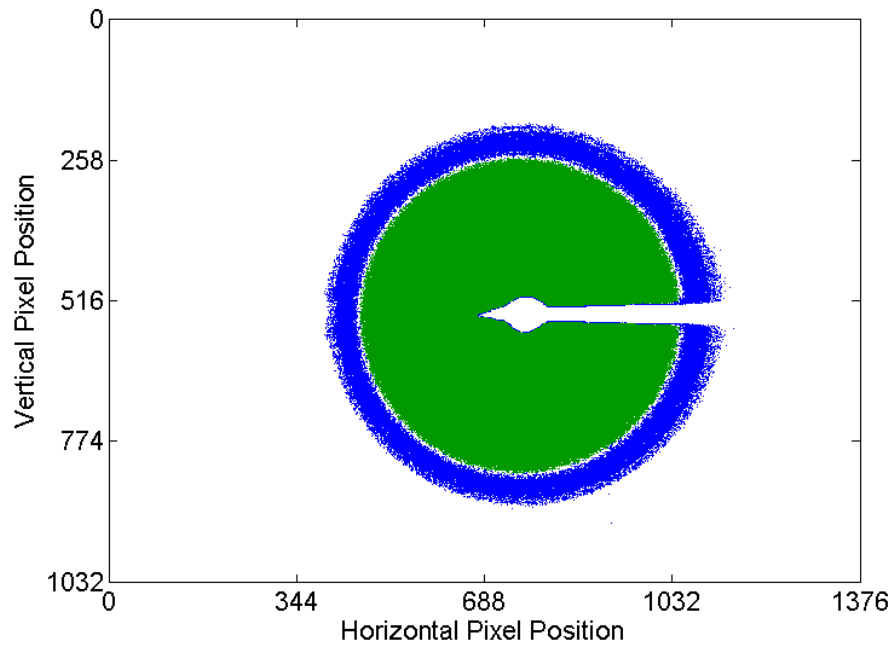


Figure 3.1.2.3. Scatter plot for the image in figure 3.1.2.1 showing the areas considered saturated (green) and the area used calculate the scaling factor (blue)

$$I_{sc} = S_c I_{rec}$$

[eq. 3.1.2.1]

Where I_{rec} is the recorded intensity of the image, S_c is the scaling factor between 2 images and I_{sc} is the scaled intensity which would have been recorded had the CCD not saturated.




Figure 3.1.2.4. HDR processed CBED image of Amorphous Silicon in false “temperature” colour scheme. This image was created using 7 images with exposure times ranging from 2-160s with the higher exposure time images being processed manually to remove any artefacts.

The images in figures 3.1.2.2 and 3.1.2.4 appear to be almost identical to the naked eye which is due to automatic gain normalisation which causes the brightest pixel to be white and the darkest to be black, scaling the rest of the image appropriately. The difference between these images can more readily be seen once the image has been radially averaged and the reduced intensity function has been calculated. This is shown in figure 3.1.2.5 in which the reduced intensity functions from figures 3.1.2.2 and 3.1.2.4 are shown alongside each other and the differences are apparent. The HDR image has a vastly improved SNR at high q , in fact its barely noticeable compared to that of the Low Dynamic Range (LDR) image. This improved SNR leads to a much more stable and much less noisy final RDF.




Figure 3.1.2.5. Reduced Intensity Functions of amorphous Silicon annealed at 450°C for 4 hours using the Low and high Dynamic Range processes. The LDR reduced intensity function shows considerably worse SNR at the extreme of the q range.

In a typical situation images obtained from the T20 using the SIS camera are processed during the acquisition stage. This involves removal of the “black” intensity and adjustments for gain imperfections. The black intensity is that recorded by the CCD which is not due to electron beam; this is dominated by (but is not limited to) random electronic noise. The black intensity is a greater issue at low exposure levels due to the inherent random nature incoming photons from external sources. Gain imperfections are due to the non-uniform reaction of the CCD due to incoming photons. This means that uniform illumination of the CCD will not produce an image with uniform intensity. This effect is more prominent at high exposure levels in automatically processed images as shown in figure 3.1.2.6 and hence it was decided to process images manually. At high exposure levels automated gain and black intensity corrections are no longer sufficient; this is because the reference images are taken at low exposures which cannot accurately describe the response at higher exposures.




Figure 3.1.2.6. CBED image of Amorphous Silicon exposed for 160s in false “temperature” colour scheme. This image was processed automatically by the ITem acquisition program. Notice the large amount of noise and artefacts caused by the sub-optimal processing this effect is most noticeable at higher exposures

The processing of the images involves recording so called “black” and “gain” reference images at each exposure level. Black images are taken with the column valves closed so that no intensity due to the electron beam is present. It is necessary to take black reference images at every exposure because the dark current that the CCD receives is random and low current. Since the dark current is random, at low exposures a black reference image is different every time. At higher exposures the noticeable difference (other than high intensity) is that the some of the random nature has been lost and the image begins to average out. For this reason a 2 second black reference image cannot be used to process a 160 second image even if the 2 second were multiplied by a factor of 80.

Gain reference images are taken with uniform illumination of the CCD using the electron beam either with the sample removed or positioned such that the electron beam passes through the hole at the centre of the sample. Gain images are necessary to gauge the linearity of the CCD response to illumination and normalised accordingly. Simply put if a thousand pixels are illuminated with the same current for the same time, the result would be 1000 different counts. Typically those thousand pixels would be within a few

percent of each other (in terms of recorded counts) which is an acceptable variance. For HDR work with very high exposure the gain reference (on the SIS) breaks down because a low exposure gain reference image cannot represent the linearity differences over long exposure times. Essentially for an image with 10 counts on average, a difference of 1 count, represents a 10% change. For an image with 100 counts on average a 1 count difference is only a 1% change. Similarly two pixels may read the same on a 1 second exposure but may differ considerably for a 100 second exposure. For this reason a gain reference image needs to be acquired for each exposure time to accurately represent the CCDs response at every exposure time.

During the data acquisition stage diffraction data, black and gain reference images are taken at a variety of exposure levels to ensure good SNR over the whole q range. The diffraction images are then mathematically manipulated as described by equation 3.1.2.2 to produce a processed image as shown in figure 3.1.2.1.

$$I_p = \frac{I_{up} - I_{bl}}{I_{ga} - I_{bl}} I_{ga}^{max}$$

[eq. 3.1.2.2]

Where $I_p, I_{up}, I_{bl}, I_{ga}$ are the processed, unprocessed, black reference and gain reference images respectively and I_{ga}^{max} is the maximum intensity of the gain reference image used to rescale the final, processed image.

It has been possible due to collaboration with the University of Warwick, to use a Gatan ORIUS SC1000 camera [76] to take diffraction data. The Orius has the ability to withstand the unscattered beam without a beam blank shown in figure 3.1.2.7, this is because it is built using a thin film scintillator [76] which does not absorb as much energy from the incoming electron as a more solid scintillator. This makes the Orius incredibly resistant to damage, it is also capable of resisting blooming [76]. The Orius also has a greater number of pixels 4008*2686 compared to the SIS cameras 1376*1032. This would

give a greater Q range at the same camera length were it mounted in place of the SIS camera. In the case at the University of Warwick the Orius was mounted at the bottom of a JEOL 2100, and at a camera length of 8cm it had a similar q resolution as the SIS at Glasgow but with a significantly increased Q range seen in figure 3.1.2.8.




Figure 3.1.2.7 High Dynamic Range image of Ti:Ta₂O₅ taken with a Gatan Orius Camera on a JEOL 2100 TEM at the university of Warwick. The camera length here is 8cm giving the image a calibration factor (ds) of $\sim 0.00305\text{\AA}^{-1}$ per pixel which is similar to the SIS camera on the T20 at 860mm camera length. A typical CBED image from the SIS camera has been scaled to match the pixel size and resolution of the Orius to provide relative scale. This gives the image an impressive q range whilst maintaining a good Q resolution.

With the larger pixel area combined with lower camera length allows for a good q resolution whilst allowing for a much greater q range. The maximum q range in figure 3.1.2.7 is $\sim 39\text{\AA}^{-1}$ however due to distortions in the image which are not visible in figure 3.1.2.7 the useable q range is $\sim 25\text{\AA}^{-1}$. When compared to the typical q range of the SIS camera $\sim 17\text{\AA}^{-1}$ the advantages of using the Orius are clear, this can be seen visually in 3.1.2.7 where a typical CBED image from the SIS has been overlapped for scale. The image in figure 3.1.2.7 has been created using HDR processing to give it a dynamic range of $\sim 10^9$. This is exaggerated because the central beam is exposed, within the useable data range the dynamic range is $\sim 10^5$ - 10^6 . The SIS camera has a useable dynamic range of $\sim 10^3$ without HDR, and $\sim 10^4$ with HDR. This difference is due the fact that the maximum

exposure time of the SIS is 160s compared to 500s for the Orius. However, even with a greater exposure time, the dynamic range of the SIS would not significantly increase as much of the camera is already saturated at a 160 second exposure (figure 3.1.2.1).




Figure 3.1.2.8 Intensity patterns for typical images from the Orius camera (blue) and the SIS (green). The maximum q for the sis at 860mm camera length is $\sim 17\text{\AA}^{-1}$ for the Orius at 8cm camera length the maximum is $\sim 39\text{\AA}^{-1}$. The maximum useable q range for the Orius is $\sim 25\text{\AA}^{-1}$ which is a significant increase over the SIS. The SNR at maximum for the Orius is also significantly better than the SIS.

In summary HDR processing from an imaging point of view, is an important step forward for diffraction studies. By allowing for a much greater intensity difference between the brightest and darkest areas weaker information can be preserved. This will lead to weaker peaks at high and low q becoming visible in the diffraction images. This newly visible information will then be carried forward into any further analysis. HDR imaging increases the dynamic range of the SIS camera by at least approximately a factor of ten significantly reducing the noise at high q .

3.1.3 Qmax increase spliced images

It is very important when calculating reduced density functions to have a large q range in the image. This is because the length of the q range translates in the resolution of the final RDF. If the q range is too small the RDF will be unresolved or may even appear blurred, worse still very short distance, prominent peaks can be cut out completely and the information lost. In figure 3.1.3.1 The effect of the q range (q max) on the positions of the 1st and 2nd peaks in a Tantalum RDF is illustrated. With increasing q range the positions converge towards their true distance reflecting the increasing accuracy due to high q . Furthermore at very low q both the first and second peaks (at 6 and 3Å⁻¹ respectively) disappear completely as the data has been cut off. A smaller, weaker, peak at 1.1Å possibly corresponding to oxygen···oxygen distances is cut off at an even higher q (~9Å⁻¹).




Figure 3.1.3.1: Position of the Tantalum oxygen (left) and tantalum-tantalum RDF peaks as a function of q max. Both plots show convergence at high q max underlining the need for high q data sets when calculating $G(r)$.

For these reasons it is necessary to obtain high q data with good SNR in order to fully resolve the peaks in the RDF (figure 3.1.3.2). As explained in Section 3.1.1 due to weaknesses in the SIS camera, high q cannot be obtained by reducing the camera length on the T20. While this is unfortunate, it is not the ideal method of increasing q range

regardless as it results in reduced q resolution. So in order to increase q max without sacrificing q resolution it was necessary to develop a method to increase q max whilst remaining at the 860mm camera length.

Toward that end two methods were considered, the first was to shift the centre of the diffraction pattern to the extreme range of the beam blank thus increasing q max. Whilst this is a simple, straight forward task, it could not produce useable results due to low SNR. As explained in section 3.1.2 the SIS camera does not possess the dynamic range necessary to provide acceptable SNR at a q of 17\AA^{-1} so increasing the q range beyond that is not feasible. This problem can of course be solved using HDRI however it was not available at the time.




Figure 3.1.3.2. Series of RDFs from the same image produced by varying q max; the RDFs appear to converge as q max is increased. The shorter distance information is affected more by the decreasing q max, this is because q is a reciprocal distance and as such high spatial frequency data is cut off first.

An alternative method was developed, which to an extent avoided the dynamic range issue by eliminating the brightest feature on the image. This is achieved by shifting the diffraction pattern off the camera so that only the outer most diffuse peak remains (figure 3.1.3.3). This allows the camera to be exposed much more than usual increasing the SNR at high q without blooming the brightest information or using HDRI. Then by calculating the centre of the diffraction pattern (relative to the image) the centred and off centred images can be effectively over lapped and analysed as one image. As with HDRI there is a scaling issue between the two images, due to different exposure times, requiring that the intensities be adjusted. This was done using a similar method as was described in section 3.1.2. It involved defining a band of pixels at some radius r to $r + \Delta r$ from the centre of the diffraction. Averaging the intensity in this band in each image and taking the ratio between the two averages to provide a scaling factor. This is similar to the scaling method used in HDRI, however this cannot be done pixel by pixel since a significant portion of the pattern is off the camera. This splicing method has the capacity to more than double q_{\max} , however this in practice was never achieved due to low SNR. It was possible to increase q_{\max} to 25\AA^{-1} from 17.5\AA^{-1} whilst still maintaining an acceptable SNR at high q .




Figure 3.1.3.3. Off centre CBED image taken at 860mm camera length, there is a distortion on the left hand side of the image which was later masked to remove its affects.

3.2 X-ray Diffraction

Although primarily work done at Glasgow has been focused on electron imaging and diffraction studies, it is important to recognise that most of these techniques were created for use with X-rays and since been adapted for use with electrons [64]. It is also important to note that X-rays have some advantages over electrons as well as some disadvantages. The main disadvantage of using X-rays is the relatively small cross section due to the fact that photons are not charged as well as the relatively high wavelength compared to electrons [72]. Due to the small cross section of X-rays, it is difficult to probe nanovolumes as the scattering signal is (practically) too weak. This forces the use of relatively large probe sizes which results in very “averaged” data from the structure. This is not always an issue (for crystal and polycrystalline studies), however for amorphous structures which inherently have little long range order, an “average” is not always the best approach. Using electrons which have a much larger cross section allows for a much smaller probe and hence nano-volumes to be investigated [117], which leads to much less averaging.

X-rays do have the advantage of higher angle scattering which allows for a greater Q resolution. Furthermore X-ray beam line setups typically allow for collection of intensity to much higher Q due to high energy x-rays and large area detectors. Higher Q leads to higher resolution on the final RDF, as well as lower noise (figure 3.1.3.1). Furthermore the afore mentioned issues with large probe sizes and “very averaged” data can be circumvented using highly focused, high intensity beam lines. However, as would be expected beam time on these specialised lines is scarce and for this work is overshadowed by the relatively convenient access to TEM time. The analysis and modelling techniques explained above are equally valid for X-ray generated RDFs and vast amounts of work have been presented in the literature [125, 126].

We were fortunate to be granted time on beam line B16 of Diamond Light Source at the Rutherford Appleton laboratory in August 2010. This beam time was used to study 500nm thick thin film Titania doped Tantalum deposited on an amorphous SiO₂ substrate.

These samples were mounted to standard SEM stubs and placed in the beam line at a grazing incidence angle of 2.9° . It was determined that an angle of 2.9° at a beam energy of 8keV was sufficient to assume that scattering from the coating dominated the scattered intensity. X-rays were collected using a large area scintillator and CCD which covered approximately 15° of the scattering angle. The detector was mounted on an arm which allowed for 99° of movement, allowing for collection over high scattering angles.

As the B16 is a non-vacuum line, and since the incident beam energy was relatively low, air scattering was an issue thus requiring long exposure times to achieve an acceptable SNR. Further despite the size of the detector, repeated exposures at multiple angles were needed to completely cover the available q range. These images were then stitched together using a similar method described in section 3.1.3 into a single diffraction pattern. Furthermore the diffraction pattern has a cyclical nature and as such it is not possible or practical to average the detector output over a significant range. It was decided that since the smallest diffraction ring was 10s of centimetres in radius, that a linear average across the axis of rotation of the detector was valid for less $< 1\text{cm}$ around the centre of the detector.

Although diffraction data was obtained in this fashion, an RDF was never calculated. As stated the detector arm was capable of collecting data over a large angular range ($\sim 100^\circ$). However due to the low energy of the beam, the range only corresponded to a q of approximately 8. As seen in section 3.1.3 it is necessary to obtain data to a much high q in order to calculate an accurate RDF. For this reason the X-ray diffraction data was considered unfit for RDF manipulation.

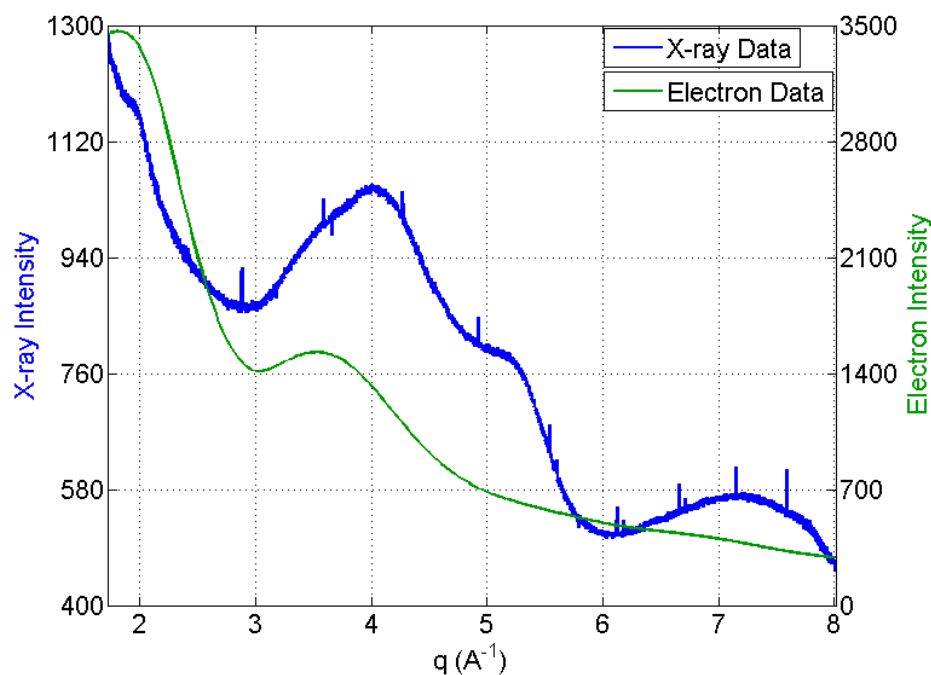


Figure 3.2.1 A Comparison between electron and X-ray diffraction data, the q range shown represents the maximum range attainable using an energy of 8KeV on beamline b16 at Diamond Lightsource. Note the splitting of peak at $4-5\text{\AA}^{-1}$ in the X-ray data and the enhanced peak at $6-8\text{\AA}^{-1}$, sadly the 2\AA^{-1} peak is almost complete obscured by noise.

Although this experiment could not lead to any accurate RDFs, it was not by any means an unfruitful venture. It is true that the low incident beam energy led a short Q range, however it also resulted in an impressive q resolution of $5 \times 10^{-4} \text{\AA}^{-1} \text{pixel}^{-1}$ compared to $2 \times 10^{-2} \text{\AA}^{-1} \text{pixel}^{-1}$ for the 860mm camera length on the T20 using the SIS camera. As seen in figure 3.2.1 the enhanced q resolution has led to a splitting of the 2^{nd} peak in the intensity pattern, a peak which corresponds to metal-oxygen nearest neighbour distances. Had an RDF been feasible from this data, it would likely have shown a similar splitting in the first peak of the RDF. It was at the time believed that this first peak was due to both Ti-O and Ta-O distances [105] but it was impossible to say without modelling or further experimental data. This x-ray data was the first indication that the theory was correct, it is not by any means conclusive as it is still impossible to say without further

work what the two peaks correspond however, at least from this work it was confirmed that there is indeed two peaks.

3.3 Improving the Quality of Reduced Density Functions

In semi-empirical modelling the greatest source of error and inaccuracy is the experimental information driving the model. It is of utmost importance to ensure that any measured data must be of the highest quality. Great care must be taken ensure the data has the highest SNR achievable within the constraints of the experimental equipment. The fitting process must then be as precise as possible to accurately describe the atomic scattering intensity leaving only the reduced intensity function. A bad fit to the intensity either through a poor fitting algorithm or poor quality data will lead to an inaccurate RDF, most likely with a large “noise” peak below 1\AA (figure 3.3.1). Typically noise peaks are more prominent at short distances because of low SNR at high spatial frequencies (high q). Bad fitting can affect short or long distance information depending on where the fit inaccurately describes the intensity (high or low q respectively).




Figure 3.3.1. Example of RDF calculated from low quality data, this represents the best fit which could be achieved to the intensity which results in a negative noise peak below 1\AA

The question then arises, what contributes a bad fit to the data and how can it be measured? There are of course many statistical means of determining the goodness of fit however the simplest option is often the best approach. Since the sub-angstrom noise peak in an RDF (figure 3.3.1 point A) is usually the most prominent indication of a bad fit it was decided to base a simple classification system on this peak. Each RDF was classified into one of 3 groups, classes 1, 2 and 3 based on the ratio of the height of the noise peak relative to first peak i.e. $R = h_{noise}/h_{first\ peak}$. RDF classifications are defined as follows, where h_{noise} and $h_{firstpeak}$ are the heights of the sub-angstrom noise peak and first RDF peak.

Class 1	$R < 0.5$
Class 2	$0.5 \leq R \leq 1$
Class 3	$R \geq 1$

These are not arbitrary limits, during the fitting process it is seen that increasing the goodness of fit beyond the class 1 limit i.e. $R \ll 0.5$ has a limited effect on the first peak (remember that noise has a greater effect on the shorter distance peaks). In the Class 2 region, the first RDF peak is affected by the goodness of fit, typically shortening at the fit worsens; therefore under-representing the abundance of the shortest nearest neighbour distance. In the class 3 region the height of the first peak is typically severely hampered by the bad fit and the width of the peak is also affected as well. For modelling purposes class 3 RDFs are rejected immediately as they would likely lead to non-physical structures. Class 2 RDFs can be used for modelling but are accepted that there is a certain ambiguity regarding the results. Class 1 RDFs are regarded to be unaffected by the noise in the data and can be used for modelling and any further RDF analysis (see section 3.3.2).

Obviously the best fit possible is always sought after but for various reasons it is sometimes not possible to achieve a class 1 fit. High SNR diffraction data usually leads to more class 1 RDFs and as previous discussed there are ways of increasing the SNR.

3.3.1 High Dynamic Range Reduced Density Functions

It has been shown that HDR processing improves the dynamic range and SNR in diffraction images. Now it is important to emphasise the improvements HDR brings to RDFs over LDR images. Since HDR images have considerably higher SNRs, it would be expected that HDR RDFs would also be less noisy than LDR RDFs. As expected this is the case, as shown in figure 3.3.1.1 where the HDR RDF has a significantly lower noise peak. It may appear that in figure 3.3.1.1 that the HDR RDF has a noise peak approximately 0.25 in height (in arbitrary units). However this peak (at $\sim 1.1\text{\AA}$) is not noise, it is most likely due to direct oxygen...oxygen nearest neighbours. These are almost certainly not oxygen...oxygen bonds, but some as yet unidentified distance possible from the planar fragments to be discussed in section 4.2.2. The sub angstrom noise peak is barely noticeable at $\sim 0.4\text{\AA}$, indicating that this was an immaculate fit. The oxygen...oxygen peak is visible in the LDR RDF as well as a shoulder to the sub-angstrom noise peak. It was not until HDRI was introduced that this peak was seen or recognised as a peak, as it is typically obscured by noise. HDR RDFs also benefit from higher (useable) q as was shown in figure 3.1.2.5 which for the reasons discussed in section 3.1.3 helps to resolve the shorter distance peaks. Under typical imaging conditions even HDR RDFs are limited to $q_{\max} \sim 18$, however it may be possible to extend this in the future. In the future combining the image splicing from section 3.1.3 and HDR processing may lead to higher q however this has not been achieved yet.

Apart from improved SNR, HDR RDFs also exhibit a greatly reduced spread between other RDFs taken from the same coating. Since a local amorphous state is non-repeatable (otherwise it would exhibit long range order) it is to be expected that there would be structural change between different areas of a material. In initial LDR work this was seen and assumed to be normal, accepting a certain amount of variance due to noise. Through HDR work a better representation of the variation in the sample has been ascertained more precisely. In figure 3.3.1.2, a series of LDR and HDR images are shown from the same sample using the same microscope conditions. The LDR images are in fact the shortest exposure images used in the processing to create the HDR images. As seen

the LDR images show a greatly enhanced spread in the peak height, FWHM as well as the sub-angstrom noise peak varying considerably. If it assumed that the variation in the RDFs represents variation in the structure and not simply measurement or processing induced error; then the question needs to be asked, what is the cause? The most likely explanation is differing amount of whatever structure motif is causing the peaks in the RDF to begin with. Since an amorphous structure lacks long range order it is reasonable to not expect the structure to “look” the same at ever point. Instead the structure as whole should be explainable in terms of an average and as such one would expect some deviation around said average. From a modelling perspective, this means a model based on a RDF will represent a possible form of the structure at that location. Alternatively, if based on an average RDF, then the model will represent a possible form of the structure on average.




Figure 3.3.1.1 RDFs from an HDR image (blue) and an LDR image (green) showing the significantly improved sub angstrom noise peak for the HDR image.

This shows that some of the spread seen in the LDR images is not due to structural variation but instead due to noise. It also leads to more accurate representation of the shorter distance peaks as high q noise more prominently affects short distance information due to their inverse relationship. As such the shoulder in noise peak in figure 3.3.1.2 which was originally believed to be also noise has been resolved. In the HDR RDFs this peak stands above the noise and as mentioned earlier is possibly due to oxygen...oxygen distances. This illustrates the need for high q data, as seen in figure 3.1.3.1 at $q \sim 18 \text{ \AA}^{-1}$ the first peak (defined as the peak at 2 \AA) is only just converging and as such the peak at 1.1 \AA^{-1} (here after called the zeroth peak) will not have converged. This means that the zeroth peak cannot be quantitatively analysed as its width, height and FWHM have almost certainly not converged. With an increased q range, perhaps $q \sim 25 \text{ \AA}^{-1}$ the zeroth may converge and be suitable for quantitative analysis as described in section 3.3.2.




Figure 3.3.1.2. A series of RDFs from HDR and LDR images of amorphous Ti:Ta₂O₅ showing the spread due to noise. Notice that the peak at approximately 1.1 Å is resolved in the HDR images considerably more than in the LDR images.

3.3.2 Reduced Density Function Analytical Techniques

The interpretation of RDFs is relatively straight forward as each peak corresponds to a nearest neighbour distance. The complexity arises trying to identify the structural features responsible for each nearest neighbour. For compound structures however, a further complication is that the constituent elements responsible for a particular nearest neighbour distance are also unknown. In the case of elemental structures this is not an issue, and instead all that is necessary is to identify the structural feature which leads to the nearest neighbour distance. Obviously the interpretation of RDFs becomes more difficult in compound structures with increasing numbers of constituent elements. This is because the number of partial RDFs increases as $n(n + 1)/2$, where n is the number of elements. For doped structures a certain amount of interpretation can be made (with an acceptable margin for error) if the parent structures RDF is known. For the case of Titania doped Tantalum, the RDF of amorphous Tantalum is known [123] and as such new peaks can be assumed to be due to the Titanium influence.

It is also possible to analyse RDFs without peak identification, in the case where structures are studied as a function of some external variable. With a large enough data set it is possible to make determinations of the change of a structure without knowing the cause of the peaks in the RDF. For instance, in the investigation of Titania doped Tantalum (in chapter 4) a data set of 6 coatings with different dopant levels was studied. This allowed dopant dependent trends to arise and estimations to be made about the change in the structure with doping. Obviously without further experimental work or modelling it cannot be said with 100% certainty what in the structure is changing. It is however reasonable to make preliminary conclusions to the cause of the changes prior to any modelling results.

In terms of what information is available in the RDF it has already been discussed that the modal nearest neighbour distances are obtainable as the position of each peak. Furthermore it is possible to estimate the relative abundances (relative to each other) of each of the nearest neighbour distances by measuring the heights of the peaks. The

coordination number is related to the area under the peaks and as such is available (however not used here) provided the area can be determined, which it is, if the peaks are modelled with Gaussians. Here however great care must be taken as the heights of the peaks are susceptible to noise [64] more so than the position of the peaks [64] which can lead to inaccuracies. It is also possible to estimate the spread in each nearest neighbour distance as the width of the peak. The width of peak being a combination of thermal and static disorder affects as well as the aforementioned resolution limit. The spread of the distances can be assumed to have some relation to the order of the structure for as in crystalline structures with strong short and long range order; there is little or no spread in the nearest distance. It is important however, not to confuse peak spread with crystallinity, while it is true that crystals have strong order at all ranges, it is not true that an amorphous structure with a narrow peak spread is approaching crystallinity. For example it has been seen that Titania doping of Tantalum increases the order (decreases the spread) of the 1st nearest neighbour distance, however it is not likely that Titania doping is causing crystallisation simply increased ordering.

3.3.3 Area Mapping

Note: To avoid confusion during the following section (up to and section 3.4) the words “depth” and “thickness” will refer to distance between the coating surface and the substrate and the thickness of the sample relative to the electron beam respectively.

As explained the RDF is a powerful tool for analysing diffraction data and determining structural change and characteristics. Like with most tools however it is only as useful as the data it is calculated from. It is known that amorphous structures change with heat treatment [105] and doping [118]. These structural changes can be determined by taking data from large heat treatment or doping based sample sets. The change in the medium range order can be determined by taking large datasets from the same sample [127, 128]. But what about the effect of thickness? It is already known that thin films and bulk structures can have quite different properties so thickness must have some effect.

Using RDFs to determine the structural change due to coating depth should be a straight forward process.

Here we define a new technique known as “area mapping” where diffraction data is taken as function of depth into the coating. Diffraction data is taken from various points on the sample and the resultant RDFs are analysed in the manor explained previously. Then by using the interface as a fixed point and the total depth of the coating, depth dependent studies can be carried out quite simply. The interface is chosen as the fixed point since it is the coating which will be eroded during sample preparation preferentially. This way structural change as a function of depth can be determined providing information of possible causes of structural change and structural properties.




Figure 3.3.3.1 Area map of diffraction study sites for the heat treated (left) and as deposited (right) samples. The red markers indicate the areas in which diffraction data was taken. The Substrate for both samples is crystalline Silicon and marked as SUB in the images.

There is of course an issue with coating depth dependent measurements which is the thickness of the area of interest. Typically the encapsulation sample preparation technique results in sample with hole in the centre, with the thickness of the sample increasing radially out from the hole. This means that sample thickness varies with the depth into the coating which is not an ideal situation. However in a typical sample a wide area of the sample is electron transparent and available for diffraction studies. This results in electron transparent areas at different depths into the coating but not necessarily in the same area.

These studies were performed on a pair of coatings of Tantalum, one with no heat treatment (as deposited) the other heat treated to 600°C for 24hours. Both these coatings were produced by MLD technologies and were slated as 1µm deep however this was experimentally determined to be in error. The total depth of the coatings was estimated via bright field TEM investigation to be ~950nm and ~1200nm for the heat treated (HT) and as deposited (AD) samples respectively. From these samples 7 and 8 (figure 3.3.3.1) different depths were sampled from the HT and AD samples respectively which is a sufficient amount of data to show good trends. The heights, widths and positions of the 1st and 2nd RDF peaks were taken and the results are shown in figures 3.3.3.2 and 3.3.3.3 for HT and AD samples respectively.

The HT sample shows strong trends in all the measured parameters indicating a significant change in the structure with depth. The heights of both peaks decrease with depth; this would suggest that closer to the substrate there is less regular structure as the nearest neighbour distances are less abundant. The ratio of the height of these peaks also changes with depth, whereas at the surface the 2nd peak is taller than the 1st, this relationship is opposite near to the substrate. The change in width of the peaks might help to explain this behaviour. The 1st peak becomes narrower with depth indicating a tighter distribution of nearest neighbour distances suggesting a more ordered structure near to the substrate. The 2nd broadens with depth, now one may be quick to think that that this indicates a decrease in order with depth however it may actually suggest the complete opposite. It has been shown previously [105] that for Tantalum the first peak is

made up entirely of Ta-O distances, but the 2nd is a mixture of Ta···Ta distances and long Ta···O. These two sub-peaks also have different nearest neighbour distances and it is their widths which cause them to overlap. Therefore if they were to narrow as is the case with the 1st peak one would see the total peak decrease in height which is seen. This narrowing would likely also decrease the width which is not seen however this behaviour maybe closely linked to the nearest neighbour distance. In the bottom plot of figure 3.3.3.2 the trend of the peak for both peaks decreases with depth suggesting a shortening of the nearest neighbour distance. If the shorter of the sub-peaks which lead to the 2nd RDF peak is also shortening with depth, this would have the effect of widening the 2nd peak despite both sub-peaks narrowing.

This is of course is speculation and cannot be confirmed without modelling; however it is a fair to assume that if the order is increasing for the 1st peak then it should for the 2nd. The above speculative analysis would suggest that the structure becomes more ordered with depth. There are a variety of factors which could lead to this, firstly that this is a heat treated sample which could affect the structure at different depths differently. If this were true, then similar analysis of a sample with no heat treatment should show little structural change (or at least a different change). Secondly the interface to the substrate could also have an effect, if it was a highly ordered crystalline material, which it is, the amorphous coating could be “attempting” to match the lattice of the crystal at the interface. In this situation heat treatment should have little or no effect and a un-heat treated sample on the same substrate should show the same trends.




Figure 3.3.2.2. Heights (top), FWHM (middle) and positions (bottom) of the 1st and 2nd RDF peaks for amorphous Tantalum heat treated to 600°C as a function of depth of the nanovolume. Both the heights and positions show clear trends with depth indicating that the peaks drop in height and move to short distances toward the interface. The trend in the peak width is less clear but it seems that the 1st peak becomes narrower and the 2nd peak broader with depth.




Figure 3.3.2.3 Heights (top), FWHM (middle) and positions (bottom) of the 1st and 2nd RDF peaks for amorphous Tantalum with no heat treatment as a function of the depth of the nanovolume. The peak heights show a trend similar to the heat treated sample, a decrease in height with depth. The FWHM and peak positions remain relatively constant (within error) with depth.

For an un-heat treated sample of Tantalum on crystalline silicon, the result of this type of analysis is quite different. As seen in figure 3.3.3.3 the heights of both peaks show a rough trend which is opposite to the heat treated sample which suggests an increase in regular structure with depth. In this context “regular structure” would refer to a structure motif (possibly multiple) which causes the peaks seen in the RDFs. Specifically, increased or decreased “regular structure” refers to the varying population of the motif(s). This is possibly due to the fact that the material closest to the substrate was deposited first and as such will have had more time during the sample deposition process to form regular structure than surface. It could also be due to atoms at the surface having dangling bonds (because there are no atoms above them) and as such cannot form the regular structure that the deeper more bulk-like material can. The peak widths and positions remain constant (with-in error) with depth which suggests no increase in order in the structure with depth. It seems that substrate is having little effect on the order of structure and that it is more likely that the trends seen in the HT sample are due to the heat treatment and not the crystallinity of the substrate. Again these results cannot lead to a definitive answer merely a hypothesis of what is likely to be changing in the structure with depth.

3.4 Coating Density Measurements from X-ray Reflectometry

Using XRR methods it has been possible to determine the density of some of the coatings studied for this thesis. Moreover in combination with previous measurements of the density of other LMA coatings the remaining densities have been obtained through interpolation. This coating run (discussed at length in chapter 4) consisted of 6 coatings with varying Ti cat. % percentages as detailed in table 3.4.1. The density of coating 0 had been previously measured in house using chemical etching [129] and was found to be 7.7gcm^{-3} . The density of coating 4 had also been measured by LMA [130] and found to be 5.5gcm^{-3} . The density of coating 2 has measured using the XRR method explained previously and found to be 7gcm^{-3} .

Coating Identifier	Ti Cat. %
0	0
1	8.26
2	19.55
3	29.13
4	53.74
5	25.69

Table 3.4.1: Ti. Cat % of LMA Titania doped Tantalum coating run as measured using EELS

With three known values of density the remaining densities were interpolated. Since the known values for coatings 0, 2 and 4 lie approximately on a straight line, a simple linear interpolation was used to determine the unknown values. Whether it is reasonable to assume a linear relationship between density and Ti % is unclear however as a simple test the density of pure amorphous Titania has been extrapolated. This value is $\sim 4\text{gcm}^{-3}$ which is slightly lower than the density of the Rutile crystalline phase of Titania ($\sim 4.2\text{gcm}^{-3}$). It would be expected that the crystalline phase would be more dense as the atoms are arranged into a close packed structure as has been with crystalline Tantalum. Therefore it was decided that the interpolated values were accurate and were later used in the modelling process (see chapter 4).




Figure 3.4.1 Density measurements for the series Ti:Ta₂O₅ coatings discussed in chapter 4.

The density for the 6.5% doped coatings was measurements using XRR. Values for the 0 and 20% doped coatings were measured using alternative methods. The remaining values were linearly interpolated from the measured values as was a value for pure amorphous Titania.

Further measurements were later taken for the remaining coatings so the densities could be determine experimentally. Data was taken for coatings 1, 2, 4 and 5 as well as coating 0, the former are shown on figure 3.4.2. However the data from these measurements was not analysed as it was determined to be in error. The basis for this can be seen in figure 3.4.2 where there seems to be a considerable variation in the critical angle for silica peak. It is assumed since these coatings were produced by the same vendor, at approximately at the same time that there should be a high degree of reproducibility. This means that the silica layers in each sample should be essentially the same and as such show the same results in XRR which is not the case.




Figure 3.4.2. Raw XRR data for coatings 1, 2, 4 and 5 showing the critical angles of silica and $\text{Ti:Ta}_2\text{O}_5$ layers at $\sim 0.2^\circ$ and 0.32° respectively.

If this assumption is followed then it is logical to attribute any variation in the silica critical angle to changes in the experimental setup. It would then be possible to essentially recalibrate the experiment by fixing the silica critical angle to some value and rescaling the data to reflect that change. This was done and preliminary results are shown in figure 3.4.3, where a problem can be identified. Visually it can be determined that the critical angles of 19.6 and 25.7% doped coatings are lower than that of the 53.7% doped coating. Since the densities of crystalline Tantalum and Titania are known to be 8.6gcm^{-3} and 4.26gcm^{-3} (in its rutile phase); then it is fair to assume that the addition of Titania will lower the density of a Tantalum coating. From equation 2.6.3 it is known that $\theta_c \propto \sqrt{\rho}$ hence it would be expected that the 53.7% coating would have the lowest critical angle. This is not the case and by extension, means that the previous assumption, that by fixing the silica critical angle any experimental discrepancies are accounted for is not valid. This leads to the conclusion that there are other experimental variables affecting these results which need to be accounted for during the experiment.




Figure 3.4.3. XRR data for coatings 1, 2, 4 and 5 rescaled such that the critical angle for silica is fixed.

Experimental determination of coating density for modelling is an important and necessary advancement and in the future XRR will provide this. It happens that the results collected from these coatings were not fit for this purpose. A possible cause of the experimental discrepancy was the inability to correctly determine the height of the sample during calibration. The experimental set up used a fixed height clamp to hold the samples in place during measurements. However the samples were not all of the same height and this had to be adjusted for, manually. In the future, a motorised stage will ensure an identical sample height for every measurement. This would remove some of the experimental error and increase reproducibility and allow for accurate density measurements.

3.5 Improving Density Functional Theorem Simulations

DFT simulations depend on a great many parameters which affect accuracy, precision and the rate of convergence. When using a plane wave basis set as is the case here, the choice of the cut off energy is very important. The kinetic energy cut off is used to reduce the infinite plane wave basis set to a finite one to make calculations feasible. This obviously leads to an error in the final energy since some of the information is lost from those planes waves which are ignored. By careful choice of the kinetic energy cut off the error can be minimised but not eliminated. Typically the energy of a system will converge towards a final value as the cut off energy is increased (figure 3.5.1). Computational expense of a simulation increases with cut off energy, i.e. more plane waves, more work; so simply setting the cut off energy to a high value is insufficient. However by setting the cut off energy to the cusp of this convergence, both accuracy and expense concerns can be alleviated.




Figure 3.5.1. Total energy of a 100 atom, Tantalum structure as a function of the kinetic energy cut off.

Further parameter which has a serious impact on the accuracy is number of k-points, those sampled in the first Brillouin zone. In reality there are an infinite number of these points and as such an infinite number of calculations. In practice a finite number of k-points are sufficient to determine the total energy of solid. This is valid since the electron wavefunctions will be near constant for neighbouring k-points, thus the wavefunction at both can be represented by that at either one. The optimum number of k-points can be determined in a similar fashion as the cut off energy, by increasing the number and taking the value at which the energy converges.

Typically when preparing to run DFT models, the first step is to “optimise” the model, i.e. determine the optimal cut off energy and k-points. Previous work in this area [105] used a cut off energy of 200eV and a single k-point, known as the gamma (Γ) point. Both of these values are sub-optimal, they were chosen simply to reduce the computational expense. The optimal cut off energy and number of k-points was found to be 600eV and 8 (in a 2x2x2 grid). This represents (for a given system, pseudopotential, XC functional) an effective increase in computational expense of a factor of ~ 200 over the sub-optimal parameters.

Arguments can be made to whether optimised parameters are necessary for this particular modelling process considering the final refinement step. The final step of the modelling process, in which the DFT model is refined back into experimental alignment, will adversely affect the energy state. It is accepted that this will happen, but that it can be reduced by limiting the maximum movement of each atom so they stay in relatively the same positions. For this reason, it might not be necessary to ensure a true groundstate energy prior to refinement.

If the optimised parameters lead to a model which is structurally closer to the experiment then the increased expense can be justified. By starting the final refinement from a model whose RDF more closely matches that of the experiment, fewer iterations will be required to achieve convergence. Fewer iterations will further limit the detrimental effect refinement has on the energy of the structure. As can be seen in figure

3.5.2 the optimised model does in fact more closely match the experimental RDF. This is especially prominent in the 2nd peak at $\sim 3.5\text{\AA}$ which both more closely matches the shaped and the position. For the un-optimised DFT RDF the 2nd peak has split into two separate peaks at ~ 3 and $\sim 3.8\text{\AA}$ respectively; the former almost certainly corresponding to an overestimation of the abundance of metal-metal distances.




Figure 3.5.2. A comparison of RDFs from experimental (green), un-optimised simulation (blue) and optimised simulation (red) sources. The optimised DFT RDF shows a closer agreement to the experimental RDF especially in the 2nd peak at $\sim 3.5\text{\AA}$.

It has been shown that using optimised parameters during the DFT calculations leads to a structure which more accurately represents the experimental RDF (figure 3.5.2). Furthermore it has also been shown that by using optimised parameters, the calculated energy converges to the true energy of the structure (figure 3.5.1). Finally the energy

evolution of the system follows a similar path to the un-optimised system but ends at a lower energy state (figure 3.5.3). The energy related results were not unexpected however, the fact that the RDF of an optimised simulation more closely matches the experimental compared to the un-optimised simulation was unknown.




Figure 3.5.3. Energy evolution for two models with the same starting configuration. The optimised model (blue) used a cut off energy of 600eV and 8 k-points, the un-optimised model used a cut off energy of 200eV and 1 k-point.

There are of course disadvantages of using optimised parameters mainly (as mentioned previous) due to the computational expense of running these simulations. CASTEP simulations scale linearly with k-points, cubically with the number of atoms and approximately cubically with cut off energy. This means that the by increasing the k-points from 1 to 8 and the cut off energy from 200 to 600eV results in a factor ~ 200 increase in computational expense. This makes a simulation which would take several days to run, take many months. CASTEP is very well parallelised and as such some of this time can be recovered by increasing the number CPU cores used. Still a simulation will

take a few months, assuming the computational resources are available for that amount of time.

In summary, the scientific benefits of using optimised parameters in modelling are clear. Optimised parameters lead to lower energy states which more closely match the experimental data allowing for more accurate analysis and further modelling. However un-optimised parameters allow for models to be created on a much short timescale at the cost of absolute precision. With sufficient access to high performance computational resources, optimised models should be prioritised over un-optimised ones. However If computational resources are limited, then the un-optimised parameters allow for models to still be created in the knowledge that any results are not as accurate as in the optimised case.

3.5.1 Decreased Simulation time

In the literature [105, 129] DFT simulations of the type used here are often allowed to run for a specific amount of time. This is to allow for the structure to reach equilibrium at each stage before continuing onto the next. Usually each stage is given a significant amount of simulation time, several picoseconds or thousands of iterations. This method, whilst simple and likely to achieve a stable equilibrium is time consuming. Time consumption is a minor detail for small simulations or where large computational resources are accessible. However for large systems, or where significant computational resources are unavailable, reducing the computational expense (and thus the simulation time) becomes very important.

To reduce the total time of a simulation one does not simply blindly reduce equilibration time at each step. Doing so would likely lead to the structure becoming trapped in a high energy state, this is especially important at high temperature as the thermal energy is required to shift the structure away from its starting point. However by observing the energy evolution over time, it is easy to determine when the structure has reached equilibrium. At this point the simulation is stopped, the parameters adjusted,

and then allowed to continue thereby reducing the total simulation time by avoiding hundreds or thousands of unnecessary iterations after the structure has reached equilibrium.

When using a method like this, there is a danger that the simulation will be stopped before reaching equilibrium and the resulting structure will be in a high energy state. As such this method was tested before implementation in future modelling by running two simulations of the same structure, from the same starting point and following the procedure except with a shortened equilibration time. The two simulations 'fast' and 'slow' were run using 100 atom model of Tantalum using an 8 stage SN process. The SN started at 3000K it was then dropped to 2000K, then at 300K steps to 500K, and finally to 300K (room temperature). The slow simulation was allowed 5ps to achieve equilibrium at each temperature; the fast simulation was allowed only 2ps. Both simulations followed almost identical paths toward their final structures and achieved the same final energy (figure 3.5.1.1). This proves that the extra 3ps per temperature step the slow model was allowed did not improve the simulation i.e. did not allow the structure to achieve a lower energy state. It was still necessary to perform over 9000 iterations in the fast simulation; however, this is, of course, a significant improvement over the slow simulation.




Figure 3.5.1.1 Total energy evolution of two separate simulations which differ only the amount of time the system is allowed to find achieve equilibrium at each temperature. The 'slow' simulation (green) was allowed 5ps (2500 iterations) at each stage, the 'fast' simulation (blue) was allowed 2ps (1000 iterations) at each stage.

Total simulation time can also be reduced by increasing the size of the temperature step thus decreasing the number of steps. Again this effect was investigated running a simulation using a 2 stage SN 3000K then immediately dropping to 300K. Of course this is an extreme example and cannot possibly represent the cooling process it did prove that the slower cooling process does lead to a lower energy structure. As seen in figure 3.5.1.2 the 2 stage SN simulation ends in a much higher energy state than the 8 stage SN process. This means that reducing the number of temperature steps has a significant (negative) effect on the final structure and as such should not be used to reduce simulation time .




Figure 3.5.1.2 Total energy evolution of two simulations with a different amount of stages in the SN process. The fast simulation (blue) from figure 3.5.1.1 uses the 8 stage SN explained early where as the other simulation (green) uses a 2 stage SN.

It was important to minimise the computation expense of the DFT modelling process as computational resources were at a premium. The majority of modelling results present in this thesis have been performed on ScotGrid, a cluster situated in Glasgow. Due to technical limitations of ScotGrid whilst running multiprocessor computations, it was not possible to run more than one simulation at a time. Were two simulations ran in parallel, they would unfortunately run at <50% efficiency compared to a single simulation. This meant that simulations needed to be run in series rather than parallel which introduced time constraints. This led to the need to reduce total simulation time by altering the method outlined by Bassiri *et al.*[105].

3.5.2 Molecular Dynamics simulations vs Geometry Optimisations

As explained in chapter 2, molecular dynamics (MD) and geometry optimisation (GeomOpt) are distinct processes with advantages and disadvantages. GeomOpt follows a path to the minimisation of the forces acting on the structure and as such, guarantees a minimum energy state. However it does not guarantee the global minimum only a minimum. This makes GeomOpt highly susceptible to the state of the initial structure, if it is close to the global minimum then GeomOpt will find it, if not it will most likely become stuck in a higher energy minimum. The advantage of MD or specifically SN MD is that it increases the chance of finding the global minimum (or at least a low energy minimum). This is because it is first given sufficient thermal energy to escape any high energy states and given time to explore the energy landscape. The disadvantage of SN MD is that it is computationally expensive and thus typically takes much longer than GeomOpt to complete.

Geometry optimisation is not unable to achieve the global minimum it is merely limited by the starting structure. Therefore if the initial structure is quite close to the groundstate, then GeomOpt would be more advantageous than MD. Moreover since this modelling process is semi-empirical, the initial structure used in the DFT simulation (whether it be GeomOpt or MD) should be a 'good guess' since it matches experimental data. By optimising a structure and putting it through a SN MD simulation separately this can be tested. If the initial structure is close to the global minima the geometry optimisation should lead to a lower energy structure than the MD. This is because the MD simulation has thermal energy, however a further geometry optimisation of the final MD structure would lead to the same as the straight geometry optimisation.




Figure 3.5.2.1 Energy evolution of a geometry optimisation (green) and molecular dynamics simulation (blue). Both simulations used the same input parameters and input structure of amorphous Tantalum.

For a 100 atom system of amorphous Tantalum heat treated to 600°C the geometry optimisation and MD simulated anneal produce very different structures. From figure 3.5.2.1 it can be seen that the geometry optimisation results in a structure with a much higher energy than the MD simulation. This indicates that the initial structure is far from the global minimum and as such the optimisation led to a high energy state. This effect is compounded by the fact that the MD energy can be further decrease by optimising its final structure. This will be explored further in the following section in which the possible advantages of using geometry optimisation in conjunction with molecular dynamics will be discussed

3.5.3 Molecular Dynamics with Geometry optimisation

It has been established now that geometry optimisation is insufficient to achieve groundstate energy for these amorphous structures. It is however still a powerful tool for modelling purposes and methods incorporating geometry optimisation will be discussed further here. The purpose of the DFT phase of the modelling process is to transition the structure from a state matching the experimental data to an energetically realistic state. MD allows the structure to escape a higher energy minimum and approach a lower energy state. Then by optimising the final MD structure it can be dropped into a local minimum (though not necessarily the groundstate) to further reduce the energy. In figure 3.5.3.1 this has been done for the structure described in section 3.5.2 to illustrate the further reduction in energy.




Figure 3.5.3.1 Energy evolution of 100 atom amorphous Tantalum structure where the final MD structure has been optimised. The time axis for the geometry optimisation is arbitrary, since the each successive iterative is an atomic step and not an time step as is the case with the MD.

Optimisation of the structure after the MD simulation should lead to a more accurate final structure after the final refinement step has been completed. It also further improves the energetic state of the structure from the MD; this provides a better starting position for the final refinement. It is accepted that by refining the structure back in experimental alignment after the MD simulation the energetic state will degrade i.e. that the structure will return to a high energy state. By starting in a more energetically favourable position, less of the energetic improvement should be lost during the refinement process. This should lead to a structure in a more energetically favourable state following the final refinement stage than a structure which was not optimised prior to refinement. Regrettably this cannot be confirmed as the structures produced during final refinement have ~2700 atoms making DFT calculations very computationally expensive.

3.5.4 Iterative Modelling Method

It has been previously mentioned that although the theoretical structure matches the experimental RDF at the RMC stage, it tends to drift away from this during the DFT stage. This is because during the DFT simulation the structure is given significant thermal energy which leads to a significant shift in the atomic positions. The degree of this drift is illustrated in figure 3.5.4.1 where the experimental, RMC and DFT RDFs are compared; note that the initial structure used in the DFT simulation was the RMC structure. Firstly neither the RMC nor the DFT structures match the experimental RDF below 1Å, however this is not inaccurate since the data below 1Å is non-physical noise.




Figure 3.5.4.1 RDFs of 25.7 Ti cat % doped Ta_2O_5 coating from the experiment (blue) the RMC generated structure (green) and the structure after the DFT MD simulation.

After the DFT stage of modelling, the structure is in a low energy state but as seen in figure 3.5.4.1 can no longer represent the experimental data. This means that the DFT structure is not representative of the true amorphous state of the coating. It has been shown previously that optimised parameters can improve the fit between the experimental RDF and that of the post DFT structure. However even with optimised parameters the fit is by no means perfect

In an ideal state the final theoretical structure of a coating would both match the experimental parameters and be in an energy optimised configuration. So far this has not been achieved through a single iteration of RMC and DFT. However by iterating between RMC and DFT simulations it may be possible converge the structure to such a state. This process, if successful, would create a structure more likely to represent the experimental structure. A limitation of this process however, is that DFT calculations are computationally expensive and that expense depends heavily of the quantity of atoms

simulated. For heavy elements such as Tantalum a 100 atom model is feasible but still expensive. Simply increasing to 200 atoms whilst being preferable would also increase the computational expense (using CASTEP) by a factor of 8 making the simulation unfeasible. This means that the final stage of the modelling process in which a 3*3*3 supercell is generated cannot be undertaken as it would make a DFT simulation would be unfeasible. Classical Molecular dynamics [REF!] could be the answer allowing for larger box sizes and many constraints and is currently being employed within the LSC collaboration and as is such is not here. Furthermore it would most likely move the structure away from the ideal state (once achieved) and as such it should not be done after the iterative process. This will limit the maximum number of atoms in the final structure to ~100 atoms.




Figure 3.5.4.2 The effect of an increased number of RMC steps on the fit between the theoretical and experimental RDFs.

In an attempt to maintain the structural features formed during the DFT stage, the following RMC simulation is performed using a reduced maximum atomic step size (0.001Å compared to 0.1Å). This would reduce the probability of the structure straying far

from its initial energy optimised state. This will also slow the iteration process and possibly lead to a vast number of iterations be required to converge to experimental alignment. As seen in figure 3.5.4.2. after 10^5 iterations the fit between the experimental and theoretical RDFs is good. However, after 10^6 the fit has improved considerably, now almost perfectly matching the experimental data.

In this state the structure will have drifted from the energetically relaxed state from which is started. In all likelihood the structure will drift further with more RMC iterations. In actuality the structure which was produced after 10^6 RMC iterations could not be geometry optimised as it was “too far” from an optimised state. The structure produced from 10^5 RMC iterations however, could be, and was optimised. The result of this optimisation was, again, a structure in an energetically relaxed state. This state shared a similar energy to the state the structure had been in following the original optimisation performed prior the 2nd round of RMC. As seen in the figure 3.5.4.3 the structure also shared a similar state in terms of its RDF.

This means that the 2nd round of RMC, intended to “tweak” the structure back into experimental alignment had undone the original optimisation. The RMC also left the structure in a state close to the final, energetically optimised structure which is why it the optimised back into said state. Overall, the result of this iterative method has shown that the processes of RMC and DFT have the opposite effect on the state of the structure. This means that even with repeated iterations the structure wouldnt be able to achieve a state which was both energetically relaxed and able to reproduce the experimental RDF.




Figure 3.5.4.2 The result of a geometry optimisation performed on a structure refit to the experimental RDF after a simulated anneal simulation.

This does not mean that the use of RMC is unnecessary since RMC provides a valuable initial structure which will, significantly, reduce the simulation time required to achieve an energetically relaxed state. Moreover it is not necessarily true that the experimental RDF is representative of the structure. There is of course, still a degree of inaccuracy in the experimental RDFs due to, low SNR, finite q_{\max} and invalidity of approximations such as single scattering and small angles. However it is also possible the experimental RDF is only representative of the structure's current configuration and not its "final" configuration. It can be said that an amorphous state is metastable, its structure is evolving in time simply at a relatively slow rate. However DFT simulations cannot account for this effect and should only be able to represent final 'stable' states of the structure. For these reasons, the fact that the DFT relaxed structures do not match the experimental RDFs is not entirely inaccurate.

3.6 Chapter Summary

Semi-empirical modelling, has the advantages over purely theoretical modelling as it is guided by direct experimental evidence. However, this strength is also a weakness as the modelling process can be lead astray by inaccurate data. It is therefore of utmost importance to use only high quality data to contain any modelling efforts. Various methods to improve the quality of CBED data have been investigated in this chapter. Factors such as SNR, q_{\max} and q resolution have to, and have been considered for these methods. HDRi has been shown to significantly improve SNR across the q range, reducing noise and increasing the accessible q range. Extended range CBED can significantly increase max q without sacrificing q resolution however, it requires HDRi in order to do so with an acceptable SNR.

Various X-ray techniques have also been investigated for the purposes of improving the modelling process. X-ray diffraction, an analogous, highly used technique can provide high q_{\max} and good SNR without sacrificing q resolution. However facilities with the correct experimental setup are rare and access is limited. Moreover, X-ray diffraction cannot (easily) image nanovolumes as readily as electron microscopes due to their longer wavelength and lower reaction cross section. A further technique of interest is X-reflectometry, with this the coating density can be determine experimentally adding a further constraint to the models. Whilst initial forays into XRR proved unfruitful, this has been attributed to operator error and lack of proper mounting platform on the experimental setup which was used.

New analytical techniques have been explored to provide new, experimental insights into structural alteration due to chosen variables e.g. doping. RDF analysis techniques which measure various properties of RDFs, such as peak positions can support and complement modelling results. It is important to emphasise the importance of high quality, high q_{\max} data for RDF analysis as they are susceptible to noise and convergence effects.

Different aspects of DFT modelling have also been investigated for the purposes of increasing accuracy and decreasing simulation expense. Simulation time can be significantly reduced simply by reducing the equilibration time allowed at each temperature step. This does not unduly affect the final state of the system in the tests, however, it will reduce the possibility of slower processes from occurring. Geometry optimisation has been shown to improve the energetic state of the structures if used in conjunction with MD as would be expected. It cannot take the place of MD, however, it works well as a complementary process. A new form of “iterative” modelling has also been investigated using RMC and Geometry Optimisation in tandem which was not able to produce the desired result, i.e. a model which fits the experimental data is in an energetically relaxed state.

Chapter 4: The effect of Titania doping on Tantalum.

4.0 Chapter Introduction

Advanced LIGO is slated to use 30% Titania doped Tantalum and Silica multilayer coatings for the reflecting surface of the test masses [30]. It is known that Titania doping reduces the mechanical loss in Tantalum at room temperature and that the reduction peaks at 30% doping [63]. The mechanism responsible for mechanical loss and thermal noise is not fully understood [43]. The current hypothesis is that mechanical loss is due to structural transitions, e.g. bond flipping [131, 132]. For the mechanical loss of silica it was suggested that a change in the bond angle of a Si-O bond was responsible for the loss [131, 132]. While silica will be used in aLIGO, it is not the dominant source of coating thermal noise, which is the Ti:Ta₂O₅ layers [63]. For this reason, Ti:Ta₂O₅ has been studied at length to measure, determine the cause of and develop methods to reduce mechanical loss [63, 105, 118, 129].

Previous work in this area has shown that heat treatment also affects the loss of Tantalum [43, 58, 59] however only subtle changes are seen in studies using electron microscopy [105, 129]. This is possibly because the heat treatment alters the medium range order more than the short range order to which our RDF studies are sensitive. A similar effect has been seen for amorphous Silicon in which short range order is near constant with heat treatment [127], whilst the medium range changes significantly. Titania doping however, has been shown to significantly change the short range order and causes significant changes in the RDF analysis [118]. In the work described in this chapter, Diffraction studies have been performed on 6 different Titania doped Tantalum coatings with varied dopant levels detailed in table 1 (section 3.4). These coatings were produced by LMA, a company specialising in high performance materials. As such this series of coatings will often be referred to as the LMA coatings hereafter. RDFs have been calculated, and the subsequent analysis and modelling are present here.

4.1 Experimental Analysis

The addition of Titania to Tantalum has been shown to change the mechanical loss [63], the density (section 3.4) and Young's Modulus [133]. This indicates (not unexpectedly) a change in the structure of Tantalum with the addition of Titania. It might also be expected that this structural change, which also leads to significant changes in the density and the mechanical properties would lead to significant changes in the diffraction data. This is however, not the case as seen in figure 4.1.1 where diffraction patterns from the 6 different doping levels show little obvious difference other than relative intensity, which is due to varying exposure times. It has been shown in the literature that the nearest neighbour distances of Ti-O and Ta-O are very similar with short (1.994Å) and long (2.039Å) [134]. This means that it is possible for Titanium atoms to replace Tantalum atoms in Tantalum without significantly changing the structure. Whilst this may lead to a subtle change in the diffraction data, this may not be noticeable by eye, even if it becomes apparent through analysis of a RDF.

Where the diffraction pattern does not appear (from a casual glance) to change with doping, the same cannot be said for the RDFs. Even from a qualitative stance, the RDF can be seen to change considerably from undoped Tantalum to 53.7% Titania doped Tantalum. This can be seen in figure 4.1.2 where RDFs for each dopant level are shown sequentially downwards. Again in a qualitative sense, variation in peak height can be seen in the 1st and 2nd peaks at ~2 and ~3.5Å respectively. With closer inspection, it can be seen that the 1st peak drifts to the right as the doping level increases which would indicate a slight lengthening in the modal nearest neighbour distance. Furthermore, the 2nd peak develops a prominent shoulder on the left side as the doping level increases, suggesting a separation of two superimposing peaks. Care must be taken here as the resolution of the RDFs is ~0.4Å and as such any shift in peak position may be an effect of this. Furthermore this resolution will make separation of the two peaks of Ti...O and Ta...O unresolvable unless it is significant. Unfortunately barring the RDF of the 29.13% doped coating, none of the RDFs show the peak at 1.1Å seen in the HDR RDFs in figure

3.3.1.1. This is most likely due to the limitations explained earlier due to low SNR and low q_{\max} .




Figure 4.1.1: Diffraction patterns for Ti:Ta₂O₅ ranging from 0% (top left) to 53.7% (bottom right) doping. These images were taken using a standard set of microscope settings, 860mm camera length, $\alpha = 1.19 \pm 0.05$ mrad and spot size 5.

While qualitative analysis provides preliminary results, quantitative analysis of the RDF is needed to fully reveal the nature of the change in the nearest neighbour properties. There are two main peaks in the Ti:Ta₂O₅ RDFs at ~ 2 and $\sim 3.5 \text{ \AA}$ respectively, and any shorter distance peaks are not considered as the q range is insufficient to properly resolve them (see section 3.1.3). Longer distance peaks are also not considered as they are a result of the superposition of many peaks making analysis problematic. For

these 6 LMA coatings, the heights, positions, widths were determined as well as the ratio between the 1st and 2nd peak height (to determine relative abundance).

In terms of the nearest neighbour distance, there is a general increase for both peaks up to 30% doping, decreasing afterwards (figure 4.1.3). As seen earlier the mechanical loss of Ti:Ta₂O₅ is minimised at 30% doping, this suggests that the nearest neighbour distance is closely related to the mechanical loss. The results suggest that an increased nearest neighbour distance leads to a lower loss structure. The increased distance could be an indication of weaker bonding, if so the structure would be more flexible. Greater flexibility would allow the structure to move toward homogeneity by lowering the energy barriers between different configurations. This move towards homogeneity might naively be expected to be accompanied by a decrease in the width of the peaks, as the spread of nearest neighbour distance would decrease. This behaviour is indeed, seen for the first peak in the RDF (figure 4.1.3), however the 2nd peak behaves differently. For the 2nd peak the spread decreases as expected up to 20% doping and then increases quite rapidly afterwards. This could suggest a decrease in homogeneity after 20% doping, however considering the behaviour of the 1st peak this is unlikely. It is more probable that as for the area mapped Tantalum (section 3.3.3), the constituent partial RDF peaks are separating and, as such, widening the peak in the RDF. This is supported by the peak height, which increases up to 20% and then decreases rapidly afterwards. The effect can also be seen visually in figure 4.1.2 as the development of a prominent shoulder on the peak indicating separation.

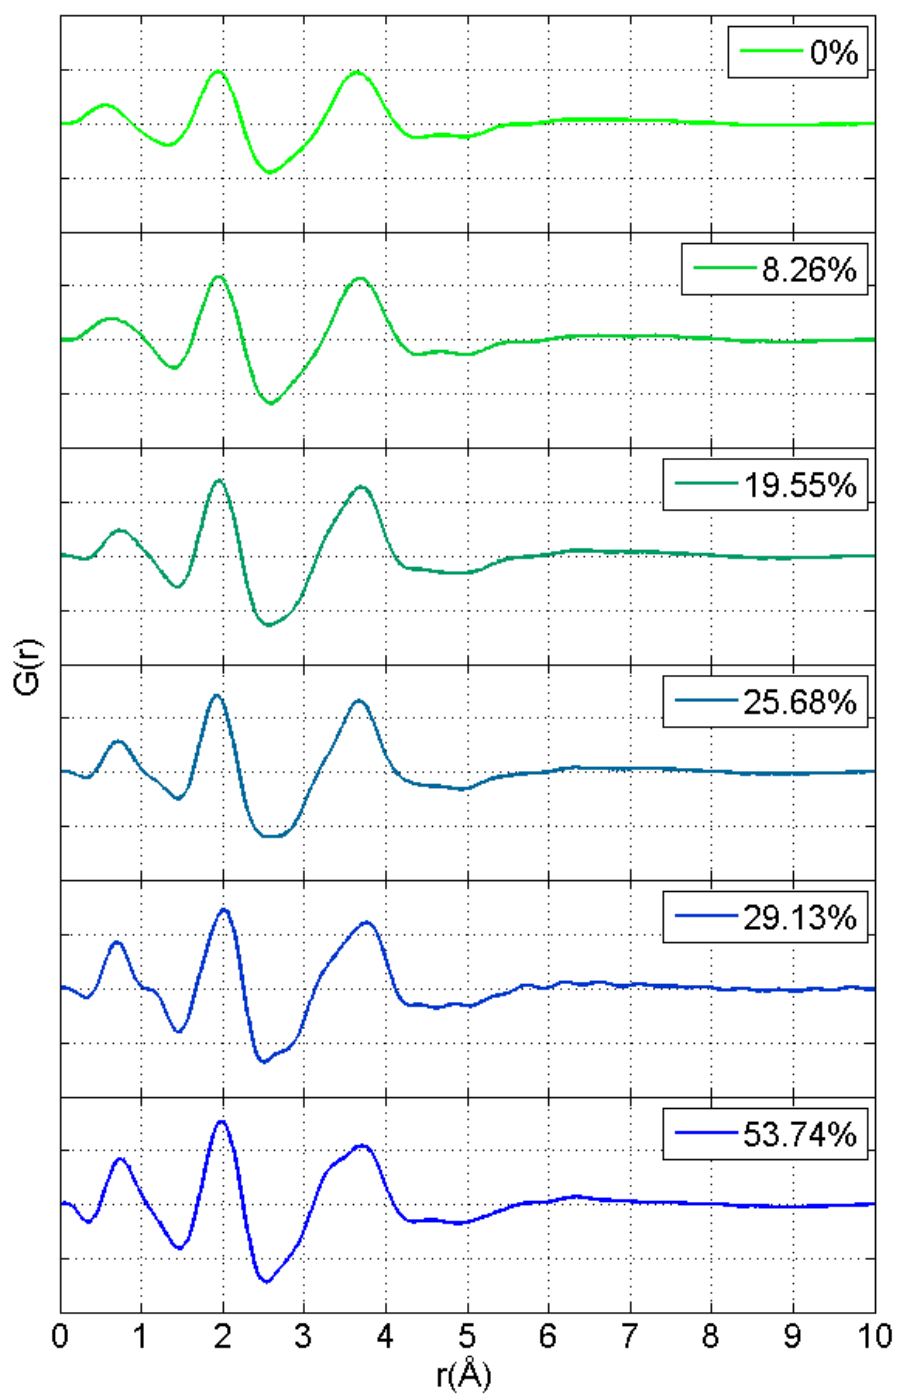


Figure 4.1.2 RDFs from the 6 Ti:Ta₂O₅ coatings, the Y axes are dimensionless and have an arbitrary scale and thus are not shown, however each plot has same Y axis.




Figure 4.1.3a. Peak Positions for the 6 LMA Titania Doped Tantara coatings against the Ti cat. %




Figure 4.1.3b. Peak Heights for the 6 LMA Titania Doped Tantara coatings against the Ti cat. %




Figure 4.1.3c. Peak Widths for the 6 LMA Titania Doped Tantalum coatings against the Ti cat. %

From the RDF analysis a clear pattern has emerged on the change in the structure Tantalum as it is doped with Titania. As the doping level is increased, the structure moves towards homogeneity and the preferred nearest neighbour distance increases. At some point between ~30 and ~55% doping, this trend reverses and the structure moves away from homogeneity. This is possibly due to a change in the dominant structural influence, i.e. that Titania becomes dominant over Tantalum at high doping.

A key result here is the behaviour of the width of the 2nd peak in the RDF as seen in figure 4.1.3. As stated it seems that peak decreases with doping indicating a move toward homogeneity. This trend is also seen in the room temperature loss as seen in figure 1.4.3.1. When shown alongside each other as in figure 4.1.4, the correlation between them becomes clear. This seems to indicate that room temperature mechanical loss benefits from a structure closer to homogeneity. This is possibly due to increased flexibility as postulated earlier which would allow the structure to reconfigure at lower energy cost thus dissipating less energy overall.




Figure 4.1.4. Room Temperature mechanical loss for the 6 LMA coatings as a function of Ti cat. %. Also plotted is the FWHM for the 1st peak ($\sim 2\text{\AA}$) as a function of the Ti Cat. %. The absolute values of these quantities are not related and as such they are shown on arbitrary Y-axes.

A series of strong trends outside of error have been seen in the experimental data from the LMA coatings. These trends have provided the basis for postulations regarding the change in the atomic structure of Tantalum as it is doped with Titania. It seems that the structure moves towards homogeneity at longer nearest neighbour distances and that this process results in less dissipative structure. It is possible that the decreased dissipative nature of the structure is due to increase flexibility and thus a lowering of energy barriers between states. These are preliminary findings based on experimental data, however it would be expected that atomic models will express the same properties as they are guided by the experimental data.

4.2 Modelling Analysis

4.2.1 Analysis of models upon completion of the DFT stage of modelling

Atomic models of the 6 LMA coatings were created using the process described in chapter 2. The density of the models was fixed to the values taken from the linearly interpolated densities in table 3.4.1. The stoichiometry was determined experimentally as described above and the results of which are given in table 4.2.1.1. The number of atoms for each structure was intended to be 100 however, due to stoichiometric constraints this was not always possible and the true number of atoms is given in table 4.2.1.1. The keen observer will note that the stoichiometries presented in table 4.2.1.1 indicate that the materials are oxygen deficient. The calculation is simply $2.5 \cdot \text{Ta}\% + 2 \cdot \text{Ti}\% = \text{O}\%$ (5 oxygens for every 2 tantalum in Tantalum and 2 oxygens for every titanium in titania) and shows coatings 4 and 5 are particularly deficient. Whilst the question of whether these coatings are truly oxygen deficient (and if so, why?) or whether this is an artifact from the EELS remains open; however, in the absence of any other stoichiometry measurements these were assumed true and used in all further work.

Coating Identifier Code	Oxygen %	Tantalum %	Titanium %	No. of Atoms
0	0.688	0.312	0.000	102
1	0.680	0.294	0.026	99
2	0.677	0.260	0.063	98
5	0.642	0.266	0.092	99
3	0.696	0.216	0.089	103
4	0.628	0.172	0.200	93

Table 4.2.1.1 EELS stoichiometry results for the LMA coatings note that the coatings are shown out of order, this is because they are arranged by Ti Cat. %.

Each model was fit to an experimental RDF of at least class 2 to ensure accurate experimental representation. The DFT simulations employed the norm conserving pseudopotentials of Lee *et al* [94] with a cutoff energy of 200eV and just sampled at the gamma point. The simulations were performed using the standard SN steps defined earlier with an iteration time step of 2fs leading to a typical total simulation time of 18-20ps. The lattice parameters were fixed during the simulations to ensure constant density and retain the initial cubic cell which is required for the final refinement step. During final refinement the maximum step size for each atom was reduced to 0.001Å in order to restrict the overall divergence from the DFT structure.

Coating Identifier Code	Energy Change per Atom (eV)
0	-1.53
1	-1.39
2	-1.32
5	-1.48
3	-1.36
4	-1.32

Table 4.2.1.2. Energy change per Atom between initial and final states of a structure having undergone a DFT MD simulation in CASTEP. Note that the change is negative which indicates a reduction in total energy.

As would be expected, the RMC process used to create the initial models resulted in structures in high energy states. This is implied by the significant reduction in the total energy of the structure as a result of the DFT simulation as shown in table 4.2.1.2. This indicates that the DFT MD is breaking up high energy configurations in favour of low energy ones which is the desired effect. The absolute values for the energy are not important as they depend on the number of atoms, the exchange functional and the

program used to perform the calculations. The important aspect of the total and cohesive energies is effect the initial parameters have i.e. the stoichiometry and RDF.

Although these simulations were performed using sub-optimal parameters it is still useful to calculate their cohesive energies. This is done using the method described in section 2.5 after the final DFT structure has under gone GeomOpt in order to ensure minimal energy. As expected the cohesive energy varies dependent on the initial modelling parameters as seen in figure 4.2.1.1. Although the correlation with the RDF statistics is poor, they do share some key characteristics. Both the cohesive energy and the 1st peak width drop sharply at 20% doping, then increase rapidly at 25% doping; furthermore at they both drop at 30% doping and rise slowly up to 55%. As stated the correlation is poor, however the presence of these key features suggests a relationship between the two quantities which has been obscured. The likely cause of this poor correlation is the sub-optimal parameters used in the simulation, specifically the cut off energy. It is likely that using optimised parameters would resolve this issue and show a clear trend.




Figure 4.2.1.1 Calculated cohesive energy of the LMA coatings against the full width at half maximum of the 1st peak in their respective RDFs. The values for the Y-axes are arbitrary and unrelated and as such have been removed.

From an energetic point of view the DFT process has performed well, and led to a more energetically favourable and likely state. Understandably the reduction in energy of the models has led to a significant change in state of the structure and its RDF. As seen in figure 4.2.1.2 the RDF of the model, which initially matched the experimental RDF has seen considerable drift. The DFT process has not led to sub-Ångstrom structure and as such the RDF below 1Å has remained flat, confirming the original hypothesis that the sub-Ångstrom peaks are due to noise or poor fitting. The DFT process has also overestimated the abundance of the nearest neighbour distance responsible for the 2Å peak. Furthermore the RDF has become bimodal around 3.5Å due to over estimation of a nearest neighbour distance at 3Å. Using the partial RDFs such as figure 4.2.1.3 as a guide these peaks can be identified. As suspected the 2Å peak in the RDF is due to Ti-O and Ta-O distance and the 3.5Å peak is a superposition of Ti...Ti, Ta...Ta, Ti...Ta although the former is weak up to 55% doping.

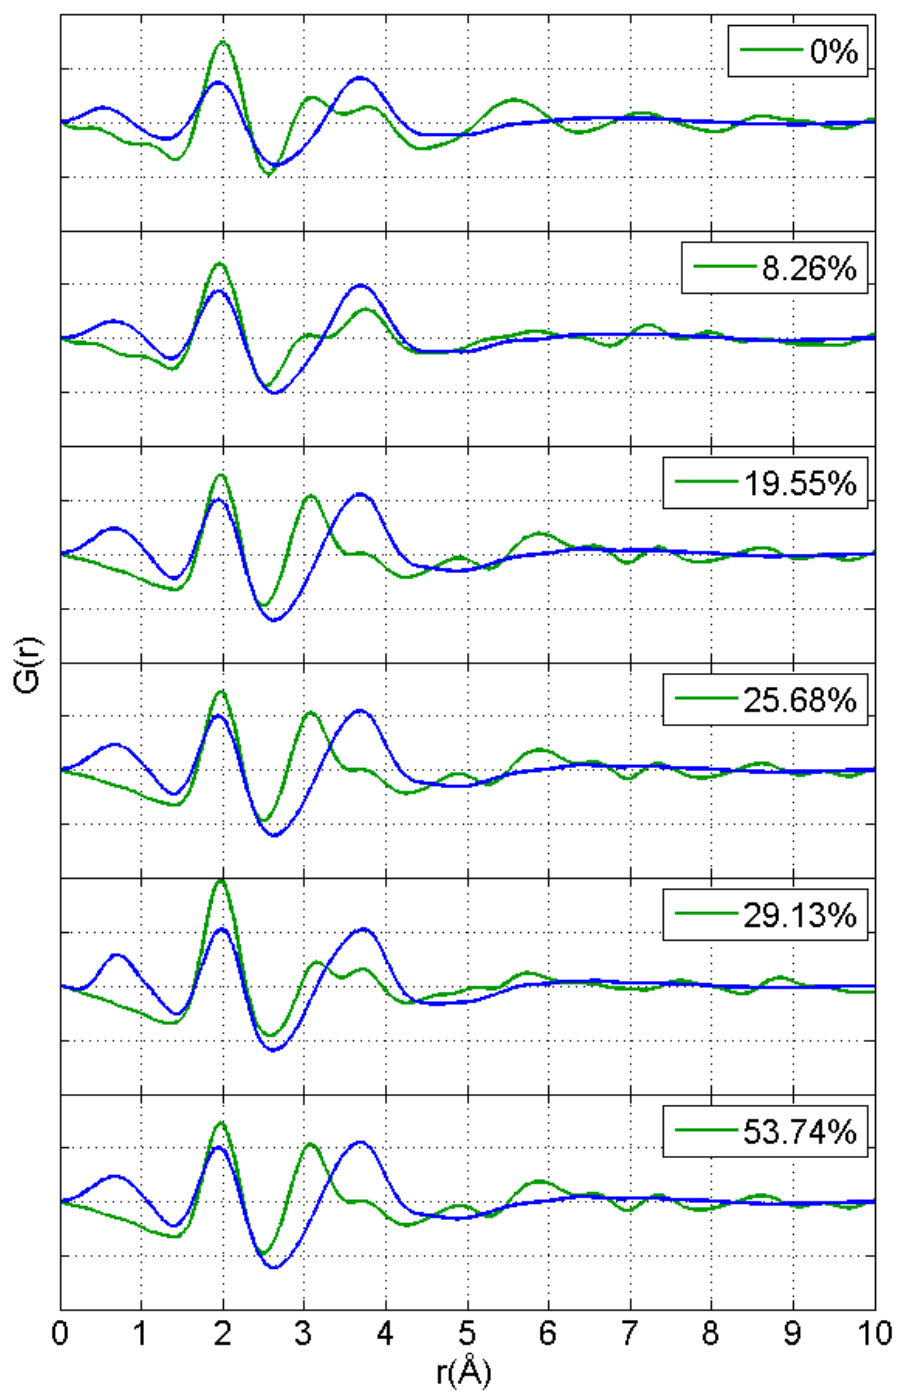


Figure 4.2.1.2 Experimental (blue) and theoretical (green) RDFs for the LMA coatings, The theoretical RDFs are calculated for the structure upon completion of the DFT stage.

Through extraction of the partial RDFs from the DFT models the evolution of each nearest neighbour distance can be determined. Throughout the doping series the O...O distance

remains a relatively weak peak occurring at $\sim 2.8\text{\AA}$. This distance does not represent direct O...O distances, instead it would be expected that they would occur in a O-Ti-O or O-Ta-O configuration. Direct O...O nearest neighbour distances or bonding would be expected around 1.1\AA as previously stated, this is not seen in the DFT models. This peak is not clearly resolved in RDFs calculated from diffraction patterns recorded under standard conditions, but only when HDR acquisition is used, as described in Section 3.3.1.




Figure 4.2.1.3 Partial RDFs for all 6 LMA coatings between 1 and 4\AA . All 6 plots share the same Y and X axis scaling but the values are arbitrary and thus have been removed.

As predicted the Ti-O and Ta-O nearest neighbour distances are centred on 2\AA due to their similar covalent radii. The Ti-O nearest neighbour distance seems to consistently increase in abundance with doping. This is confirmed by quantifying the Ta-O to Ti-O as seen in figure 4.2.1.4 where the ratio follows a linear trend with Ti Cat %. This confirms the expectation that the Ti Cat. % is increased the peaks become dominated by Ti distances over Ta. Furthermore Ti-O distances appear to be longer than Ta-O as seen in figure 4.2.1.5. Combining the results shown in figure 4.2.1.4 and 4.2.1.5 the increase in the 1st peak position seen experimentally in figure 4.1.3a can be explained. As the doping

level is increased the longer Ti-O nearest neighbours become increasingly dominant, shifting the overall peak to longer distances.




Figure 4.2.1.4. Ratio of abundances of Ta-O to Ti-O nearest neighbours against Ti Cat % as measured as the heights of the peaks in the figure 4.2.1.3.




Figure 4.2.1.5. Nearest neighbour distances of the Ta-O and Ti-O nearest neighbours against Ti Cat. %

The metal-metal distances occur approximately where expected from previous modelling efforts [105] i.e. $\sim 3\text{-}3.7\text{\AA}$. Ti...Ti distances seem to be rare up to $\sim 55\%$ doping suggesting Ti...Ta and Ti...O are energetically preferable. These distances become prominent at $\sim 55\%$ doping, likely due to a lack of Ta in the sample which would also explain significant drop in abundance of the Ta...Ta and Ta...O distances. The bimodal nature of the RDF seen in figure 4.2.1.2 is indeed an overestimation of a 3\AA distance however the results are somewhat unexpected. Where it was speculated that some unique nearest neighbour would occur at this distance, it is in fact populated by all the metal-metal distances.

As stated, during the DFT process the RDF becomes bimodal at $\sim 3.5\text{\AA}$. These two nearest neighbour distances occur at 3 and 3.8\AA , which alludes to a possible cause. The sum of the covalent radii of Ta...Ta, Ti...Ta and Ti...Ti is approximately 3\AA for each case suggesting a close approach nearest neighbour distance is the cause of the 3\AA peak. The

question of whether the metal atoms are forming bonds at this distance would require further work not covered in this thesis. The longer 3.8Å distance however is what would be expected of a Metal-Oxygen-Metal distance. This could explain why the 2nd peak is bimodal, as it is split into direct and indirect Metal-Metal distances.

4.2.2 The role of Planar Fragments in Titania doped Tantalum

In crystalline materials, there exists a configuration of atoms which can be tessellated to create the overall structure; this is called a unit cell. For polycrystalline materials a similar configuration exists however, simple tessellation is no longer sufficient to create the overall structure as cells can be orientated differently. As a unit cell exists for both crystalline and polycrystalline structures it stands to reason that some analogous structural motif should exist for amorphous materials. Whether said motif is capable of describing the overall structure of an amorphous material through tessellation and re-orientation (almost certainly) whilst allowing for variation in the length scale is unclear.

In previous studies of undoped Tantalum, a repeating atomic configuration was found to occur in semi-empirically driven models [105]. This motif took the form of a 4 atom planar fragment consisting of 2 Tantalum and 2 oxygen atoms (Ta_2O_2) in a diamond shape configuration (figure 4.2.2.1). These planar fragments were found to form during the DFT stage of modelling suggesting they represent a low energy configuration. They are also found in Titania doped Tantalum models along with a few variations. A undoped Titanium (Ti_2O_2) fragment is found as well as a 'Hybrid' fragment containing both Titanium and Tantalum (TaTiO_2) shown in figure 4.2.2.1. This indicates that in the doped structures the Titanium atoms are capable of sitting at previously Tantalum dominated position. This is possibly due to the covalent radii of Titanium being similar to that of Tantalum. Thus allowing it to replace a Tantalum atom without significantly changing the dimensions of the structure.




Figure 4.2.2.1. 3 variants of the 4 atom planar fragments found in Tantalum and Titania doped Tantalum models after DFT stage.

It would be expected that as the dopant level is increased then the number of Ti doped fragments would similarly increase. Through analysis of the partial RDFs it has already been determined that the abundance of Ti...Ti distances is low up to 55% doping this would indicate a lack of pure Titanium planar fragments. There is also a general decrease in mechanical loss up to 55% doping, which suggests a link between planar fragment abundance and mechanical loss.

An assumption can be made that all the nearest neighbours shown in the PRDFs are due to these planar fragments. This is a fair assumption, if it is accepted that not all the fragments will be fully formed, some may have been in a transitional state when the DFT simulation finished. The abundances of Ta...Ta, Ti...Ta and Ti...Ti distances would then represent the abundances of each of their respective planar fragments. Since there is little or no Ti...Ti distances up to 55% these fragments can be ignored up to that dopant level. It is therefore most likely that the decrease in room temperature mechanical loss is controlled by the ratio of Ta...Ta to Ti...Ta fragments.

To determine the abundances of the planar fragments, a series of Gaussians were fit to the 3\AA peak in the PRDFs. As seen earlier these peaks are bimodal and as such can

be represented by the sum of 2 Gaussians. Since the relative abundances of a nearest neighbour distances are determined from the peak height, then the total abundance is just the sum of the heights of the 2 Gaussians. An additional parameter was introduced to account for the non-zero starting point of the Gaussians i.e. $G = g_1 + g_2 + C$, where G is the total Gaussian fit, g_1 and g_2 are the individual Gaussians and $\frac{1}{2} C$ is adjustment of the intercept for each Gaussian. The parameter $\frac{1}{2} C$ was kept constant for each PRDF but allowed to vary between different PRDFs as the zero point is not constant between PRDFs.




Figure 4.2.2.2. Ratio of Ta-Ta to Ti-Ta distances against Ti Cat. % calculated from Gaussian parameters. Also shown is the room temperature mechanical loss on an appropriate scale.

In figure 4.2.2.2 the results of this analysis can be seen alongside the room temperature mechanical loss. The results presented are that of the ratio of the

abundance of Ta...Ta to Ti...Ta distances. As can be seen there is a strong correlation between the ratio and the loss up to 55% doping. As stated earlier, there are few Ti...Ti fragments up to 55%, which means that the properties of the structure are determined by the Ta...Ta and Ti...Ta distances. At 55% doping, there is a significant abundance of Ti...Ti fragments and these fragments are not taken into account in this analysis. Up to 55% doping, the loss and TaTa:TiTa ratio have a 97% linearly correlation implying that the loss is indeed controlled by this ratio. This suggests that Ti...Ta fragments are less dissipative than Ta...Ta fragments.

In order to determine why a structure containing more Ti...Ta fragments is less dissipative, the planar fragments were modelling to calculate their zero Kelvin vibrational modes. These calculations determined a series of frequencies for vibrational modes in single and double fragment systems (figure 4.2.2.3) the 3 lowest frequencies for each system are shown in tables 4.2.2.1 and 4.2.2.2. From these vibrational results, in general it can be seen that the Ti...Ta system tends to have lower vibrational frequencies for its low order modes. The Ti...Ti systems seem to have the highest vibrational frequencies, with the Ta...Ta systems somewhere in between.

Species	1 st mode (cm ⁻¹)	2 nd mode (cm ⁻¹)	3 rd mode (cm ⁻¹)
TaTa	104	138	156
TiTa	68	204	264
TiTi	184	271	325

Table 4.2.2.1. Vibrational frequencies for the first 3 zero Kelvin modes for a system containing a single planar fragment.

Species	1 st mode (cm ⁻¹)	2 nd mode (cm ⁻¹)	3 rd mode (cm ⁻¹)
TaTa	32	34	69
TiTa	16	36	49
TiTi	64	73	175

Table 4.2.2.1. Vibrational frequencies for the first 3 zero Kelvin modes for a system containing two planar fragments.

The vibrational frequencies can elude to (in relative terms) the stiffness of the structure. A flexible structure would typically require less energy to excite than a stiffer one and this would be reflected by a lower vibrational frequency. This suggests that the Ti-Ta systems are more flexible than the Ta-Ta and Ti-Ti systems and therefore a relative decrease in TaTa/TiTa ratio would increase the overall flexibility of the structure. The picture then emerges that with increasing Ti doping, a greater percentage of planar fragments become (or form) doped leading to an increased flexibility in the overall structure. From experimental results (figure 4.1.4) it is known that the structure more closely approaches homogeneity with higher doping. This suggests that the increased flexibility allows the structure to more closely approach homogeneity.

This now, allows for speculation on the origins of room temperature mechanical loss in Tantalum and to why Titania doping decreases the loss. As the structure becomes more flexible it is likely, that it would more easily be able to accommodate defects into the structure. This would lead to a decrease in the amount of discontinuity within the structure moving it toward homogeneity, thus reducing energy dissipation and hence mechanical loss. In order to experimentally test this, Tantalum coatings would need to be doped with an element which significantly increases or decreases the flexibility of the planar fragment (assuming it can form a new hybrid variant). Indirect evidence which supports structural flexibility as a loss mechanism has been seen in the studies of amorphous Si and SiO₂. The theory would suggest that room temperature mechanical loss of Si would be greater than SiO₂ as Si is further from homogeneity (thus more stiff) than SiO₂. This has previously been investigated as part of general coatings research in the LSC and IGR and indeed, Si does show greater loss than SiO₂ at room temperature [131, 132, 135].

This is by no means a proof of concept, as it is not certain that flexibility is the mechanism of loss for all structures, or indeed just for Tantalum. However it does lead to road map for future experimental and modelling work. Firstly a new dopant must be determined which further increases flexibility either as a co-dopant to, or replacement of Ti. This new dopant must then be modelled using the modelling process key to this thesis

in order to determine whether it can still form planar fragments with Ta. If this is the case, then new coatings can be created using these new dopants to experimentally determine the mechanical loss.

4.2.3 Analysis of the final models

As explained in the previous sections the DFT stage of modelling is not the final stage. Upon completion of the DFT simulated anneal, a supercell of 3*3*3 DFT cells is created (with randomised orientations). This supercell is then refined, using RMC, back into alignment with the experimental RDF, from which it drifted during the DFT stage. The purpose of this stage of modelling is to ensure that the model matches the experimental data taken from the TEM. The increased size of the cell also significantly increases the statistics of the model i.e. significantly more atoms over which to average, as well as including longer range order. Furthermore the larger box size allows use and analysis of longer distances within the structure limited by the necessarily smaller DFT models. To clarify, one would not expect distances within a model of greater than half the simulation size to be accurately represented as the size of the simulation then becomes relevant. A typical Tantalum DFT model is an $\sim 11\text{\AA}$ cube and thus will accurately represent distance below 5.5\AA which is perfect for Tantalum as the longest distance of particular interest is no more than $\sim 4.2\text{\AA}$. To reduce the possibility of shifting the structure far from the energetically relaxed state it will have achieved in DFT, the step size of the RMC is reduced considerably, a step size of 0.001\AA is a typical. In essence this stage of modelling is intended to “tweak” the model back into experimental alignment as opposed to the “brute force” approach used during the initial RMC stage of modelling.

The RDFs of these final models have not been presented as they match the experimental RDFs. However, the partial RDFs are shown in figure 4.2.3.1, the first thing to notice is the refined nature of the PRDFs. Each function is smoother than the PRDFs from the DFT stage shown previously, due to the increased number of atoms. This is especially noticeable at longer distances, however, typically for Tantalum only distances up to $\sim 4\text{-}5\text{\AA}$ are of any importance since after this point the RDFs and PRDFs become quite

complicated. Again, as with the experimental RDFs distances below 1Å are non-physical and typically the RDFs are flat within these regions confirming said. This creates a region of interest in the RDFs of 4Å between 1 and 5Å which typically contains the shortest distance for each atomic pairing.




Figure 4.2.3.1 Partial RDFs for the LMA series of coatings after the final refinement stage of model, each model here had ~ 2700 atoms. All 6 plots share the same Y axis scale however the absolute values are arbitrary and so have been omitted. The X-axes have the limits 0 – 5 Å and have a spacing of 1Å.

The PRDFs of these final models seem to be able to explain the behaviour of the experimental RDFs seen in section 4.1. As seen in figure 4.1.2 and 4.1.3 the first peak of the RDF, which is due to metal-Oxygen distances, shifts to longer distances with doping. It was hypothesised that this was due to an increasing abundance of a longer distance nearest neighbours pulling the peak to a longer distance. However according to these PRDFs the shift in the RDF peak is due to a general increase in the nearest neighbour distances of both the Ti-O and Ta-O distances. This is not likely to be what is actually happening in the structure, it is still more likely that the Ti-O and Ta-O nearest neighbours have slightly different distances and the shift is due to an increasing relative abundance of the Ti-O over the Ta-O as was seen in section 4.2.1. This is based on previous research in the literature which stated that the Ti-O have slightly longer nearest neighbour distances than Ta-O in the crystalline form of Titania doped Tantalum [134].

The PRDFs do explain, in an expectable way why a “shoulder” develops on the 2nd peak in the RDFs seen in figure 4.1.3. It was believed that this was due to an increasing abundance of the 3Å peak causing a slight separation and thus creating a bimodal distribution centred on 3.5Å. The peak in question is not the Ti···Ti peak which it was expected to be, it is in fact the Ti···Ta peak. It was expected that the Titanium would form slightly shorter metal-metal nearest neighbour distances than Tantalum and as such the 3 metal-metal peaks would occur in ascending order of length Ti···Ti, Ti···Ta and Ta···Ta. However this is not the case, in most of the PRDFs extracted to date, the Ti···Ta distance have appeared at a shorter distance to that of the Ti···Ti and Ta···Ta. This could suggest that Ti atoms form tighter bonds to Ta than they do to Ti if the peak is due to direct metal-metal bonding. Alternatively, it could also suggest that Ti-O-Ta bonds form with more acute angles than Ti-O-Ti or Ta-O-Ta thus leaving the metal ions closer together. A third possibility is that the TiTaO₂ fragments form smaller, i.e. shorter bond distances than their pure element counterparts, this would also bring the Ti···Ta atoms closer together.

It is unlikely that direct metal-metal bonding will occur in this material like this. This invalidates the first of the three possibilities; the latter most explanation is difficult to test as it would involve identification and extraction of all planar fragments within the

structure. The remaining explanation could be tested by examining the structure directly and extracting the bonding statistics i.e. bond angles, dihedral angles etc. While this type of analysis is possible, and has been performed in the past [129] it will not be covered in this thesis.

The final piece of the analysis to consider is the peak height ratios of these PRDFs as was presented in section 4.2.1 for the DFT models. This analysis is performed for a several reasons, firstly to confirm whether or not the trend seen in the PRDF analysis for the DFT models was still present after final refinement. This will indirectly confirm whether the final refinement stage essentially undoes the energetic relaxation achieved by the DFT stage. Furthermore with a greater number of atoms (greater statistics), and access to longer range structure the results of such analysis are less open to scrutiny as they are averaged over a much greater number of occurrences.

As was done previously, each PRDF peak for the Ta-Ta, Ti-Ta and Ti-Ti was fit with a Gaussian distribution and the heights of the peaks was read from the Gaussian parameters. As would be expected from figure 4.2.3.1, for the most part single Gaussians were sufficient to accurately model the PRDF peaks. Unlike in section 4.2.2 where only the contributions from Ta-Ta and Ti-Ta nearest neighbours were taken into account, here the contributions from Ti-Ti neighbours were also included. This introduced a complication, which was whether the Ti-Ti and Ta-Ta contributions should be weighted equally in the ratio. Where previously the ratio presented was simply that of the Ta-Ta and Ti-Ta contributions, here the ratio becomes $R = (TaTa + (A * TiTi)) / TiTa$. Where $TaTa$, $TiTi$ and $TiTa$ are their respective contributions and A is a scaling factor for the $TiTi$ contributions.

The question becomes, what is value of A ? With no other evidence, it would be easiest to assume $A = 1$ i.e. that the $TiTi$ and $TaTa$ contributions be weighted equally. However as seen previously, the vibrational frequencies of $TiTi$ and $TaTa$ fragments are not equal, the $TiTi$ fragments are stiffer. This suggests that they would have a greater detrimental effect on the overall flexibility of the structure than $TaTa$ fragments. If it were

assumed that the flexibility of the structure essentially scales linearly with the vibrational frequencies then $A \approx 1.9$. However, it is also possible that the flexibility contributions scale as the square or cube of the vibrational frequencies leading to values of A of 3.61 and 6.58 respectively. There are no proposed physical bases for various contribution scaling's, they are merely to determine a best fit, from which a physical basis could be hypothesised.

Since it is difficult to determine without further study which of these scaling is correct all were used separately. In figure 4.2.3.2 the results of these different scaling values on the fragment ratio is shown. For the unweighted contribution ratio, the final point at 53.76% Ti Cat. % is calculated to be lowest of all the doping levels. This would suggest that a coating doped to 53% would have lower loss than at 30% doping, has already been proven untrue. With linearly weighted TiTi contributions, the ratio is found to be lowest at 30% which coincides with the mechanical loss measurements. This is also true for both the subsequent scaling, so it is likely that equally weighting the TiTi and TaTa contributions is invalid.

For each scaling value the Pearson linear correlation coefficient [136] was determined, which is a measure of the likelihood of functions being linearly related. The Pearson correlation results are shown in table 4.2.3.1 and show that the squared relationship achieves the best fit. This is of course, by no means proof that mechanical loss scales as the square of the vibrational frequencies of the crystalline fragments which make up the structure. However it is further evidence that modelling and experimental processes are moving towards greater understanding of the structural cause of mechanical loss.




Figure 4.2.3.2. The ratio of Ta...Ta distances summed with Ti...Ti distance against Ti...Ta Distances. This is an indication of the ratio between the pure and doped planar fragment variants. Shown are four functions, each has a different scaling of the TiTi contributions, using $A = 1, 1.9, 1.9^2$ and 1.9^3 . Shown alongside for comparison is the room temperature mechanical loss of Titania Doped Tantalum.

A	Correlation
1.00	0.9423
1.90	0.9556
3.61	0.9721
6.86	0.9588

Table 4.2.3.1 Calculated Pearson linear correlation coefficient for each TiTi contribution scaling

4.2.4 Analysis of Density Functional Theory Models Performed Using Optimised Parameters

As explained in the previous chapter, the choice of parameters in a DFT simulation determines whether or not the model can be considered “optimised”. Optimisation in this case is dependent primarily on kinetic energy cut off and k-point grid variables. For the majority of models performed for this thesis these values were set to 200eV and [1 1 1] respectively, values below optimisation. Typical “optimised” values for these parameters using CASTEP and the norm conserving pseudopotentials by Lee *et al* [94] are 600eV and [2 2 2] respectively. The suboptimal parameters as used simply because the computational expense of using optimised parameters is unfeasible with the computational resources available.

However, through the use of the HECToR supercomputer with the aid of funding from the EPSRC, optimised models have been performed for the LMA coatings. These simulations used the optimised parameters stated above however, despite the significantly increased computational resources of HECToR, still took over 3 months of compute time to complete. As would be expected the initial and final energies of these models were far lower than that of the unoptimised models. This is due to more accurate representation of the wavefunctions due to the inclusion of higher energy plane waves.

The result of these optimised models has been, so far, encouraging. As explained and presented previously, after the DFT stage of modelling the structures are typically in a state in which they no longer match the experimental RDFs as they did following the RMC stage. This has been attributed to the fact that during the RMC stage the energy of the structure is not taken into account and as such it is likely to form energetically unfavourable configurations. When using optimised parameters however, although the structures still drift from experimental RDF alignment, they appear to drift less. This can be seen in figure 4.2.4.1 where the RDFs from the unoptimised models seem to have universally drifted further from experimental than the optimised models. This means that there likely exist, energetically relaxed states which can also accurately reproduce the

experiment RDFs on which they are based. What is likely holding the simulations back from achieving such states is the choice of functional, pseudopotentials and the various approximations on which such calculations exist. Though the previous statement is not unexpected, it does however, warrant further study aimed to determine the optimal pseudopotentials, functionals for achieving experimental RDF representivity.

Examining the RDFs of the optimised models against the unoptimised models reveals various inaccuracies brought about from sub-optimal parameters. Firstly it appears that the overestimation of the metal-oxygen distances (first RDF peak) is due, in part to sub-optimal parameters. For the optimised models, the abundance of metal-oxygen distances appears to be closer to that seen experimentally. It must be noted that shorter distance peaks are more evidently affected by a limited q range than longer distance peaks, and they are also more greatly affected by noise than longer distance peaks. This means that the true experimental peak height of the shorter, metal-oxygen distance may not be what appears here, it may well be greater and thus closer to the estimations of the DFT simulations.

It also seems that the optimised models favour a singular metal-metal distance at $\sim 3.5\text{\AA}$ which was not seen in the unoptimised models. As seen in figure 4.2.4.1 the sub-optimal parameters lead to a bimodal distribution at $\sim 3-3.5\text{\AA}$ possibly due to a combination of direct metal-metal bonding at 3\AA and metal-oxygen-metal bonding at 3.5\AA . It also seems that the optimised models prefer shorter distances than both the unoptimised models and the experimental RDFs. This is possibly due to stronger bonding in the models than would be expected in nature. However, it is also likely to be a miscalibration of the lengthscale of the CBED images. None of the models, optimised or otherwise seems to show a peak at $\sim 1.1\text{\AA}$ which has been seen experimental in HDR RDFs. This could be an indication that this peak is a result of noise, or a contaminant in the sample, trapped water for example. It could also indicate a weakness in the oxygen pseudopotential as it a 1.1\AA distance is about right for direct Oxygen...Oxygen bonding.

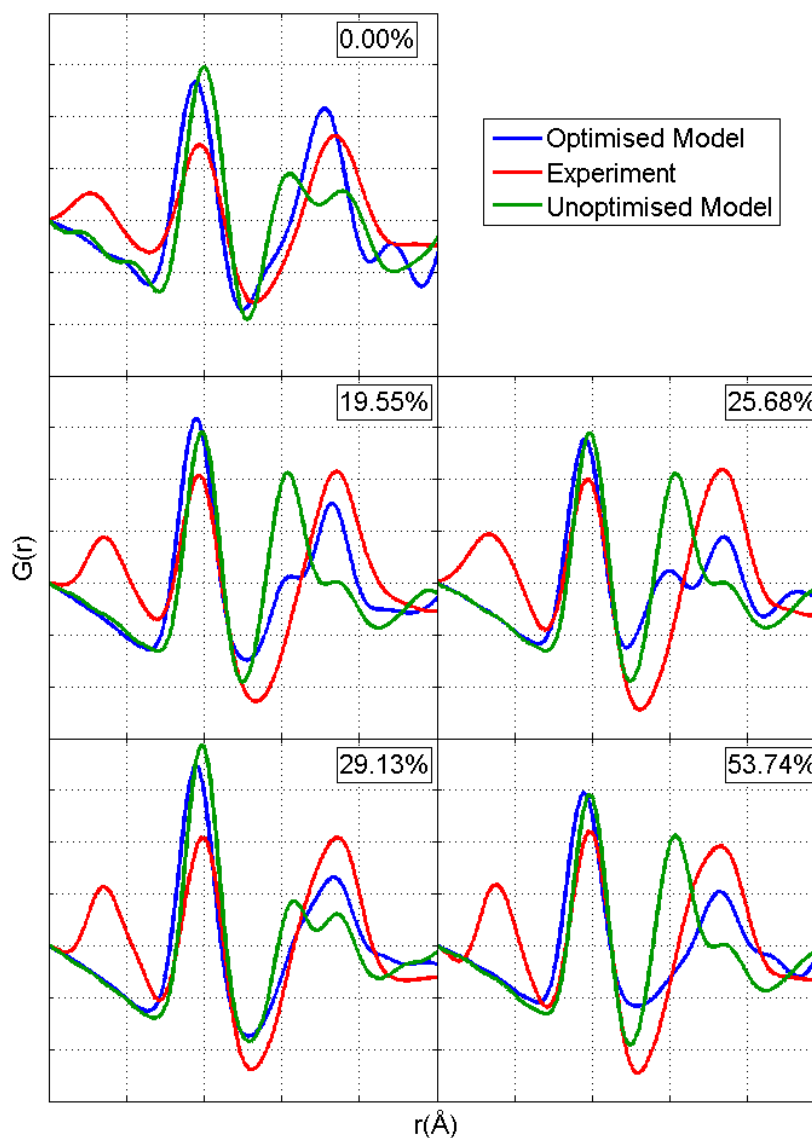


Figure 4.2.4.1 RDFs from generated from the experimental data (red), the models which used Unoptimised parameters (green) and optimised parameters (blue). Notice that for the most part, the RDFs from optimised models more closely match the experimental than the unoptimised models. All 5 plots share the same Y axis scale however the absolute values are arbitrary and so have been omitted. The X-axes have the limits 0 – 5 Å and have a spacing of 1Å.

After extraction of the Partial RDFs from the optimised models some interesting results become clear. It has been seen previously in the experimental analysis of the RDFs

(section 4.1) that both the first and second peaks in RDF move to longer distances with doping (figure 4.1.3). The PRDFs seem to explain this behaviour well for the first peak, considering the behaviour of the Ti-O and Ta-O distances. As seen in figure 4.2.4.2 the Ta-O peak remains essentially constant in position with doping remaining just slightly less than 2Å. The Ti-O peak also remains essentially constant in position at just slightly more than 2Å. What changes is the ratio between both peaks, with increasing doping the Ti-O peak becomes taller. This will have the effect of shifting the first peak of the RDF to longer distances, and this is postulated at the time and seems to be have been confirmed.

For the second peak, the relationship is more uncertain as would be expected due to the increased number of constituent peaks. However, by concentrating on the two strongest constituent peaks, that of the Ta···Ti and Ta···Ta distances the behaviour of the RDF peak can be implied. As the structure doped with a greater amount of Titanium the Ta···Ta peak seems to remain in a constant position and simply changes height indicating the expected change in abundance of Ta···Ta distances. The Ta···Ti distances increase in abundance, as would be expected but at the same time shift from mostly 3Å distances to a near equal distribution of 3 and 3.7Å distances. This again, like with the first peak would have the effect of shifting the RDF peak to longer distances with doping.

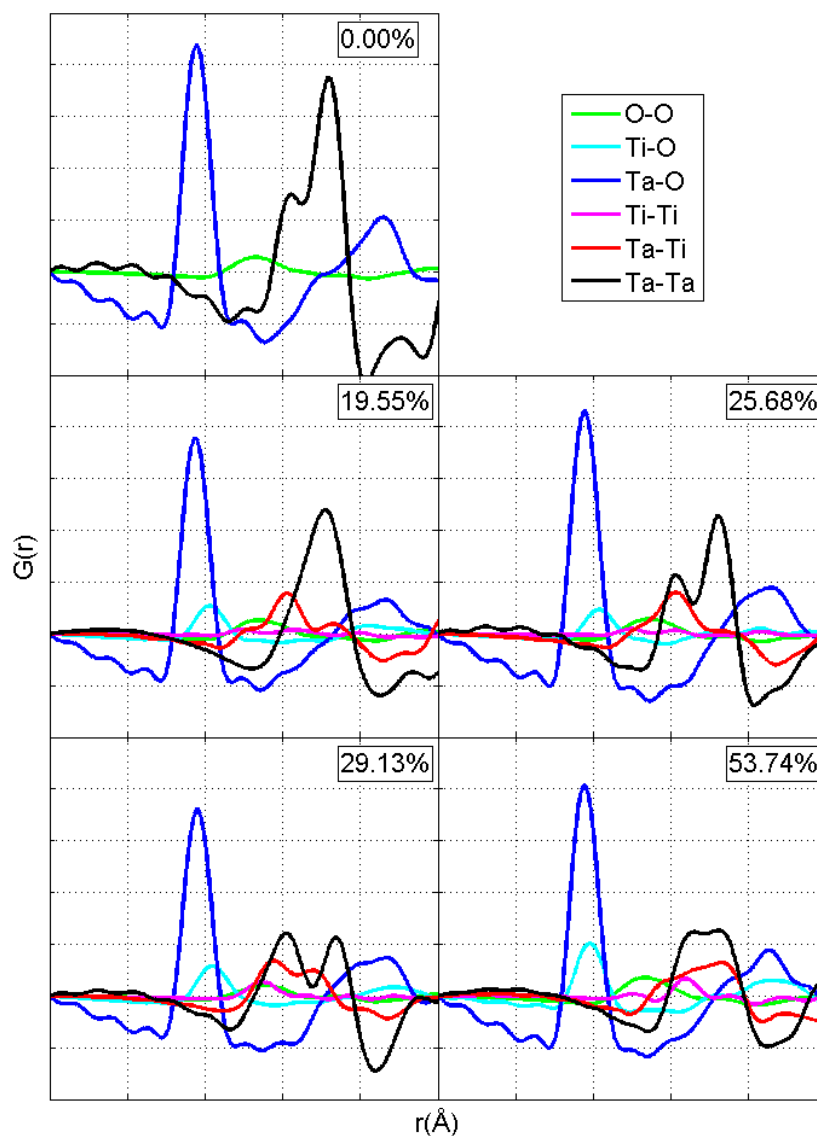


Figure 4.2.4.2 Partial RDFs for the completed “optimised” models for the LMA coating series. The model for the 8.26% model is missing as it was corrupt and unusable. All 5 plots share the same Y axis scale however the absolute values are arbitrary and so have been omitted. The X-axes have the limits 0 – 5 \AA and have a spacing of 1 \AA .

These are of course only qualitative analyses however, quantitatively the analysis performed on previous modelling efforts would be quite difficult. For previous results, single or double Gaussians were used to model the peaks in these PRDFs this is unfortunately not sufficient in this case. In order to correctly represent the peaks in these

PRDFs 3 and 4 Gaussians are required, this is perhaps a true representation of the structure. However, it is also likely that the peaks are bi modal (and thus representable by a double Gaussian) and are being skewed and obscured by approximations. This makes quantitative analysis unreliable in this case until such time as a more robust method of peak modelling can be explored.

4.3 Chapter Summary

Titania doped Tantalum is an important material for future gravitational wave detectors for use as a high index mirror coating [30]. It is therefore necessary to understand its mechanical, optical and structural properties and the relationship each have with each other. Currently work is underway to determine the mechanical and optical properties of this material through experimental and analytical means through the LSC [31, 40, 43, 49, 52, 63, 133]. Studies are also in progress to determine the structure of amorphous Titania doped Tantalum and its properties [105, 118, 129]. Here a series of 6 Titania doped Tantalum coatings each with a different level of doping have been experimentally processed and theoretically modelled.

From an experimental perspective, Titania doping of Tantalum does not cause phase separation within the structure, instead Titania seems to blend seamlessly into the Tantalum matrix. Significant changes are seen quantitatively in the diffraction patterns and the RDFs. These changes suggest that as Tantalum is doped more heavily with Titania, the nearest neighbour distances lengthen whilst overall the structure moves towards homogeneity. Not only are these changes detectable through RDF analysis but also that the progress towards homogeneity correlates strongly with the mechanical loss.

Through RMC and DFT modelling the structures of these coatings have been determined, as has the behaviour of the nearest neighbours with doping. It was found that the crystalline building block of Tantalum, a 4 atom diamond shaped planar fragment is also found in Titania doped Tantalum. Furthermore it was found that these planar fragments can be doped directly without destroying them, possibly due to Titanium's

similar covalent radius to Tantalum. Moreover, it appears that the ratio of these different types of planar fragment seem to have a significant effect on the mechanical response of the coating. It is indeed possible that inherent flexibility of each planar fragment variant is directly linked the room temperature mechanical loss. A very high correlation has been found between the ratio of the stiff pure element fragments over hybrid fragments and the mechanical loss. Though it warrants further study and experimentation to confirm, it is possible that by simply increasing the flexibility of the Tantalum planar fragment or by increases the abundance of more flexible variants the mechanical loss of the structure can be lowered.

Finally, through use of HECToR and with funding from the EPSRC, new realms of modelling have been opened up with regards to the use of optimal modelling parameters. The use of these parameters has been shown to significantly decrease the drift of the structure from experimental alignment during DFT simulations. This suggests that with further experimentation with different pseudopotentials, exchange functionals and other modelling parameters; it will be possible to create models which are both energetically relaxed, and yet still match the experimental RDFs on which they are based. It is also necessary to improve the data acquisition process and post processing steps to include HDR imaging and High q_{\max} diffraction data to further resolve shorter peaks in the RDFs specifically the 2Å Metal-Oxygen peak and the recently discovered 1.1Å Oxygen...Oxygen Peak.

Chapter 5: Modelling of Annealed Titanium Dioxide doped Tantalum Pentoxide

5.0 Chapter Introduction

The high index material used in the multilayer reflective coatings for aLIGO will be 30% Titania Doped Tantalum [30]. These coatings will be heat treated to $\sim 600^{\circ}\text{C}$ for ~ 24 hours and each layer will be $\frac{1}{4}$ wavelengths thick [102]. The reason the coatings are heat treated is because heat treatment has been shown to improve the optical properties [62]. Heat treatment has also been shown to affect the mechanical loss in pure Tantalum [43, 58]. It is therefore important to determine the effect of heat treatment on Titania doped Tantalum both in terms of mechanical loss and the micro-structural change. For this reason a series of 8 coatings have been investigated experimentally and through modelling. These 8 coatings are divided into two sample sets, each doped to a different level 25% and 55%. The 4 coatings within each sample set were annealed at different temperatures, 100°C , 300°C , 400°C and 600°C . This coating set should allow for the characterisation of structural change due annealing and determine how that change is affected by annealing temperature.

5.1 Experimental Analysis

5.1.1 Image Analysis

These 8 coatings, here-after are referred to as the CSIRO coatings, were investigated experimentally using a FEI Tecnai F20 TEM. The main differences between this TEM and the T20 used for previous work are the electron source, and the camera and its position in the microscope. The F20 uses a thermally assisted (Schottky) FEG to produce a more coherent, more intense beam of electrons than the thermal source on the T20. The camera, a Gatan MSC 794 model, CCD Camera is mounted beneath the viewing port on the microscope and in line with the beam. The main advantage of using the F20 over the T20, is the camera and specifically the interconnect between the scintillator and the CCD. For the SIS that interconnect is in the form of the prism

mentioned earlier in chapter 3, this causes the defects also shown in chapter 3 for the reasons described there. The 794 does not suffer from this issue, as each pixel on the CCD is connected to the scintillator directly via a fibre optic cable. This therefore does not allow for any distortion due to reflections as is the case with the SIS camera. This allows for the use of the entire CCD instead of the need to discard large quantities of data. Furthermore the 794 has also been seen in experience to be more resistant to blooming compared to the SIS allowing it to represent more truly, the intensity due to the diffraction. A slight disadvantage of the F20 is that the 794 camera has only 1024×1024 pixels compared to the SIS which has 1376×1032 . This gives the F20 a slightly decreased q range at the 190mm camera length typically used for this work, compared to that used with the SIS camera. The 190mm camera length was chosen as it most closely matches the q resolution possible on the SIS camera on the T20. A greater q range could of course be achieved by lowering the camera length; however doing this would also decrease the Q resolution which is unwanted.

As with the LMA coatings, a large number of diffraction patterns were recorded from each of the CSIRO coatings, in order to create an average diffraction pattern. When compared qualitatively, the diffraction patterns from each of the 8 coatings appear similar as seen in figures 5.1.1 and 5.1.2. This again indicates that any changes in the structure are subtle. The Diffraction patterns also confirm that crystallisation has not occurred in the structure as they remain amorphous. As can be seen in the top right of each diffraction image, the beam blank appears quite large with respect to the camera size. This is due the CCD being mounted lower in the F20 compared to the SIS in the T20, this causes the beam blank to cover more of the diffraction pattern at the lower camera lengths used in F20. The beam blank results in the loss of approximately $\frac{1}{4}$ of the data, which is unfortunate but difficult to avoid.

Stated Ti %	O %	Ti %	Ta %	Ti Cat. %
25%	0.653	0.047	0.300	13.03%
55%	0.607	0.209	0.184	52.87%

Table 5.1.1.1 Stoichiometry for 2 coatings in the CSIRO series doped to 25% and 55% and annealed for 24 hours at 400°C.

Further experimental work included EELS in order to constrain the model and RDF stoichiometry. Due to time constraints, only 2 of the 8 coatings had EELS performed on them, these were the 25% and 55% both annealed at 400°C. It has been seen in earlier work [129] that stoichiometry of Tantalum coatings does not change significantly with heat treatment. Thus it was decided to assume that the stoichiometry of the remaining 3 heat treatment temperatures is similar to that of the 400°C coatings. These stoichiometries are shown in table 5.1.1.1.




Figure 5.1.1. Diffraction patterns for CSIRO 25% Titania doped Tantalum in increasing annealing temperature left to right

The RDFs are, as would be expected, very similar to that of those seen in the LMA series of coatings. Each of the RDFs contain two main peaks at ~ 2 and $\sim 3.5\text{\AA}$, most likely corresponding to metal-oxygen and metal-metal distances respectively. The 100°C coating with 25% doping and the 600°C coating with 55% doping both have a series of HDR RDFs and as such show a prominent peak at $\sim 1.1\text{\AA}$. As explained earlier it is believed

that this peak corresponds to direct Oxygen...Oxygen distances since the oxygen covalent radii is $\sim 0.56\text{\AA}$ [137].




Figure 5.1.2 Diffraction patterns for CSIRO 55% Titania doped Tantalum Pentoxide in increasing annealing temperature left to right

In previous studies, analysis of RDFs for coating series with varying annealing temperatures has proven unfruitful [129]. This is most likely because the change in the short range order is subtle for heat treatment. Major changes are seen due to heat treatment in the medium range order as has been recently reported for Silicon [127]. It is possible that with thin samples, high SNR data and immaculate background fitting that these subtle changes can be seen in the short range order as well. It was believed that using the TF20 microscope with the Gatan MSC 794 camera combined with HDR imaging this could be achieved.

The diffraction patterns were obtained and RDFs calculated in the usual fashion. Bright field images were used to record the positions from which each pattern was recorded to allow for area mapping and later repetition of results if necessary. All images were taken using the 190mm camera length which set $q_{\max} \approx 15\text{\AA}^{-1}$. Camera exposures were kept relatively short to avoid over exposure, however acceptable SNR was still achieved (in general). Due to technical concerns regarding the robustness of the camera HDR images were not acquired for all coatings. The inability to use HDR as standard was unfortunate although was in the best interest of the equipment. The loss of

HDR was tempered by the ability of the Gatan MSC to resist blooming one of the primary reasons HDR is essential whilst using the SIS camera. Furthermore the MSC has a much higher saturation limit than the SIS allowing for great SNR even in LDR images.




Figure 5.1.1.3 Experimentally determined RDFs for the four CSIRO coatings doped at 25%, as can be seen, only slight changes can be seen between the different heat treatment temperatures.




Figure 5.1.1.4 Experimentally determined RDFs for the four CSIRO coatings doped at 55%, as can be seen, again, only slight changes can be seen between the different heat treatment temperatures.

The method of analysis of the RDFs from the CSIRO coatings whilst in general, analogous to that used for the LMA coatings did have some key differences. Typically the analysis of the RDFs was done using absolute values read directly from the RDFs. However in the interests of the accuracy and better representation of the true relationships a new method was adopted. This method involved isolating and modelling each peak using Gaussians as was done for the partial RDFs in section 4.2.2. In this case, however, a single Gaussian was determined for each peak, again using a factor C to correct of the non-zero starting point. It was decided to remove this value C from the total calculated height prior

to any averaging. This was justified as it would be expected that multiple RDFs of same coating should have very similar (if not identical) values of C . Hence any major variation in C would be due to noise in the RDF and not change in the structure.

By reducing the fitted Gaussian height by C , the non-zero offset adjustment, the peak height is reduced to the positive peak and not the total height. This is valid for experimental RDFs for the reasons which have already been stated; however it is also unnecessary for the PRDF analysis discussed in section 4.2.2. The CSIRO RDFs and the LMA PDRFs differ in the fact that the former are experimental and the latter are theoretical. Experimental RDFs are affected by noise and as such this must be accounted for. The theoretical PRDFs are not affected by noise as they are calculated directly from the structure and as such the total height of the peaks will be neither under or over estimated by C . Modelling of the coatings was also performed (discussed in section 5.2) and an average RDF was taken to be the basis of those models.

As stated the change due to heat treatment is expected to be subtle in the short range order and that high quality RDFs are expected to be required to detect it. For this reason the RDFs were filtered to remove outliers and anomalous data to hopefully reduce noise and erroneous influences in the final results. This was done by a fairly standard method, which involved excluding data points outside a 2 standard deviation (from the mean) acceptance criteria. This assumes that the data points have a distribution similar to that of a normal distribution. As such ~90% of the data will be within 2 standard deviations of the mean justifying the acceptance criteria. Whether or not RDF statistics do indeed follow a normal distribution is unclear however this is still believed to be a valid assumption.

The analysis of the CSIRO coating RDFs has produced strong trends outside the bounds of error. Although the error is proportionally large, (see figures 5.1.3-5.1.5) is not sufficient to account for the trends. This is mainly due to the filtering explained earlier which was able to significantly reduce the magnitude of the error.




Figure 5.1.3. Heights of the first and second RDF peaks for the 25 (bottom) and 55% (top) doping CSIRO coatings as a function of heat treatment temperature

The heights of the first two peaks show a peak at heat treatment temperature of 300°C. This peak is carried across to both peaks and both doping levels, however for the first peak of the 55% doped coating the peak is within error. This suggests that at 300°C heat treatment the abundance of regular bonding is at its maximum. It has already been seen in undoped Tantalum that further heat treatment past 300°C has a significant effect on the mechanical loss most significantly at cryogenic temperatures [58]. More recently it has a similar trend has been seen in heat treated, Titania doped Tantalum though this is unpublished. This could suggest that the mechanical loss is not related to the abundance of regular bonding in the structure. However at cryogenic temperatures different peaks in mechanical loss appear at heat treatment temperatures. At 300°C there is a single identifiable peak at 40K, however at 600°C heat treatment a 2nd peak is also identifiable at 25K. This suggests that the abundance of the regular bonding is responsible for the peak at 40K in the 300°C heat treated samples but is not the cause of the 25K peak seen in the 600°C sample. It is possible and likely that multiple processes contribute to the mechanical loss and that each will peak at different cryogenic temperatures and be accentuated by different heat treatment temperatures.

The peak positions and peak widths also show similar peaks at 300°C however they are less prominent. If the position and width statistics are to be believed then the total picture would be that at 300°C the structure has the greatest number of nearest neighbours but also that the nearest neighbour distances are at their longest, whilst simultaneously being furthest from homogeneity. As we have already seen with the LMA coatings, mechanical loss seems to decrease as the structure approaches homogeneity. This would explain the peak in the loss at cryogenic temperatures (40K) however it would not explain the peak at 25K seen in the 600°C coating.

There are no clear trends in these results which explain why a peak would appear in coatings which have been heat treated to 600°C. It is possible that these results are still obscured by noise due to low SNR or low q_{max} . However it is more likely as has been seen with Silicon [127] that the peak at 25K is caused by a change in the medium range order which would not be visible in these studies. As previously stated, it would not be valid to

assume that a single mechanism is responsible for all mechanical loss. This is especially true for systems which show prominent features due to doping/annealing at various temperatures or pressures as is the case here. It is more likely that many separate mechanisms exist, and are each responsible for a single loss feature such as the peaks at 25 and 40K.

Changes in the medium range order is suggested as the cause of the peaks seen due to annealing at 600°C+. The reasoning behind this hypothesis is that from previous studies it has been determined that Tantalum crystallises in the region of 800°C in air [63]. Crystallisation changes a structure with strong short range but little long range order to a structure with both strong short and long range order. Up to 600°C no nano-crystals have been found to exist in Tantalum (through bright and dark field studies) [129]. This suggests that the change in the short range order between the amorphous and crystalline states is quite sudden, possibly occurring in as little as a 50°C window. However it is very possible that the medium/long range order begins its change toward crystallinity at a much lower temperature and thus changing the overall mechanical response of the structure. This could be investigated through use of fluctuation electron microscopy (FEM) [138], a method not discussed or used in this thesis. FEM essentially works by mapping the change in the diffraction data of nano volumes across the sample. This can then be used to constrain a large typically classically driven modelling process to create a model of the structure over a large range. This method is therefore sensitive to changes in longer range order which standard RDF measurements are not.




Figure 5.1.4. Positions of the first and second RDF peaks for the 25 and 55% doping CSIRO coatings as a function of heat treatment temperature




Figure 5.1.5. Heights of the first and second RDF peaks for the 25 and 55% doping CSIRO coatings as a function of heat treatment temperature

5.2 Modelling Analysis

Atomic models of these coatings have been generated using in the usual manor, using an average RDF input to remove noise. The stoichiometry for each coating for both the RDF and the model was taken from table 5.1.1.1 assuming that the coatings for each dopant level remained constant across the range of annealing temperatures. In the interests of time and computational expense the models were run using sub-optimal parameters. The cut off energy was set to 200eV and the gamma point was used for the k-point grid.

An 8 stage SN with ~3ps equilibration time per temperature step was used to simulate the annealing process. Finally the internal geometry of each model was optimised in order to ensure that the models sat in a potential well. It was necessary to fix the lattice parameters during both stages of modelling as a cubic lattice is required for the RMC code to work. The initial and final energies of the models indicate a modest change in the energy of the models; however the final energies do not differ significantly. (table 5.2.1). It is possible that the despite the change in the material properties the structural change is still quite subtle. However when the initial energies are shown alongside the energy changes (figure 5.2.1) the true meaning of the results become clear.

Heat Treatment Temperature (°C)	Initial Energy (eV)	Final Energy (eV)
100	-30889	-31016
300	-30893	-31011
400	-30899	-31014
600	-30848	-31014

Table 5.2.1 Initial and final energies for the 4 25% doped CSIRO coating models calculated by CASTEP.




Figure 5.2.1 The Initial energies and energy changes due to due to the simulated anneal for the 25% doped models.

There is an almost identical trend seen in both the initial energies of the models and the energy change which resulted from the simulated anneal. This leads to the conclusion that despite starting from different initial structures, all 4 models converged (or came close to converging) to the same final structure. Which means that the subtle changes seen in the RDFs earlier, were not sufficient to lead to unique relaxed models. It also means that the assumption used earlier that the stoichiometry is not a significant factor is not valid. This is an unfortunate result however, it does underline a previous statement that it is imperative to use high precision input data in semi-empirical modelling.

This conclusion is further emphasised by the average cohesive energy calculated from these models (figure 5.2.2) where the y-axis scale has been exaggerated to show the variance. The cohesive energies of these 4 models shows that they fell into two separate potential wells during geometry optimisation. If these 4 models were unique it would be expected that the cohesive energy would show trends similar to that seen for the LMA

coatings. This not the case, it is likely that whilst all 4 models covered (roughly) to the same structure, they differed just enough to optimise in to nearby potential wells thus explaining the differing cohesive energies.

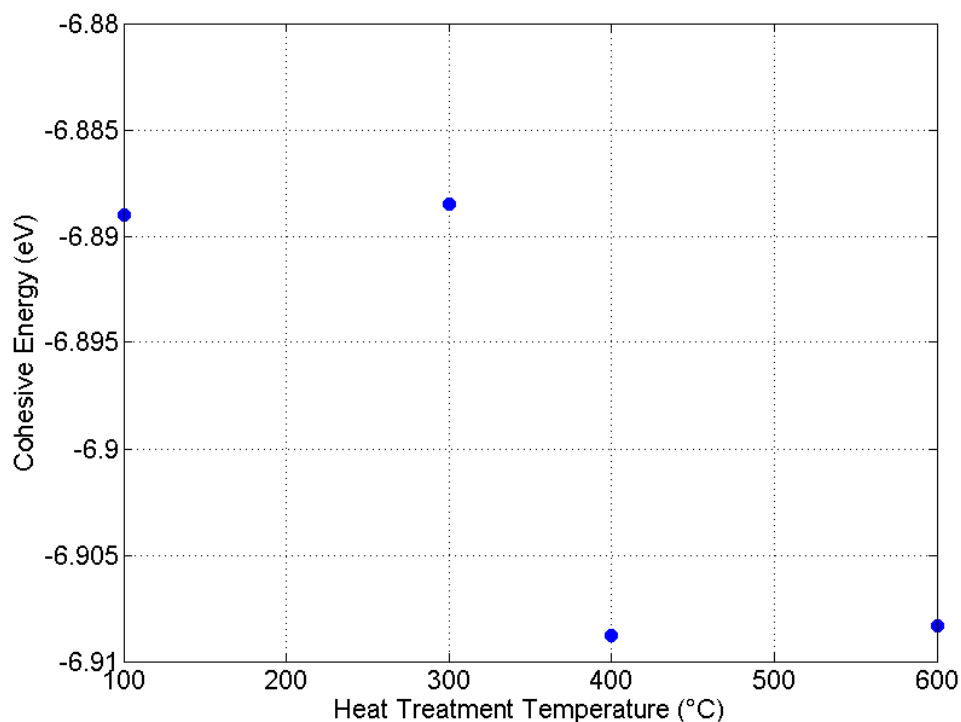


Figure 5.2.2. Average cohesive energies for the 4 25% doped CSIRO coatings as a function of heat treatment temperature.

5.3 Chapter Summary

Heat treatment of amorphous Tantalum, Titania doped Tantalum and a host of other coating materials has been shown to improve their optical properties [43, 55, 56, 58]. It has also been shown to alter the mechanical [58] and structural [105] properties as well. It is of utmost importance to the field of gravitational wave research to understand the nature of the metamorphosis undergone by coating materials during heat treatment. This must be achieved from a mechanical and structural perspective so as to understand not just how the material properties change but also why.

Previous work on Tantalum has shown that the mechanical loss of the coating changes considerably at cryogenic temperatures [58]. However structural studies have shown little change in the short range order of Tantalum due to heat treatment [105]. This suggests that either the change in the short range order is subtle or that there is no change in short range order and that the change actually occurs at longer range order. Recent studies on amorphous Silicon has revealed that although the short range order is virtually unaffected by heat treatment, while the medium range order changes significantly [127].

It is believed that with high quality diffraction data even subtle changes in short range order are detectable. Experimental studies using the TF20 microscope and Gatan 794 CCD camera have led to a higher quality dataset than has previously been achieved. Significant trends have been seen throughout the RDF analysis despite a decreased SNR due to lack of HDR data. These trends in the data have led to a possible explanation for a peak seen in the mechanical loss of Tantalum measured at 40K for coatings heat treated to 300°C. RDF studies have not revealed any clues to the cause of 25K peak in mechanical loss seen in coatings heat treated to 600°C. It is possible that separate mechanisms are responsible for the 25K and 40K mechanical loss peaks and that they are due to changes in the medium and short range orders respectively.

An Atomic modelling study of heat treated coatings has unfortunately proven fruitless. Atomic models for heat treated coatings have seemingly converged from unique initial parameters to a single structure. It is possible that for the moment the data simply is not of sufficient quality to provide sufficiently different initial parameter to avoid converging to a single final structure. Furthermore these models were run using the assumption that stoichiometry remains constant with heat treatment temperature. This was believed to be a valid assumption, however in light of new evidence may have not have been.

In the future, further studies of heat treatments effects on Tantalum and other coatings will require HDR to ensure high quality data. Moreover the stoichiometry of

each coating must be experimentally determined as well as the coating density so as to further improve the accuracy of both the RDFs and the atomic models. In conclusion, experimental studies have provided insight into the structural effect of heat treatment of Titania doped Tantalum. However, currently the data is of insufficient quality for accurate atomic modelling.

Chapter 6: General Conclusions

Gravitational waves, which are predicted by General Relativity, will become a valuable and integral member of the multiple messenger astronomy array. Moreover, they will allow astronomers and astrophysicists to probe the most spectacular events in the Universe. It is therefore essential to develop detectors capable of frequent and accuracy detection of gravitational waves. This will be achieved through reduction of the noise sources which limit current detectors, and further limiting those which will affect future detectors. The key noise source (and focus of this thesis) is thermal noise, caused by fluctuations due to finite temperature. Through mechanical excitation, electron diffraction studies and semi-empirical modelling it is expected that the mechanisms behind this noise can be identified and thus reduced. The key material of interest for this thesis is Titania doped Tantalum, since it has been shown that Titania doping lowers mechanical loss in Tantalum [63] however, how this is achieved was not understood.

Here an established modelling method has been employed for the purpose of rooting out the structural cause of mechanical loss. Various limitations of the method have been pointed out and methods to improve it have been developed and will be employed in the future. Structural changes in Titania doped Tantalum as a function of doping percentage and heat treatment temperature have been investigated and some possible causes of mechanical loss in these material have been suggested from the results.

Conclusions on the improvement of the techniques

Experimental constraints have been set on this work due to the equipment available and the materials and properties to be studied. Some of those constraints have been identified in this thesis, and methods to address them have been employed. A limitation has been seen in the SIS CCD camera which limits q_{\max} and hampers the maximum achievable SNR at q_{\max} . Without replacing the camera, a technique to increase q_{\max} and the SNR at q_{\max} has been employed. This technique is called HDRi, and whilst the

technique was not developed during this project as it is a standard imaging technique and widely used in conventional photography, it has hitherto not been extensively used in electron diffraction. It has, however, been adopted and employed and the advantages of using HDRi are clear. HDRi significantly increased the maximum SNR and dynamic range leading to much less variation in the RDFs due to detector noise, and therefore improved reliability. This will lead to more precise results in the future and HDRi will be employed as standard in future work.

Methods of increasing q_{\max} have been investigated and could be employed in the future as standard, though these typically require HDRi in order to produce viable results. The necessity of high q diffraction data has been shown here, which thus outlined the need to use methods capable of achieving high q_{\max} . The method outline in chapter 3 could be employed in future work, however simpler methods such as off-centre diffraction or lower camera lengths are also viable though the latter presents other issues. From the results presented in chapter 3 an estimated $q_{\max} \sim 25$ should be sufficient to converge the RDF peaks and with HDRi an acceptable SNR could be achieved at this range. In the future, experimental techniques should incorporate this conclusion and aim to achieve $q_{\max} \sim 25$. A further increase in q_{\max} could be achieved however, it is preferable to maintain a good q resolution and high SNR, and to sacrifice these in order to achieve $q_{\max} > 25$ would be ill advised.

Brief forays into X-ray diffraction and X-ray reflectometry have been undertaken and the results suggest that these methods would greatly complement current experimental techniques. In this thesis neither of these techniques produced viable data, however, and the reasons as to why this happened are clear, although the advantages of these techniques are obvious and future work will need to include these tools. X-ray diffraction is routinely used for the same type of semi-empirical modelling as is detailed here with electrons. Electrons however, have some distinct advantages over X-rays, the most important being the ability to focus the electron beam to a nanoprobe. Facilities exist for further X-ray diffraction work and, the increased q_{\max} theoretically achieved through these setups will provide further insight in to the coating atomic structure. X-ray

reflectometry, properly executed and analysed, will allow future work to be constrained by experimentally-determined coating densities.

Density functional theory modelling is a well-established tool for the atomistic simulation of materials. Through very simple methods, the computational expense of the DFT modelling has been significantly decreased compared to previous work in this field. By decreasing the simulation time intelligently, the simulations can be run ~2.5 times faster without loss of accuracy in the final result. This has also allowed, for the first time, models to be ran with fully optimised parameters. These parameters, which effectively increase the computational expense by a factor of 200, have proved very beneficial to the reliability and believability of the final result. Optimised parameter DFT simulations employing a simulated anneal technique have been shown to result in a structure much closer to that of the initial RDF than those done with the previous non-optimised parameters, as used in the first published work from Glasgow on the atomic structure of amorphous Ta₂O₅ [105]. In nature, of course, a structure must exist which is both energetically favourable and which matches the RDF, as this is what is measured experimentally. However, in the past experimental alignment and energetic relaxation seemed to work against each other. Using optimised parameters, it has been shown that the natural state of the structure may be achievable through simulation. Assuming appropriate computational resources are available in the future, these optimised parameters will be able to be employed as standard.

Conclusions on the structure of Titania doped Tantal

Advanced LIGO is designed to use multilayer coatings of alternating layers of SiO₂ and TiO₂ doped Ta₂O₅, as low and high refractive index materials respectively. Tantal has been shown to have good optical properties within the design specifications of an advanced detector such as aLIGO. It has also been shown that the dominant source of loss will also be from the Tantal layers of the coating. Titania has been employed as a dopant, as it was discovered that it can lower the mechanical loss of the structure without

adversely affecting the optical properties. These are all known quantities, what remains unknown is the source of the mechanical loss in Tantalum and Titania doped Tantalum.

This thesis has focused on the study of Tantalum and the effect of doping and heat treatment on the structure and how it relates to mechanical loss. What has been found is that local order changes significantly with doping, nearest neighbour bonds change in length and abundance as well as more closely approaching homogeneity. Experimentally, these attributes follow the mechanical loss very closely, though the picture is not clear without modelling. The modelling results have shown that as doping increases, Ti atoms occupy previously Ta-dominated positions in the planar fragments. Moreover the percentage of the doped planar fragments increases (relative to undoped fragments) with doping. The relative abundance of the doped planar fragments follows the loss very closely, suggesting a strong relationship. In structural terms, the overall structure of the coating becomes more flexible with doping and thus will more readily, accept defects, becoming less dissipative. This is due to the increased flexibility of the doped fragment, which, with its increased abundance, will increase the overall flexibility. It seems therefore, that the mechanical response of a material is related to the overall flexibility of its structure and thus it may be possible to decrease loss by increasing flexibility further.

Heat treatment of Tantalum and Titania doped Tantalum also has a strong influence on the mechanical loss. Past studies of Tantalum loss as a function of heat treatment have been unfruitful because, it seems, that the change in the structure (at least in the short range) is subtle. Here, Titania doped Tantalum, heat-treated to various temperatures was investigated in the hopes of determining their relationship. Strong trends were seen in the RDF analysis which could suggest a source of a peak seen at 40K for samples heat treated to 300°C. Certainly, there is a subtle change in the short range order of Titania doped Tantalum due to heat treatment, however it seems that even with data of the highest quality this change is difficult to detect and characterise. Unfortunately modelling of these coatings did not prove fruitful, however the cause of this has been identified and with further work, it should be possible to model these structures.

Possible Future Work

From the results presented in this thesis it seems probable that a mechanism for loss in Titania doped Tantalum has been identified. This is, of course, not necessarily the only mechanism of loss but it is most likely an important one. In the future, steps should be made to confirm these conclusions and begin to engineer low-loss coatings. Firstly a larger range of coating compositions should be used, 10-15 coatings with doping ranging from 0-100%. These coatings should be heat treated to a standard temperature, such as 600°C for a fixed amount of time, perhaps 24 hours, as was the case of the LMA coatings. This could lead to crystallisation in the highest dopant specimens and this would have to be investigated and an optimum temperature/heat treatment time would need to be determined. An increased dataset would allow the trends seen in the LMA coatings to be resolved further.

In response to the possible loss mechanism in the LMA coatings, new ways of reducing the loss should be investigated. This should be done by creating new coatings doped to 30% with alternative dopants to Titania. So far two possible replacement dopants have been suggested, these being Zirconia (ZrO_2) and Hafnia (HfO_2). These have been determined through 0 K vibrational simulations to create more flexible doped fragments than Titania. Therefore these coatings should be produced, and their mechanical losses and atomic structures characterised. If the loss follows the predicted trend then it should be lower in these coatings than in Titania doped Tantalum. If this is so, and the structure shows $ZrTaO_2$ or $HfTaO_2$ fragments, then it would confirm the hypothesis that structural flexibility is a controlling factor in the mechanical response. Finally, in all further investigations, new experimental techniques should be adopted to ensure $q_{max} \sim 25$ is achieved as standard and that HDRi is also used as standard.

References

1. Maxwell, J.C., *A Treatise on Electricity and Magnetism* Vol. 1 and 2. 1904.
2. Ellis, R.K., W.J. Stirling, and B.R. Webber, *Qcd and Collider Physics*. 2003: Cambridge University Press.
3. Fermi, E., *Versuch Einre Theorie Der B-Strahlen. I.* Zeitschrift fur Physik, 1934. **88**(3-4): p. 161-177.
4. Newton, I., *Philosophiæ Naturalis Principia Mathematica*. 1st ed. 1686.
5. B R Stella, H.-J.M., *Y(9.46 Gev) and the Gluon Discovery (a Critical Recollection of Pluto Results*. European Physical Journal H 2011. **36**: p. 203-243.
6. Gell-Mann, M., *Symmetries of Baryons and Mesons*. Physical Review, 1962. **125**: p. 1067-1084.
7. Weinberg, S., *A Model of Leptons*. Physics Review Letters, 1967. **19**: p. 1264-1266.
8. Einstein, A., *Die Grundlage Der Allgemeinen Relativitatstheorie*. Annalen der Physik, 1916. **49**.
9. D I Blokhintsev, F.M.G.P., *Neutrino Hypothesis and Conversation of Energy*. Pod Znamenem Marxisma, 1934. **6**: p. 147-157.
10. Saulson, P.R., *Fundamentals of Interferometric Gravitational Wave Detectors* World Scientific, 1994. **Singapore**.
11. Weber, J., *Detection and Generation of Gravitational Waves*. Physical Review, 1960. **117**: p. 306-313.
12. Weber, J., *Observation of the Thermal Fluctuations of a Gravitational-Wave Detector*. Physical Review Letters, 1966. **17**: p. 1228-1230.
13. Weber, J., *Gravitational-Wave-Detector Events*. Physical Review Letters, 1968. **20**: p. 1307-1308.

14. Weber, J., *Evidence for the Discovery of Gravitational Radiation*. Physical Review Letters, 1969. **22**: p. 1320-1324.
15. Braginskii, A.B., Manukin, E.I., Popov, V.N., Rudenko, and A.A. Khorev, *Upper Limit of Density of Extraterrestrial Gravitational Radiation*. Zhurnal Eksperimentalnoi i Teoreticheskoi Fiziki, 1974. **66**(3): p. 801-812.
16. Drever, R.W.P., J. Hough, R. Bland, and G.W. Lessnoff, *Search for Short Bursts of Gravitational Radiation*. Nature, 1973. **246**(5432): p. 340-344.
17. Levine, J.L. and R.L. Garwin, *Absence of Gravity-Wave Signals in a Bar at 1695 Hz*. Physical Review Letters, 1973. **31**(3): p. 173-176.
18. Tyson, J.A., *Null Search for Bursts of Gravitational Radiation*. Physical Review Letters, 1973. **31**(5): p. 326-329.
19. Gertsenshtein, M.E., *On the Detection of Low Frequency Gravitational Waves*. Soviet Physics, JETP, 1962(16).
20. Weiss, R., *Electromagnetically Coupled Broadband Gravitational Antenna*. . Quarterly Progress Report 105, Electronics Research Laboratory, 1972.
21. Saulson, P.R., *Thermal Noise in Mechanical Experiments*. Physical Review D, 1990. **42**(8): p. 2437-2445.
22. Albert Michelson, E.M., *On the Relative Motion of the Earth and the Luminiferous Ether*. American Journal of Science, 1887. **34**: p. 333-345.
23. Hewitson, M., G. Heinzel, J.R. Smith, K.A. Strain, and H. Ward, *Principles of Calibrating the Dual-Recycled Geo 600*. Review of Scientific Instruments, 2004. **75**(11): p. 4702-4709.
24. Abbott, B.P., R. Abbott, R. Adhikari, P. Ajith, B. Allen, G. Allen, . . . L.S. Collaboration, *Ligo: The Laser Interferometer Gravitational-Wave Observatory*. Reports on Progress in Physics, 2009. **72**(7).
25. Grote, H. and L.S. Collaboration, *The Geo 600 Status*. Classical and Quantum Gravity, 2010. **27**(8).
26. Accadia, T., F. Acernese, F. Antonucci, S. Aoudia, K.G. Arun, P. Astone, . . . M. Yvert, *Status and Perspectives of the Virgo Gravitational Wave Detector*, in *Topics in*

Astroparticle and Underground Physics, E. Cocchia, et al., Editors. 2010, IOP Publishing Ltd: Bristol.

27. Arai, K., R. Takahashi, D. Tatsumi, K. Izumi, Y. Wakabayashi, H. Ishizaki, . . . L. Collaboration, *Status of Japanese Gravitational Wave Detectors*. *Classical and Quantum Gravity*, 2009. **26**(20).
28. Kawabe, K. and L.S. Collaboration, *Status of Ligo*, in *Taup2007: Tenth International Conference on Topics in Astroparticle and Underground Physics*, K. Inoue, A. Suzuki, and T. Mitsuri, Editors. 2008.
29. Adhikari, R., *Second Generation Gravitational Wave Detectors*. In *Laser Science, 2007*: p. Page LMA3.
30. *Advanced Ligo Reference Design*, in *LIGO document LIGO M060056-08-M2007*, Advanced LIGO Team.
31. Penn, S.D., P.H. Sneddon, H. Armandula, J.C. Betzwieser, G. Cagnoli, J. Camp, . . . C.C. Vassiliou, *Mechanical Loss in Tantalum/Silica Dielectric Mirror Coatings*. *Classical and Quantum Gravity*, 2003. **20**(13): p. 2917-2928.
32. Fresnel, A., *Memoir on the Diffraction of Light*. *The Wave Theory of Light - Memoirs by Huygens, Tound and Fresnel*. 1819: American Book Company.
33. Drever, R., *A Gravity-Wave Detector Using Optical Cavity Sensing*, in *Proceedings of the 9th International Conference of General Relativity and Gravitation*. 1983, Cambridge University Press. p. 265.
34. Aufmuth, P. and K. Danzmann, *Gravitational Wave Detectors*. *New Journal of Physics*, 2005. **7**.
35. McClelland, D.E., S.M. Scott, M.B. Gray, D.A. Shaddock, B.J. Slagmolen, A. Searle, . . . C. Hollitt, *Second-Generation Laser Interferometry for Gravitational Wave Detection: Aciga Progress*. *Classical and Quantum Gravity*, 2001. **18**(19): p. 4121-4126.
36. Pitkin, M., S. Reid, S. Rowan, and J. Hough, *Gravitational Wave Detection by Interferometry (Ground and Space)*. *Living Reviews in Relativity*, 2011. **14**.
37. Beccaria, M., M. Bernardini, S. Braccini, C. Bradaschia, A. Bozzi, C. Casciano, . . . Z. Zhang, *Relevance of Newtonian Seismic Noise for the Virgo Interferometer Sensitivity*. *Classical and Quantum Gravity*, 1998. **15**(11): p. 3339-3362.

38. Hughes, S.A. and K.S. Thorne, *Seismic Gravity-Gradient Noise in Interferometric Gravitational-Wave Detectors*. Physical Review D, 1998. **58**(12).
39. Braginsky, V.B., M.L. Gorodetsky, and S.P. Vyatchanin, *Thermodynamical Fluctuations and Photo-Thermal Shot Noise in Gravitational Wave Antennae*. Physics Letters A, 1999. **264**(1): p. 1-10.
40. Rowan, S., J. Hough, and D.R.M. Crooks, *Thermal Noise and Material Issues for Gravitational Wave Detectors*. Physics Letters A, 2005. **347**(1-3): p. 25-32.
41. Robertson, N.A., G. Cagnoli, D.R.M. Crooks, E. Elliffe, J.E. Faller, P. Fritschel, . . . P. Willems, *Quadruple Suspension Design for Advanced Ligo*. Classical and Quantum Gravity, 2002. **19**(15): p. 4043-4058.
42. Punturo, M., M. Abernathy, F. Acernese, B. Allen, N. Andersson, K. Arun, . . . K. Yamamoto, *The Third Generation of Gravitational Wave Observatories and Their Science Reach*. Classical and Quantum Gravity, 2010. **27**(8).
43. Martin, I.W., *Studies of Materials for Use in Future Interferometric Gravitational Wave Detectors*, in *Department of Physics and Astronomy 2009*, University of Glasgow.
44. Edelstein, W.A., J. Hough, J.R. Pugh, and W. Martin, *Limits to Measurement of Displacement in an Interferometric Gravitational Radiation Detector*. Journal of Physics E-Scientific Instruments, 1978. **11**(7): p. 710-711.
45. Einstein, A., *Ist Die Tragheit Eines Körpers Von Seinem Energieinhalt Abhängig?* Annalen der Physik, 1905. **18**: p. 639-643.
46. Caves, C.M., *Quantum-Mechanical Noise in an Interferometer*. Physical Review D, 1981. **23**(8): p. 1693-1708.
47. Buonanno, A. and Y.B. Chen, *Quantum Noise in Second Generation, Signal-Recycled Laser Interferometric Gravitational-Wave Detectors*. Physical Review D, 2001. **64**(4).
48. Harms, J., Y.B. Chen, S. Chelkowski, A. Franzen, H. Vahlbruch, K. Danzmann, and R. Schnabel, *Squeezed-Input, Optical-Spring, Signal-Recycled Gravitational-Wave Detectors*. Physical Review D, 2003. **68**(4).
49. Harry, G.M., H. Armandula, E. Black, D.R.M. Crooks, G. Cagnoli, J. Hough, . . . S.D. Penn, *Thermal Noise from Optical Coatings in Gravitational Wave Detectors*. Applied Optics, 2006. **45**(7): p. 1569-1574.

50. Brown, R., *A Brief Account of Microscopical Observations Made in the Months of June, July and August, 1827, on the Particles Contained in the Pollen of Plants; and on General Existence of Active Molecules in Organic and Inorganic Bodies* in *Philosophical Magazine* 1828. p. 161-173.
51. Einstein, A., *On the Movement of Small Particles Suspended in a Stationary Liquid Demanded by the Molecular-Kinetic Theory of Heat*. *Annalen der Physik*, 1900. **17**.
52. Evans, M., 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*, 2008. **78**(10).
53. Basti, F., F. Frasconi, E. Majorana, L. Naticchioni, M. Perciballi, P. Puppo, . . . F. Ricci, *A Cryogenic Payload for the 3rd Generation of Gravitational Wave Interferometers*. *Astroparticle Physics*, 2011. **35**(2): p. 67-75.
54. Abernathy, M.R., S. Reid, E. Chalkley, R. Bassiri, I.W. Martin, K. Evans, . . . P. Seidel, *Cryogenic Mechanical Loss Measurements of Heat-Treated Hafnium Dioxide*. *Classical and Quantum Gravity*, 2011. **28**(19).
55. Reid, S., *Studies of Materials for Future Ground-Based and Space Based Interferometric Gravitational Wave Detectors* in *Department of Physics and Astronomy* 2006, University of Glasgow.
56. Crooks, D.R.M., *Mechanical Loss and Its Significance in the Test Mass Mirrors of Gravitational Wave Detectors* in *Department of Physics and Astronomy* 2003, University of Glasgow.
57. Tywford, S.M., *Developements Towards Low Loss Suspensions for Laser Interferometric Gravitational Wave Detectors* in *Department of Physics and Astronomy* 1998, University of Glasgow.
58. Martin, I.W., R. Bassiri, R. Nawrodt, M.M. Fejer, A. Gretarsson, E. Gustafson, . . . A.L. Woodcraft, *Effect of Heat Treatment on Mechanical Dissipation in Ta₂O₅ Coatings*. *Classical and Quantum Gravity*, 2010. **27**(22).
59. Martin, I.W., E. Chalkley, R. Nawrodt, H. Armandula, R. Bassiri, C. Comtet, . . . A.L. Woodcraft, *Comparison of the Temperature Dependence of the Mechanical Dissipation in Thin Films of Ta₂O₅ and Ta₂O₅ Doped with TiO₂*. *Classical and Quantum Gravity*, 2009. **26**(15).

60. Punturo, M., M. Abernathy, F. Acernese, B. Allen, N. Andersson, K. Arun, . . . K. Yamamoto, *The Einstein Telescope: A Third-Generation Gravitational Wave Observatory*. Classical and Quantum Gravity, 2010. **27**(19).
61. Ligo. *Einstein Gravitational Wave Telescope, Proposal to the European Commission, Frame Work Programme 7*. 2007.
62. Netterfield, R.P., M. Gross, F.N. Baynes, K.L. Green, G.M. Harry, H. Armandula, . . . S.D. Penn, *Low Mechanical Loss Coatings for Ligo Optics: Progress Report*. 2005: p. 58700H-58700H.
63. Harry, G.M., M.R. Abernathy, A.E. Becerra-Toledo, H. Armandula, E. Black, K. Dooley, . . . A. Remillieux, *Titania-Doped Tantalum/Silica Coatings for Gravitational-Wave Detection*. Classical and Quantum Gravity, 2007. **24**(2): p. 405-415.
64. Cockayne, D.J.H., *The Study of Nanovolumes of Amorphous Materials Using Electron Scattering*, in *Annual Review of Materials Research*. 2007. p. 159-187.
65. Thomson, J.J., *Cathode Rays*. Philosophical Magazine 1897. **44**.
66. Chadwick, J., *Bakerian Lecture, the Neutron*. Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences, 1933. **142**.
67. Chadwick, J., *Possible Existence of a Neutron*. Nature, 1932. **129**.
68. Compton, A., *A Quantum Theory of the Scattering of X-Rays by Light Elements*. Physical Review, 1923. **21**: p. 483-502.
69. W Friedrich, P.K., M Von Laue. *Interferenz-Erscheinungen Bei Röntgenstrahlen*. in *Mathematisch-Physikalischen Classe der Königlich-Bayerischen Akademie der Wissenschaften*. 1912. Munich.
70. Guinier, A., Ann. Phys. Paris, 1939. **12**: p. 161-237.
71. J Hillier, R.F.B., *Microanalysis by Means of Electrons* J. Appl. Phys, 1944. **15**: p. 663-675.
72. M Knoll, E.R., *Beitrag Zur Geometrischen Elektronenoptik. I*. Annalen der Physik, 1932. **404**.
73. D E Jesson, S.J.P., Proc. Roy. Soc. A., 1995. **449**.

74. David B Williams, C.B.C., *Transmission Electron Microscopy*. 2nd ed. 2009: Springer.
75. Olympus Soft Imagin Solutions (Sis) GmbH, G.; Available from: <http://www.olympus-sis.com>.
76. Inc., G.; Available from: http://www.gatan.com/products/digital_imaging/products/K2/index.php.
77. Amelio, G.F., M.F. Tompsett, and G.E. Smith, *Experimental Verification of Charge Coupled Device Concept*. Bell System Technical Journal, 1970. **49**(4): p. 593-+.
78. Boyle, W.S. and G.E. Smith, *Charge Coupled Semiconductor Devices*. Bell System Technical Journal, 1970. **49**(4): p. 587-+.
79. Netterfield, R. *Review of Thin Film Deposition Techniques*. 2001. Optical Society of America.
80. J Ayache, L.B., J Boumendil, G Ehret, D Laub, *Sample Preparation Handbook for Transmission Electron Microscopy*. 2010: Springer Verlag.
81. Ahn, C.C., *Transmission Electron Energy Loss Spectroscopy in Materials Science and the Eels Atlas* 2nd ed. 2004: Wiley-VCH.
82. Roditi International, L.; Available from: <http://www.roditi.com/>.
83. Warren, B.E., *X-Ray Diffraction*. 1990: Dover Publications.
84. Kirkland, E.J., *Advanced Computing in Electron Microscopy*. 1998: Springer.
85. Doyle, P.A. and P.S. Turner, *Relativistic Hartree-Fock X-Ray and Electron Scattering Factors*. Acta Crystallographica Section a-Crystal Physics Diffraction Theoretical and General Crystallography, 1968. **A 24**: p. 390-&.
86. Chen, Y., *Electron Diffraction Analysis of Amorphous Ge₂Sb₂Te₅*, in *Materials Department* 2010, University of Oxford: Oxford.
87. Warren, B., *X-Ray Diffraction*. 1969: Addison-Wesley: MA.
88. Lorch, E., *Neutron Diffraction by Germania, Silica and Radiation-Damaged Silica Glasses*. Journal of Physics C: Solid State Physics, 1969. **2**(2): p. 229.

89. Mcgreevy, R.L. and L. Pusztai, *Reverse Monte Carlo Simulation: A New Technique for the Determination of Disordered Structures*. *Molecular Simulation*, 1988. **1**(6): p. 359-367.
90. Borisenko, K.B., Y. Chen, S.A. Song, and D.J.H. Cockayne, *Nanoscale Phase Separation and Building Blocks of Ge₂Sb₂Te₅n and Ge₂Sb₂Te₅n₂ Thin Films*. *Chemistry of Materials*, 2009. **21**(21): p. 5244-5251.
91. Borisenko, K.B., Y. Chen, S.A. Song, D. Nguyen-Manh, and D.J.H. Cockayne, *A Concerted Rational Crystallization/Amorphization Mechanism of Ge₂Sb₂Te₅*. *Journal of Non-Crystalline Solids*, 2009. **355**(43-44): p. 2122-2126.
92. Cordero, B., V. Gomez, A.E. Platero-Prats, M. Reves, J. Echeverria, E. Cremades, . . . S. Alvarez, *Covalent Radii Revisited*. *Dalton Transactions*, 2008. **0**(21): p. 2832-2838.
93. Drabold, D.A., *Topics in the Theory of Amorphous Materials*. *European Physical Journal B*, 2009. **68**(1): p. 1-21.
94. Clark, S.J., M.D. Segall, C.J. Pickard, P.J. Hasnip, M.J. Probert, K. Refson, and M.C. Payne, *First Principles Methods Using Castep*. *Zeitschrift Fur Kristallographie*, 2005. **220**(5-6): p. 567-570.
95. Rae, A., *Quantum Mechanics*. 2007.
96. W Kohn, L.J.S., *Physical Review*, 1965. **140**.
97. Ziesche, P., S. Kurth, and J.P. Perdew, *Density Functionals from Lda to Gga*. *Computational Materials Science*, 1998. **11**(2): p. 122-127.
98. Perdew, J.P., K. Burke, and M. Ernzerhof, *Generalized Gradient Approximation Made Simple*. *Physical Review Letters*, 1996. **77**(18): p. 3865-3868.
99. Pedroza, L.S., A.J.R. Da Silva, and K. Capelle, *Gradient-Dependent Density Functionals of the Perdew-Burke-Ernzerhof Type for Atoms, Molecules, and Solids*. *Physical Review B*, 2009. **79**(20).
100. Cheung, D.L.G., *Structures and Properties of Liquid Crystals and Related Molecules from Computer Simulation*, in *Chemistry Department 2002*, University of Durham.
101. Yin, M.T. and M.L. Cohen, *Theory of Abinitio Pseudopotential Calculations*. *Physical Review B*, 1982. **25**(12): p. 7403-7412.

102. Kleinman, L. and D.M. Bylander, *Efficacious Form for Model Pseudopotentials*. Physical Review Letters, 1982. **48**(20): p. 1425-1428.
103. Vanderbilt, D., *Soft Self-Consistent Pseudopotentials in a Generalized Eigenvalue Formalism*. Physical Review B, 1990. **41**(11): p. 7892-7895.
104. Marzari, N., D. Vanderbilt, and M.C. Payne, *Ensemble Density-Functional Theory for Ab Initio Molecular Dynamics of Metals and Finite-Temperature Insulators*. Physical Review Letters, 1997. **79**(7): p. 1337-1340.
105. Bassiri, R., K.B. Borisenko, D.J.H. Cockayne, J. Hough, I. Maclaren, and S. Rowan, *Probing the Atomic Structure of Amorphous Ta₂O₅ Coatings*. Applied Physics Letters, 2011. **98**(3).
106. Feynman, R.P., Physical Review, 1939. **56**.
107. Rushton, P.P., *Towards a Non-Local Density Functional Description of Exchange and Correlation*, in *Departments of Chemistry and Physics 2002*, University of Durham.
108. A Dal Corso, A.P., A Baldereschi, R Car, Physical Review B, 1996. **53**.
109. M Fuchs, M.B., E Pehlke, M Scheffler, Physical Review B, 1998.
110. C M Bertoni, L.R., F Finocchi, Physica B, 1993. **185**.
111. Juan, Y.M., E. Kaxiras, and R.G. Gordon, *Use of the Generalized Gradient Approximation in Pseudopotential Calculations of Solids*. Physical Review B, 1995. **51**(15): p. 9521-9525.
112. Ortiz, G., *Gradient-Corrected Pseudopotential Calculations in Semiconductors*. Physical Review B, 1992. **45**(19): p. 11328-11331.
113. Holy, V., J. Kubena, I. Ohlidal, K. Lischka, and W. Plotz, *X-Ray Reflection from Rough Layered Systems*. Physical Review B, 1993. **47**(23): p. 15896-15903.
114. Als-Nielsen, J., *Elements of Modern X-Ray Physics*. 2001, New York: Wiley.
115. J Daillant, A.G., *X-Ray and Neutron Reflectivity: Principles and Applications*. 1999, Springer.
116. Tolan, M., *X-Ray Scattering from Soft-Matter Thin Films*. 1999: Springer.

117. Cockayne, D., D. Mckenzie, and D. Muller, *Electron-Diffraction of Amorphous Thin-Films Using Peels*. Microscopy Microanalysis Microstructures, 1991. **2**(2-3): p. 359-366.
118. Evans, K., R. Bassiri, I. Maclaren, S. Rowan, I. Martin, J. Hough, . . . Iop, *Reduced Density Function Analysis of Titanium Dioxide Doped Tantalum Pentoxide*, in *Electron Microscopy and Analysis Group Conference 2011*. 2012.
119. Nakamae, E., K. Kaneda, T. Okamoto, and T. Nishita, *A Lighting Model Aiming at Drive Simulators*. SIGGRAPH Comput. Graph., 1990. **24**(4): p. 395-404.
120. Spencer, G., P. Shirley, K. Zimmerman, and D.P. Greenberg, *Physically-Based Glare Effects for Digital Images*, in *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*1995, ACM. p. 325-334.
121. Ward, G. *High Dynamic Range Imaging*. in *Ninth Color Imaging Conference*. 2001. Scottsdale, Arizona: Ninth Color Imaging Conference: Color Science and Engineering: Systems, Technologies, and Applications.
122. Cebenoyan, S.G.a.C., *High Dynamic Range Rendering (on the Geforce 6800)*, 2004, Nvidia.
123. Reinhard, E.G.W., Sumanta Pattanaik, Paul Debevec, *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*. 2005, Westport, Connecticut: Morgan Kaufmann.
124. Cohen, J.a.T., Chris and Hawkins, Tim and Debevec, Paul E. *Real-Time High Dynammic Range Texture Mapping*. in *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*. 2001. London, UK.
125. Warren, B.E., *X-Ray Determination of the Structure of Glass*. J. Am. Ceram. Soc. , 1934. **17**.
126. J E Enderby, D.M.N., P a Egelstaff, Phil. Mag, 1966. **14**.
127. Treacy, M.M.J. and K.B. Borisenko, *The Local Structure of Amorphous Silicon*. Science, 2012. **335**(6071): p. 950-953.
128. Borisenko, K.B., B. Haberl, A.C.Y. Liu, Y. Chen, G. Li, J.S. Williams, . . . M.M.J. Treacy, *Medium-Range Order in Amorphous Silicon Investigated by Constrained Structural Relaxation of Two-Body and Four-Body Electron Diffraction Data*. Acta Materialia, 2012. **60**(1): p. 359-375.

129. Bassiri, R., *The Atomic Structure and Properties of Mirror Coatings for Use in Gravitational Wave Detectors*, in *School of Physics and Astronomy 2012*, University of Glasgow.
130. Laboratoire Des Materiaux Avances, L., Cnrs-In2p3, France. Available from: <http://lma.in2p3.fr>.
131. O L Anderson, H.E.B., *Ultrasonic Absorption in Fused Silica at Low Temperature and High Frequencies*. *Journal of the American Ceramic Society*, 1955. **38**: p. 125-131.
132. V B Braginsky, V.P.M., V I Panov, *Systems with Small Dissipation*, in *University of Chicago Press*1985: Chicago.
133. Abernathy, M.R., *Mechanical Properties of Coating Materials for Use in the Mirrors of Interferometric Gravitational Wave Detectors*, in *School of Physics and Astronomy*2012, University of Glasgow.
134. Peters, E. and H. Mullerbuschbaum, *A Titanium Tantalum Oxide of Ti-li - $Ti_{0.33}Ta_{0.67}O_2$ (Tita₂o₆)*. *Zeitschrift Fur Naturforschung Section B-a Journal of Chemical Sciences*, 1995. **50**(8): p. 1167-1170.
135. Mcguigan, D.F., C.C. Lam, R.Q. Gram, A.W. Hoffman, D.H. Douglass, and H.W. Gutche, *Measurements of Mechanical Q of Single-Crystal Silicon at Low-Temperatures*. *Journal of Low Temperature Physics*, 1978. **30**(5-6): p. 621-629.
136. J L Rodgers, W.a.N., *Thirteen Ways to Look at the Correlation Coefficient* *The American Statistician*, 1988. **42**: p. 59-66.
137. Cordero, B., V. Gomez, A.E. Platero-Prats, M. Reves, J. Echeverria, E. Cremades, . . . S. Alvarez, *Covalent Radii Revisited*. *Dalton Transactions*, 2008(21): p. 2832-2838.
138. Treacy, M.M.J. and J.M. Gibson, *Variable Coherence Microscopy: A Rich Source of Structural Information from Disordered Materials*. *Acta Crystallographica Section A*, 1996. **52**(2): p. 212-220.