



McGhee, Graeme Ian (2024) *Studies of new materials with low mechanical and optical losses for the coatings of future gravitational wave detectors*. PhD thesis.

<https://theses.gla.ac.uk/84711/>

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

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

This work cannot be reproduced or quoted extensively from without first obtaining permission 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

Enlighten: Theses

<https://theses.gla.ac.uk/>
research-enlighten@glasgow.ac.uk

Studies of new materials with low mechanical and optical losses for the coatings of future gravitational wave detectors

Graeme Ian McGhee

MSci

Submitted in fulfilment of the requirements for the
Degree of
Doctor of Philosophy

SCHOOL OF PHYSICS AND ASTRONOMY
COLLEGE OF SCIENCE AND ENGINEERING
UNIVERSITY OF GLASGOW



University
of Glasgow

February 2024

*To my parents,
Eleanor and Ian, who inspire me every day,
and to my younger self who decided to embark on this
challenging life experience.*

*“I am among those who think that science has great beauty. A
scientist in his laboratory is not only a technician: he is also
a child placed before natural phenomena which impress him
like a fairy tale.”*

– Marie Curie

Contents

| | |
|--|--------------|
| Acknowledgements | xvii |
| Preface | xx |
| Summary | xxiii |
| Declaration | xxvii |
| 1 Gravitational wave detection | 1 |
| 1.1 Introduction to gravitational radiation | 1 |
| 1.2 Gravitational wave strain amplitudes | 2 |
| 1.3 Why study gravitational waves? | 4 |
| 1.4 Sources of gravitational waves | 5 |
| 1.4.1 Compact binary coalescences | 5 |
| 1.4.2 Bursts | 6 |
| 1.4.3 Continuous sources | 7 |
| 1.4.4 Stochastic background | 7 |
| 1.5 Interferometric gravitational wave detectors | 9 |
| 1.6 Noise sources in current detectors | 12 |
| 1.6.1 Newtonian noise | 13 |
| 1.6.2 Seismic noise | 14 |
| 1.6.3 Quantum noise | 15 |
| 1.6.4 Thermal noise | 18 |
| 1.7 Current gravitational wave detectors | 19 |
| 1.8 Future gravitational wave detectors | 21 |
| 1.8.1 Near-future upgrades | 21 |
| 1.8.2 Additional terrestrial detectors | 24 |
| 1.8.3 Space-based detectors | 30 |

| | | |
|----------|---|-----------|
| 1.9 | Conclusions | 32 |
| 2 | The mirrors of gravitational wave detectors & models of thermal noise | 33 |
| 2.1 | The mirror design | 34 |
| 2.1.1 | Coating production | 36 |
| 2.2 | Optical loss requirements of GWD mirrors | 37 |
| 2.3 | Coating thermal noise | 38 |
| 2.3.1 | Brownian motion | 38 |
| 2.3.2 | Fluctuation-Dissipation theorem | 39 |
| 2.3.2.1 | Sources of internal dissipation | 39 |
| 2.3.2.2 | Harmonic oscillators with internal damping | 41 |
| 2.3.2.3 | Thermal noise associated with a single resonant mode | 43 |
| 2.3.3 | Thermal noise associated with multiple resonant modes | 45 |
| 2.3.4 | Inhomogeneous loss applied to coating thermal noise | 46 |
| 2.3.5 | Calculating coating thermal noise | 47 |
| 2.3.6 | Bulk and shear loss dependent coating thermal noise | 49 |
| 2.3.7 | Accounting for layer penetration | 51 |
| 2.3.8 | Accounting for effective applied stresses and strains | 52 |
| 2.3.9 | How CTN calculations shall be performed in this thesis | 54 |
| 2.4 | Some notable coating material candidates | 54 |
| 2.4.1 | Titania mixed with germania ($\text{TiO}_2\text{:GeO}_2$) | 55 |
| 2.4.2 | Titania mixed with silica ($\text{TiO}_2\text{:SiO}_2$) | 56 |
| 2.4.3 | Amorphous silicon (aSi) | 57 |
| 2.4.4 | Silicon nitride ($\text{SiN}_X \text{H}_Y$) | 58 |
| 2.4.5 | Aluminium gallium arsenide (AlGaAs) | 58 |
| 3 | Studies of titania-silica mixtures as a low thermal noise coating material | 59 |
| 3.1 | Introduction | 59 |
| 3.2 | Why titania-silica mixes? | 59 |
| 3.3 | Titania-silica coating deposition and composition | 61 |
| 3.4 | Titania-silica mechanical loss studies | 66 |
| 3.4.1 | Measuring loss via gentle nodal suspension | 67 |
| 3.4.1.1 | External factor to consider: gas damping | 71 |

| | | |
|----------|---|------------|
| 3.4.2 | Fitting ringdown data to extract sample loss | 72 |
| 3.4.2.1 | The modelled amplitude decay | 74 |
| 3.4.2.2 | Factors to consider before accepting the fitted losses | 78 |
| 3.4.2.3 | Measured mechanical losses | 79 |
| 3.4.3 | Modelling vibrating disks | 81 |
| 3.4.3.1 | Modelling a coated HR stack topology | 85 |
| 3.4.3.2 | Inferring doped material elastic properties | 86 |
| 3.4.4 | Extracting coating loss from measured losses | 89 |
| 3.4.5 | Titania-silica coating mechanical loss results | 92 |
| 3.4.5.1 | Coating loss of the 63.2% Ti HR stack through heat treatment | 95 |
| 3.4.5.2 | Coating loss of the 69.5% Ti HR stack through heat treatment | 98 |
| 3.4.5.3 | Titania-silica HR stack losses discussion | 101 |
| 3.5 | Directly measured CTN of the 69.5% Ti titania-silica HR stack | 103 |
| 3.5.1 | The directly measured CTN setup briefly explained | 104 |
| 3.5.2 | Direct CTN measurement results of 69.5% Ti HR stack through heat treatment | 105 |
| 3.5.3 | Estimating the loss of the pure silica layers and inferring an optimistic CTN | 109 |
| 3.6 | Table of defect formation | 113 |
| 3.7 | Summary of the titania-silica loss and CTN investigations | 114 |
| 4 | Optical loss studies of titania-silica HR coatings | 116 |
| 4.1 | Introduction | 116 |
| 4.2 | Optical absorption studies of the titania-silica HR coatings | 116 |
| 4.2.1 | Photothermal common-path interferometry | 117 |
| 4.2.2 | Absorption of the 69.5% titania-silica HR coating obtained on two PCI systems | 119 |
| 4.2.2.1 | Discussions of the equivalence of different sampling techniques and samples | 120 |

| | | |
|----------|--|------------|
| 4.2.2.2 | Absorption of the 69.5% titania-silica HR coating through heat treatment | 124 |
| 4.2.2.3 | Summary of the 69.5% titania-silica HR PCI results | 127 |
| 4.2.3 | Comparison with cavity absorption measurements | 128 |
| 4.2.4 | Comparison of two titania-silica HR coating compositions | 130 |
| 4.2.5 | Absorption summary and further discussions | 132 |
| 4.3 | Optical scattering studies of the 69.5% titania-silica HR coating | 134 |
| 4.3.1 | Factors which give rise to optical scattering | 134 |
| 4.3.2 | Using an integrating sphere to measure optical scatter | 135 |
| 4.3.3 | Integrating sphere scatter measurements | 141 |
| 4.3.3.1 | Integrating sphere scatter results for the titania-silica HR coatings | 143 |
| 4.3.4 | Angle-resolved scatterometer studies | 146 |
| 4.3.5 | Angle-resolved scatterometer technique | 147 |
| 4.3.6 | Angle-resolved scatterometer results | 153 |
| 4.3.6.1 | Analysis of images after different cleaning methods | 154 |
| 4.3.6.2 | Comparison of ARS images and integrating sphere maps | 156 |
| 4.3.6.3 | ARS image high scatter features at larger polar angles | 158 |
| 4.3.6.4 | ARS integrated scatter results for the titania-silica HR coatings | 160 |
| 4.3.7 | Comparison of the measured scatter from the different techniques | 164 |
| 4.4 | Titania-silica optical loss study conclusions | 167 |
| 5 | Titania-silica HR coating crystallisation investigations | 169 |
| 5.1 | Introduction | 169 |
| 5.1.1 | Raman spectroscopy brief principles | 170 |
| 5.2 | Titania-silica mix Raman spectroscopy experiment | 172 |
| 5.2.1 | Titania-silica mix Raman spectroscopy heat treatment study | 175 |
| 5.2.2 | Titania-silica mix Raman spectroscopy - extended discussions | 177 |
| 5.2.2.1 | Correction of spurious Raman spectra features | 181 |
| 5.3 | Titania-silica mix supplementary GI-XRD measurements | 185 |
| 5.4 | Titania-silica mix crystallisation conclusions | 185 |

| | | |
|----------|--|------------|
| 6 | Mechanical loss studies of titania-germania coatings | 187 |
| 6.1 | Introduction | 187 |
| 6.2 | Comparisons of loss measurement procedures at Caltech and U of G | 189 |
| 6.3 | Mechanical loss studies of LMA titania-germania | 195 |
| 6.4 | Studies of titania-germania single layers made via different processes | 199 |
| 6.5 | Mechanical loss studies of titania-germania HR stacks | 204 |
| 6.5.1 | Titania-germania HR stack loss results | 207 |
| 6.6 | Various topology stacks excess loss study | 210 |
| 6.6.1 | Comparing two 8-layer stacks of different material content | 211 |
| 6.6.2 | Investigating the effects of increasing interfaces on loss | 213 |
| 6.6.3 | Comparisons of the HR coating losses to a nominally identical material-split 10-layer stack | 215 |
| 6.7 | Titania-germania HR CTN estimates | 217 |
| 6.8 | Titania-germania loss study conclusions | 221 |
| 7 | aLIGO test mass characterisation studies | 224 |
| 7.1 | Introduction | 224 |
| 7.2 | Characterisation of LLO ETM point absorbers | 225 |
| 7.2.1 | The impact of point absorbers on gravitational wave detection | 226 |
| 7.2.2 | PCI absorption maps of LLO ETMs | 228 |
| 7.2.3 | Cleaning off some of the highest absorption point absorbers | 235 |
| 7.2.4 | LLO ETM point absorber conclusions | 238 |
| 7.3 | ITM figure and radius of curvature studies | 239 |
| 7.3.1 | Zygo interferometer measurement setup procedure | 241 |
| 7.3.2 | Interpreting the test optic surface profile from Zygo wavefront data | 245 |
| 7.3.3 | Removing artefacts in the Zygo data | 250 |
| 7.3.4 | ITM surface profile results | 252 |
| 7.3.5 | ITM radius of curvature results | 254 |
| 7.4 | Conclusions from aLIGO test mass studies | 259 |
| 8 | Thesis conclusions | 260 |
| | Bibliography | 266 |

List of Abbreviations

| | |
|----------------|--|
| adV | Advanced Virgo |
| aLIGO | Advanced Laser Interferometer Gravitational-Wave Observatory |
| AR | Anti-reflecting |
| ARS | Angle resolved scatterometer |
| BBH | Binary black hole |
| BNS | Binary neutron star |
| BRDF | Bidirectional reflectance distribution function |
| BS | Beam splitter |
| BWD | Binary white dwarf |
| Caltech | California Institute of Technology |
| CASI | Complete Angle Scatter Instrument |
| CBC | Compact binary coalescence |
| CE | Cosmic Explorer |
| CMB | Cosmic microwave background |
| CSU | Colorado State University |
| CSUF | California State University, Fullerton |
| CTN | Coating thermal noise |
| DECIGO | DECI-hertz Interferometer Gravitational wave Observatory |
| DS | Double-side |
| DW | Drag wipe |
| EFI | Electric field intensity |
| ERD-TOF | Elastic recoil detection with time-of-flight measurements |
| ET | Einstein Telescope |
| ETM | End test mass |
| ETMs | End test masses |
| FC | First Contact |

| | |
|---------------|---|
| FEA | Finite element analysis |
| FWHM | Full width half maximum |
| GeNS | Gentle Nodal Suspension |
| GEO600 | Gravitational European Observatory |
| GWB | Gravitational wave background |
| GWD | Gravitational wave detector |
| GWDs | Gravitational wave detectors |
| HF | High frequency |
| HOM | Higher order modes |
| HR | Highly reflective |
| IAD | Ion assisted deposition |
| IBS | Ion beam sputtering |
| IP | Ion plating |
| IPA | Isopropanol |
| IS | Integrating sphere |
| ITM | Input test mass |
| ITMs | Input test masses |
| KAGRA | Kamioka Gravitational Wave Detector |
| LF | Low frequency |
| LIGO | Laser Interferometer Gravitational-Wave Observatory |
| LISA | Laser Interferometric Space Antenna |
| LLO | LIGO Livingston Observatory |
| LMA | Laboratoire des Matériaux Avancés |
| MIT | Massachusetts Institute of Technology |
| NA | Numerical aperture |
| NEMO | Neutron Star Extreme Matter Observatory |
| NS-BH | Neutron-star black-hole |
| O1 | Observing run 1 |
| O3 | Observing run 3 |
| O4 | Observing run 4 |
| O5 | Observing run 5 |
| PCI | Photothermal common-path interferometry |

| | |
|-----------------|--|
| PD | Photodiode |
| ppm | parts per million |
| PSD | Power spectral density |
| PTA | Pulsar timing array |
| QWL | Quarter wave layer |
| RBS | Rutherford backscattering spectrometry |
| rms | root-mean-squared |
| ROI | Region of interest |
| RPN | Radiation pressure noise |
| SN | Shot noise |
| SQL | Standard quantum limit |
| SS | Single-side |
| TCS | Thermal compensation system |
| TIS | Total integrated scatter |
| WD-NS/BH | White dwarf-neutron star/black hole |

List of Figures

| | | |
|------|--|----|
| 1.1 | The effect of a gravitational wave on a ring of test particles. | 2 |
| 1.2 | Schematic depiction of a basic Michelson interferometer. | 9 |
| 1.3 | Schematic depiction of a modified dual recycling Fabry-Pérot Michelson interferometer scheme adopted in current GWD. | 10 |
| 1.4 | Design sensitivity curve of Advanced LIGO. | 13 |
| 1.5 | Visual representation of the constituents of quantum noise. | 15 |
| 1.6 | Timelines of the first three gravitational wave observing runs. | 21 |
| 1.7 | Strain sensitivity plots for the first 3 observing runs, and the projected sensitivities for O4 and O5. | 22 |
| 1.8 | Noise curves for Advanced LIGO, A+ and Cosmic Explorer. | 23 |
| 1.9 | The Einstein Telescope observatory. | 26 |
| 1.10 | Predicted strain sensitivities of future terrestrial detectors. | 27 |
| 1.11 | Astrophysical horizon of current and proposed future detectors for compact binary systems. | 28 |
| 1.12 | Sky localisation and electromagnetic counterparts of GW170817. | 29 |
| 1.13 | Orbit and spacecraft configuration for LISA. | 30 |
| 1.14 | Expected strain sensitivities of LISA and DECIGO in comparison to a representative terrestrial detector curve (LCGT). This figure is adapted from [111]. | 31 |
| 2.1 | Photographs of two ETMs previously installed at LIGO Livingston. | 33 |
| 2.2 | A schematic diagram of a multi-layer reflective coating stack. | 35 |
| 2.3 | A simple schematic diagram of an ion-beam sputtering (IBS) chamber. | 36 |
| 2.4 | Example thermal noise spectra showing the effect of the magnitude of the mechanical loss. | 44 |
| 3.1 | ERD-TOF atomic content data for the as-deposited HR stacks. | 64 |

| | | |
|------|--|-----|
| 3.2 | ERD-TOF silicon and titanium cation concentrations for the as-deposited HR stacks. | 65 |
| 3.3 | Diagram of a disk of thickness t balanced on a hemisphere lens of diameter D . Geometrical parameters useful to calculate the equilibrium condition in GeNS are shown. | 68 |
| 3.4 | Schematic illustrating the stability of a disk balanced on a sphere. | 69 |
| 3.5 | Top-down illustration of the GeNS setup. | 70 |
| 3.6 | Example ring-up and ringdown of the fundamental resonance of a disk sample. | 72 |
| 3.7 | Flow diagram showing the ringdown analysis process. | 76 |
| 3.8 | Comparison between the measured and best-fit ringdown data of an example mode of a disk. | 77 |
| 3.9 | Mechanical loss values measured for eight modes of a blank SiO ₂ disk. | 80 |
| 3.10 | Finite element analysis representation of the first twenty-two resonant mode shapes. | 83 |
| 3.11 | Finite element analysis representation of disk mode shapes. | 85 |
| 3.12 | Effective Young's moduli and Poisson ratios calculated for mixed TiO ₂ :SiO ₂ materials | 89 |
| 3.13 | Elastic strain energy ratios. | 92 |
| 3.14 | As-deposited coating mechanical losses as a function of frequency. | 93 |
| 3.15 | Coating mechanical losses as a function of frequency for the 63.2% Ti TiO ₂ :SiO ₂ /SiO ₂ HR stack. | 95 |
| 3.16 | Coating mechanical losses as a function of frequency for sample HR-1 B | 97 |
| 3.17 | Effect of heat-treatment on coating loss for the 69.5% Ti TiO ₂ :SiO ₂ /SiO ₂ HR stack. | 99 |
| 3.18 | A selection of coating mechanical losses as a function of frequency for HR-2 B. | 100 |
| 3.19 | Coating mechanical losses for both TiO ₂ :SiO ₂ /SiO ₂ stack designs as functions of post-deposition heat treatment temperature. | 102 |
| 3.20 | Directly measured CTN of the 69.5% Ti HR coating. | 106 |
| 3.21 | Images of TiO ₂ :SiO ₂ HR CTN samples after final heat treatments. | 107 |
| 3.22 | Coating mechanical losses for both TiO ₂ :SiO ₂ /SiO ₂ stack designs as functions of post-deposition heat treatment temperature. | 110 |
| 3.23 | Mechanical loss estimates of the SiO ₂ layers of the HR coating | 111 |

| | | |
|------|--|-----|
| 3.24 | Directly measured CTN at 100 Hz of the 69.5 % Ti HR coating (HR-2). | 112 |
| 4.1 | Schematic diagram of the PCI system at the University of Glasgow. | 118 |
| 4.2 | Absorption measurements of the HR-2 850°C, 100 hour heat treated sample. | 122 |
| 4.3 | Absorption map of the central 4 mm circular diameter of the HR-2 850°C, 100 hour heat treated sample | 122 |
| 4.4 | Optical absorption versus heat treatment temperature for 69.5% Ti HR stack. | 125 |
| 4.5 | Optical microscope images of the HR-2 sample annealed at 600°C exhibiting many macroscopic defects. | 126 |
| 4.6 | Optical absorption of the 69.5% Ti HR stack with heat treatment. | 129 |
| 4.7 | Optical absorption versus heat treatment temperature of the 69.5% and 63.2%Ti HR stack coatings. | 131 |
| 4.8 | Image of an integrating sphere. | 136 |
| 4.9 | Surface radiation exchange diagrams. | 136 |
| 4.10 | The Caltech integrating sphere setup. | 139 |
| 4.11 | Schematic of Caltech integrating sphere optical scatter set-up. | 142 |
| 4.12 | Surface maps of the measured scatter. | 143 |
| 4.13 | Histograms of the measured scatter | 144 |
| 4.14 | Angle-resolved scatterometer setup. | 148 |
| 4.15 | ARS images taken at different polar scattering angles. | 150 |
| 4.16 | Diagram showing the conical section of the hemisphere of backscatter that the camera samples during ARS measurement. | 152 |
| 4.17 | ARS images for the two titania-silica 69.5% high-n layer HR stack samples before and after cleaning with First Contact. | 154 |
| 4.18 | Comparison of ARS and integrating sphere scattering measurements | 157 |
| 4.19 | ARS image of the as-deposited First Contact cleaned sample. | 159 |
| 4.20 | BRDF data 69.5% HR stack samples before and after cleaning with First Contact. | 161 |
| 4.21 | Integrated scatter data for the as-deposited and 850°C heat treated samples taken with the ARS after cleaning with First Contact and an ionised nitrogen. | 163 |
| 5.1 | Comparison of the Raman spectrum of 69.5% Ti HR stack coating and liter- ature data. | 173 |

| | | |
|------|--|-----|
| 5.2 | Comparison of the Raman spectra of the 69.5% and 63.2% Ti HR coatings. . . | 175 |
| 5.3 | Raman spectra for the 69.5% Ti HR coating at different heat treatment temperatures. | 176 |
| 5.4 | Raman spectra for the 63.2% Ti HR coating at different heat treatment temperatures. | 176 |
| 5.5 | Modelled transmission spectra for wavelength range 400-550 nm, for the 63.2% HR coating. | 182 |
| 5.6 | Raman spectra acquired using different objective lenses. | 183 |
| 6.1 | Images detailing the Caltech GeNS setup. | 190 |
| 6.2 | Measured frequency spectra of an uncoated silica disk showing excited modes. | 191 |
| 6.3 | Flow diagram detailing the ringdown fitting procedure used. | 193 |
| 6.4 | Images of the ‘halo’ feature found on one TiO ₂ :GeO ₂ coated disk. | 196 |
| 6.5 | Graphs of the coating losses measured via GeNS of the four samples, as deposited (left), and after final heat treatment at 600°C for 108 hr (right) of all modes. | 196 |
| 6.6 | Individual graphs showing the coating losses measured via GeNS of each of the four samples. | 197 |
| 6.7 | Comparison of the coating losses of the TiO ₂ :GeO ₂ single layer coatings with different deposition parameters. | 200 |
| 6.8 | Comparison of the coating losses of the TiO ₂ :GeO ₂ single layer coatings with different deposition parameters, that were able to be annealed at 600°C for 108 hr | 201 |
| 6.9 | Image showing the haziness the LMA-2 single layer developed. | 202 |
| 6.10 | Layer structure of the 52L-A coating. | 206 |
| 6.11 | Layer structure of the 52L-B coating. | 206 |
| 6.12 | Coating mechanical losses of the two TiO ₂ :GeO ₂ /SiO ₂ HR stacks. | 207 |
| 6.13 | Coating mechanical losses of the two TiO ₂ :GeO ₂ /SiO ₂ 8 layer stacks. | 212 |
| 6.14 | Coating losses of TiO ₂ :GeO ₂ /SiO ₂ coatings with varying numbers of interfaces. | 213 |
| 6.15 | Coating loss comparison for the different stack designs studied. | 216 |
| 6.16 | Coating losses of the two 52-layer stacks and the 10-layer stack of similar material split. | 217 |

| | | |
|------|---|-----|
| 7.1 | Images of the Caltech PCI system with LLO ETM-Y installed. | 229 |
| 7.2 | Histograms of the measured absorption of each map. | 230 |
| 7.3 | Spatial coordinates of the LLO ETM-Y point absorbers. | 233 |
| 7.4 | Spatial coordinates of the LLO ETM-X point absorbers. | 233 |
| 7.5 | Bright-field and dark-field microscope images of the three large crater point absorbers | 234 |
| 7.6 | Bright-field microscope images showing cleaning method trials | 236 |
| 7.7 | Diagram showing the Zygo surface profile measurement configuration. | 239 |
| 7.8 | Diagrams showing the alignment stages to facilitate the surface profile measurement. | 243 |
| 7.9 | Images of the interference pattern produced by a test optic's surface reflections interacting with the transmission sphere reference wavefront. | 244 |
| 7.10 | Image showing the subtraction profile of the average of the first 20 surface profiles and average of second 20 surface profiles of ITM 11 orientated at 270°. | 249 |
| 7.11 | Diagram of the mount used to hold test mass optics for Zygo measurement. | 250 |
| 7.12 | Surface images taken of the three ITM optics obtained with the Zygo interferometer. | 252 |
| 7.13 | Zoomed in surface profiles from Figure 7.12 of the inner most 160 mm of each optic | 252 |
| 7.14 | Zoomed in image of the central 8 mm of ITM 11 shown previously in Figure 7.12. The faint residual interferometer artefact, its clones, and the interference they cause can be seen. | 253 |
| 7.15 | Fitted power values from the central 160 mm diameter of the three ITM optics. | 256 |
| 7.16 | Images of the surfaces of ITM 11. | 258 |

List of Tables

| | | |
|-----|---|-----|
| 3.1 | Properties of the $\text{TiO}_2\text{:SiO}_2$ coating samples. | 62 |
| 3.2 | Material properties of pure and mixed IBS SiO_2 and TiO_2 , as well as a composite stack property estimates. | 89 |
| 3.3 | Energy ratios calculated in COMSOL Multiphysics for the four coated $\text{TiO}_2\text{:SiO}_2\text{/SiO}_2$ HR stack samples approximating the stacks as a composite mono layer. | 91 |
| 3.4 | Defect formation in the $\text{TiO}_2\text{:SiO}_2$ coatings. | 113 |
| 4.1 | Optical absorption values obtained via PCI for the titania-silica mix high-n layer HR sack samples through heat treatment. | 132 |
| 4.2 | Comparisons of the measured scatter of the 69.5% Ti titania-silica mix HR stack coatings. | 165 |
| 5.1 | Raman shift peak positions recorded for the two HR coating stacks after crystallisation emergence. | 178 |
| 6.1 | List of all single layer coatings and one stack design measured in the initial $\text{TiO}_2\text{:GeO}_2$ investigations. | 204 |
| 6.2 | List of all $\text{TiO}_2\text{:GeO}_2\text{/SiO}_2$ coating stack designs explored in this study. | 210 |
| 6.3 | Coating thermal noise estimates from mechanical loss results of the 52-layer $\text{TiO}_2\text{:GeO}_2$ and SiO_2 based coating stacks. | 219 |
| 7.1 | All LLO ETM-Y and ETM-X point absorbers found via PCI mapping. | 232 |
| 7.2 | The absorption measured via PCI of the LLO ETM-Y point absorbers after cleaning with a cotton-bud and IPA/acetone. | 237 |
| 7.3 | List of the first fifteen Zernike polynomials, accompanied by a visual representation of each. | 247 |
| 7.4 | Radii of curvature R of the three ITM optics acquired with the Zygo interferometer. | 257 |

Acknowledgements

I began my PhD on October 1st, 2019, shortly after the epoch of gravitational wave astronomy had begun. It was an incredibly exciting time to embark on this chapter of my life, a time of immense promise, where aspirations to further improve the sensitivity of our detectors were matched by the excitement of potential new discoveries. These detectors, already among the most sensitive scientific instruments ever built, stood on the threshold of unveiling deeper mysteries of our Universe. To have had the opportunity to contribute meaningfully to the effort of expanding our reach into the cosmos, has been both a great privilege and an honour. Now, as I reach the end of this journey, and get ready to turn the page, there are many people to whom I owe my gratitude – those who have supported, guided, and inspired me to reach this point.

First and foremost, I would like to express my deep and sincere gratitude to my primary supervisor and mentor, Dr Iain Martin. His expertise, invaluable guidance, and unwavering support throughout my PhD were the foundation upon which my academic growth was built. His dedication, encouragement, and vision have not only enhanced my research, but have also inspired me to become a better scientist. There is only one Lord of the Coatings, and it turns out he does share power!

I also wish to acknowledge, with deep appreciation, the support of Dr Peter Murray and Dr Jessica Steinlechner, both of whom went above and beyond for me throughout my studies. I have learned so much from both of you, and I would not be the scientist I am today without your guidance. I would also like to thank Prof Sheila Rowan for her encouragement and support throughout this journey. I am deeply grateful to you all.

My time as a student at the Institute for Gravitational Research has led to some of the most stressful, as well as some of the most enjoyable, experiences of my life thus far, and I have made many lasting memories that I will cherish forever. I would particularly like to thank Dr Simon Tait for our many hours of “work” and “non-work” discussions: the former (usually) leading to lots of productive research. I’d also like to thank him for his dedication and hard work on the many projects we undertook together. I am also sincerely thankful to Mr Ross Johnston and Mr Stuart Hill for their help, support, and camaraderie throughout our studies.

I wish to express my gratitude to all of my office mates and friends who have supported and motivated me throughout my studies here at the University of Glasgow, and who made working here such a joy. Special shout-outs go to Dr Laurence Datrier, Dr Grégoire Lacaille, Dr Kristian Anastasiou, Dr Martin Sinclair, Dr Jennie Wright, Dr Michael Williams, Mr Jack Callaghan, Mr Christian Chapman Bird, Mr Naren Nagarajan, and Mr Thomas Wallace. I’m looking forward to celebrating with many of you after submission!

Throughout my studies, I was very fortunate to embark on two research fellowships, each lasting roughly three months. In my time at Claude Bernard University Lyon 1, France, I was able to sharpen my skills and develop as a researcher under the guidance of Prof Geppo Cagnoli. I wish to extend my sincere gratitude to him, Dr Valérie Martinez, and Mr Paul Dabadie for welcoming me so warmly into their group. I also undertook a fellowship at the California Institute of Technology, USA, where, through my work under Dr Gabriele Vajente and Dr GariLynn Billingsley, I was able to contribute to many projects vital for future gravitational wave detector upgrades. In addition, I would like to express my heartfelt thanks to Dr Camille Makarem and Dr Liyuan Zhang for their contributions. Thank you all for your assistance and mentorship in these endeavours, as well as your friendship. I am indebted to all of you – particularly Gari, who fronted the deposit for my road-trip rental car when the rental company wouldn’t accept it from my UK debit card!

I was also lucky to work with many wonderful collaborators from all over the world. I'd like to give a particular mention to Dr Jessica Steinlechner, Ms Viola Spagnuolo, and all the folks at Maastricht University, for all the collaborative projects, assistance, journal clubs, and remote tutorials we held together throughout the global pandemic and beyond. I'd also like to extend my heartfelt thanks to Prof Joshua Smith and Mx Oli Patane at California State University Fullerton for hosting my short research visit and welcoming our continued collaboration in the months that followed.

Finally, of course, there are my friends and family back home, whose support and love throughout this process have been unwavering, even when I would occasionally drop off the grid for several days during this thesis write up – I love you all. In particular, I wish to thank my mum, my dad, and my partner. Thank you for keeping me grounded while I dreamed of the stars.

Preface

This thesis is an account of work carried out between October 2019 and October 2023 at the Institute for Gravitational Research at the University of Glasgow, under the guidance of Dr Iain Martin, Dr Peter Murray and Professor Sheila Rowan. Two and a half months in early 2022 was spent at Claude Bernard University Lyon 1, under the supervision of Professor Geppo Cagnoli. Four months between August and December 2022 were spent at the California Institute of Technology, under the guidance of Dr Gabriele Vajente and Dr GariLynn Billingsley, with a short visit to California State University, Fullerton hosted by Professor Josh Smith.

Chapter 1 describes the nature of gravitational waves and their astrophysical sources. The principles of a gravitational wave detection using laser interferometry are outlined, the major noise sources in gravitational wave detectors are discussed.

Chapter 2 describes the design of the primary mirrors of gravitational detectors, the origin of coating thermal noise, one of the most significant noise sources in current gravitational wave detectors, limiting their performance at their most sensitive frequencies. The requirement to develop improved mirror coatings with lower thermal noise motivates the bulk of the work presented in this thesis.

Chapter 3 describes research into titania-silica mixtures, showing that this novel coating material has great promise for improving thermal noise in future detectors and detector upgrades. Mechanical loss measurements were carried out mainly by the author and Ms Viola Spagnuolo, with contributions from Dr Simon Tait, and Dr Peter Murray, and additional guidance from Dr Jessica Steinlechner. Analysis of the loss measurements made use of a code written by Dr Gabriele Vajente, modified by Dr Simon Tait. Composition studies using Rutherford Backscattering Spectroscopy and Elastic Recoil Detection with

time of flight measurements were carried out by Dr Martin Chicoine and Professor François Schiettekatte at the University of Montreal. Direct thermal noise measurements of the coatings were carried out by Dr Slawek Gras and Mr Nick Demos at the Massachusetts Institute of Technology. Dr Lukas Terkowski carried out spectrophotometry to measure the thickness of the coatings.

Chapter 4 describes investigations into the optical absorption and optical scattering of the titania-silica mixed coatings. Optical absorption was measured by the author with the assistance of Dr Liyuan Zhang at Caltech. Measurements of absorption at Glasgow were shared between the author, Dr Simon Tait, Mr Ross Johnston, and Dr Iain Martin. Scattering was measured at Caltech with the assistance of Dr Liyuan Zhang and at Fullerton with the assistance of Professor Josh Smith and Mx Oli Patane, who also carried out more measurements after the author's visit.

Chapter 5 presents studies of the crystallisation of titania-silica mixed coatings using Raman spectroscopy. These measurements were carried out by the author and Mr Paul Dabadie with input from Dr Iain Martin and Mr Andrew Monaghan in Glasgow and Dr Valerie Martinez in Lyon. Dr Mariana Fazio and Ms Chalisa Gier carried out grazing incidence x-ray diffraction measurements to further characterise the crystal structure in some of the coatings after heat treatment.

Chapter 6 contains an account of research carried out at Caltech into another promising candidate material: titania-germania. This work was carried out with the assistance and guidance of Dr Gabriele Vajente, with input from Professor Marty Fejer. The coatings were deposited by collaborators at Colorado State University and LMA. The author carried the mechanical loss measurements jointly with Dr Gabriele Vajente. The author analysed all of the data using, and further developing, code written by Dr Gabriele Vajente.

Chapter 7 describes studies undertaken at Caltech into mapping the absorption of the end test mass mirrors in operation at the LIGO Livingston observatory during the first-through-third observing runs of Advanced LIGO. The so-called point absorbers of these optics limiting detector performance were located and new cleaning method to mitigate

their impact was trialled. Alongside this, the surface profiles and radii of curvature of a selection of uncoated Advanced LIGO mirror substrates were determined. This work was carried out under the guidance of Dr GariLynn Billingsley and undertaken jointly with Dr Liyuan Zhang and Dr Camille Makarem.

Summary

Gravitational radiation was predicted as a consequence of Albert Einstein's general theory of relativity, according to which an accelerating mass should radiate energy in the form of gravitational waves. Einstein presented a new notion of gravity, postulating the idea that space-time curves in the presence of objects with mass. Gravitational waves arise due to any asymmetric acceleration of mass, and are the resulting fluctuations in the curvature of space-time.

A characteristic of gravitational radiation is its extremely weak interaction with matter, making observing it extremely challenging. The first evidence for their existence would not come for a further 66 years where Taylor, Weisberg, and Fowler, after the discovery of the first known binary pulsar system, tracked its orbital decay over 6 years to be in precise agreement with the loss of energy due to gravitational radiation. It would not be until nearly 100 years after Einstein published his theory, that gravitational waves would be directly observed for the first time in 2015 with the detection of the inspiral and merger of a binary black hole system. This detection was made by the twin Laser Interferometer Gravitational-wave Observatories (LIGO), detecting a change in their 4 km long cavity lengths induced by the passing gravitational wave of a thousandth the width of a proton, and would open a window into the otherwise invisible universe.

After almost 9 years, we are now firmly in the advent of gravitational wave astronomy with over 90 confirmed detections from the first three observing runs of the Advanced LIGO and Virgo detectors, with answers to many previous questions regarding the existence of certain astrophysical systems, and their dynamics, being provided. Through the study of gravitational waves we have also gained insights into the formation of heavy nuclei in the Universe, cosmological insights about its expansion rate, and have been able to provide further tests of general relativity. However, many questions still remain unanswered and for

a more in depth exploration of the cosmological and astrophysical sources of gravitational waves, and in order to facilitate detections from signals originating at greater distances, and further our understanding of the Universe, it is crucial to increase the sensitivity of current detectors.

Current gravitational wave detectors are kilometre-scale laser interferometers, which measure changes in their length as a gravitational wave passes. One of the dominant noise sources at the frequencies where the detectors are most sensitive is the thermal noise of the interferometer mirror coatings. This coating thermal noise must be reduced in order for fainter and more distant sources to be observed. To reduce coating thermal noise suitable materials with lower mechanical losses than current materials are required, and they must meet this requirement while still meeting the high optical standards demanded for detection.

The primary objective of the research presented in this thesis was to investigate new coating material candidates which could lead to improvements in sensitivity of future interferometric gravitational wave detectors via a reduction of their mechanical loss.

Chapter 1 introduces the nature of gravitational waves and their astrophysical sources. The basics of interferometric gravitational wave detectors and how these detectors operate is outlined. The major noise sources in current gravitational wave detectors are discussed, as are future detector designs and their potential scientific implications.

Chapter 2 describes the design of the primary mirrors of gravitational detectors, the origin of coating thermal noise, one of the most significant noise sources in current gravitational wave detectors, limiting their performance at their most sensitive frequencies. Various methods for calculating thermal noise are explained. A brief history of the development of the coatings of gravitational wave detector mirrors is given and list of potential new coating material candidates is presented. The requirement to develop improved mirror coatings with lower thermal noise motivates the bulk the work presented in this thesis.

Chapter 3 describes mechanical loss investigations of a novel coating material titania-silica and its potential to meet the thermal noise requirements is evaluated. Different compositions were studied at length, and a prime candidate material ratio was identified, with a full reflectivity stack design being commissioned and tested. The coating thermal noise estimates from the losses and directly measured coating thermal noise values are compared, showing good agreement. The results suggest that this material has the potential to significantly reduce coating thermal noise, to approximately 75% of the coating noise in the Advanced LIGO detectors. The use of this coating would require suppression of defect formation and demonstration of acceptable optical properties: investigations of these issues are presented in the following chapters.

Chapter 4 outlines investigations of the optical losses of the promising titania-silica mix coatings identified to have promising mechanical loss. Specifically the optical absorption was measured through varying stages of heat treatment via photothermal-commonpath interferometry, and found to be below 1 part per million after optimum heat treatment. This absorption is very close to meeting the requirements for a gravitational wave detector mirror coating, and is an excellent result for a new material with, as yet, no optimisation of the coating deposition parameters. The optical scatter of the coatings was also characterised in its as-deposited state, and after the heat treatment regime that produce the lowest coating thermal noise. These measurements were carried out with a high precision setup making use of an integrating sphere, and through measurement of the samples with an angle resolved scatterometer. The results obtained from both setups indicated remarkably low scattering of these coatings, reaching potentially lower levels than the coatings used in current detectors.

Chapter 5 details subsequent investigations of the crystallisation of titania-silica based coating stacks as a result of heat treatment. The heat treatment temperatures at which each coating crystallised was identified via Raman spectroscopy, and follow-up investigations into the nature of the crystallisation via grazing-incidence x-ray diffraction are also presented.

Chapter 6 presents extensive investigations into titania-germania based coatings. Titania-germania is another promising candidate material for future gravitational wave detectors showing remarkably low loss for an amorphous-oxide material. Investigations of the loss of initial prototype samples of this material, deposited by the current gravitational wave detector coating manufacturer, are reported. In addition to this, coatings made with deposition recipes which were altered so as to suppress defect manifestation were also characterised in terms of their mechanical loss. Full gravitational wave detector reflectivity coating stacks consisting of titania-germania and silica were produced with this optimised recipe, with their mechanical losses measured by the author. Differences between these multi-layer stack results and predictions from single layer were subsequently investigated. Similarly to in Chapter 3, estimates of the coating thermal noise improvement they would yield versus current detector coatings was quantified.

Chapter 7 reports on investigations by the author into characterising the topographic features and defects arising on the Advanced LIGO optics. Specifically, the point absorbers seen to arise randomly across the surface of the test masses were identified, with absorption maps of both end test masses installed at the LIGO Livingston Observatory in the first-through-third observing run being produced. From these, follow-up optical microscope surveys were conducted so as to better understand their nature. New cleaning techniques were also trialled, showing the potential to mitigate the influence of many point absorbers which arise on the on the test masses throughout their lifespan. The curvature and surface profiles of uncoated input test mass substrates were also characterised, with all being found to have their curvature lie within design specifications for Advanced LIGO.

Declaration

I declare that, except where explicit reference is made to the contribution of others, that this dissertation is the result of my own work and has not been submitted for any other degree at the University of Glasgow or any other institution.

Graeme Ian McGhee

Chapter 1

Gravitational wave detection

1.1 Introduction to gravitational radiation

Einstein’s general theory of relativity [1] outlined in a series of publications from 1907 – 1916 [1, 2, 3, 4, 5, 6, 7, 8], fundamentally changed the way we understand gravity and the Universe in general. To date, it has withstood a rigorous series of tests and serves as the current theory of gravitation. Both time and space are no longer thought of as flat and rigid in Einstein’s universe (as was implicitly assumed in Newton’s theory of gravity) but instead are distorted and warped by matter [3]. If the position of a mass changes, so does the curvature of the space-time in which it resides. This change is not instantaneous but travels at the speed of light. As a consequence, gravity vanishes in flat space-time, and is strongest where it is most curved, dictating the paths in which objects will travel. This is often summarised in the words of John Wheeler: “space-time tells matter how to move; matter tells space-time how to curve” [9, 10].

Gravitational radiation was predicted as a consequence of general relativity, wherein the acceleration of an asymmetric distribution of mass will cause fluctuations in the curvature of space-time. These “gravitational waves” are ripples in the fabric of space-time itself and carry this radiation throughout the Universe, and do so practically unimpeded by matter in their path. However, masses they encounter will experience fluctuating tidal forces orthogonal to the direction of travel of the gravitational wave as a result of the local space-time distortions as they pass. Gravitational waves couple weakly to matter, making these tidal forces imperceptible for all except those produced during the most catastrophic of astronomical events involving large masses and/or accelerations [11]. This

makes gravitational radiation exceptionally challenging to detect, but as a consequence of the same weak interaction with matter, gravitational radiation, unlike electromagnetic radiation, experiences next to no absorption or scattering [11], making it an invaluable tool for astronomy - the Universe is almost completely transparent to gravitational waves.

1.2 Gravitational wave strain amplitudes

In general, all systems with accelerating non-rotationally symmetric masses (for example waving ones hand) emit gravitational waves. They will manifest themselves by alternately compressing and stretching the space-time in the plane perpendicular to their direction of propagation, resulting in fluctuating length changes in space-time. Gravitational waves are quadrupolar, meaning that as space is stretched in one direction, it is compressed in the orthogonal one. This forced length compression/stretching is known as strain. The strain-amplitude of a gravitational wave (h) can be represented by the fractional change in distance (ΔL) between two points initially separated by some distance (L). Its effects are visualised in Figure 1.1 and it can be defined explicitly as [12, 13]

$$h = \frac{2\Delta L}{L}. \quad (1.1)$$

To calculate the strain of travelling gravitational radiation on space-time, the conservation of linear and angular momentum, as well as mass energy, must be applied. For our purposes, let us imagine the in-spiralling of two neutron stars. As they orbit around each other general relativity predicts that they will emit gravitational waves, which carry




Figure 1.1: The effect of a gravitational wave with the h_+ (top) and h_\times (bottom) polarisations, incident normal to the page, on a ring of test particles.

energy away from the binary system in all directions, consequently causing the stars to lose orbital energy and spiral closer together. At first the system starts out with certain amounts of gravitational potential energy, and orbital kinetic energy, and as the distance between the stars decreases, the potential energy is converted to kinetic energy, leading to an increase in orbital velocity of the stars. But at the same time the system is also generating more and more gravitational radiation which is carried away as gravitational waves, acting to reduce the total amount of energy in the in-spiraling system over time. The intensity of these waves, as with electromagnetic waves, follows an inverse-square law with distance D where they get weaker proportionally with $\frac{1}{D^2}$. As such, the amplitude of these waves decrease as $\frac{1}{D}$. As a result, the amplitude of the strain they impart is found to be [11]

$$h \approx \frac{G(E_{\text{GW}}/c^2)}{c^2 D}. \quad (1.2)$$

Here G is the gravitational constant, c is the speed of light, and E_{GW} is the part of the kinetic energy of the source, attributed to asymmetric motion of the particles in the source, leading to gravitational radiation. From $M = \frac{E}{c^2}$ [2] one understands mass is a form of energy. With an interaction such as this, the orbital properties and the masses of the original two neutron stars determine the fraction of the total original mass in the system that is emitted as gravitational radiation (as well as the fraction that becomes gravitational binding energy). The final bound mass will be larger than both original individual masses, but smaller than the total mass of both initial masses combined. The stars themselves are not losing mass during the inspiral though; the total energy in space-time is transforming from one form - i.e. two separate unbound masses - to another form - a single tightly bound mass, plus emitted gravitational radiation. In units of megaparsecs and solar masses, the magnitude of the strain imparted can be approximated from Equation 1.2 as [11]

$$h \sim 10^{-20} \left(\frac{E_{\text{GW}}/c^2}{M_{\odot}} \right) \left(\frac{\text{Mpc}}{D} \right). \quad (1.3)$$

In other words, to directly observe gravitational waves emitted with a solar mass worth of energy, at the distance of our closest neighbouring galaxy Andromeda (~ 0.9 Mpc away), one needs detectors sensitive enough to measure fractional length changes with better precision than 1 part in 10^{20} . Thorough and more mathematically in depth descriptions

of gravitational waves and the strain they impart can be found for example in [13], though this description presented should more than suffice to illustrate the observational challenge. The length variations in space-time created by gravitational waves are so minuscule, that only very large, rapidly accelerating masses can emit gravitational waves that are strong enough to be measured on Earth.

1.3 Why study gravitational waves?

The first indirect evidence for the existence of gravitational radiation would come with the discovery of the first binary pulsar system in 1974 [14], and the subsequent studies of its orbital dynamics in the decade that followed. Taylor, Weisberg, and Fowler [15, 16] tracked the system’s orbital decay over six years of measurements and observed it to be in precise agreement with the predicted loss of energy from general relativity due to gravitational radiation, with a corresponding reduction in the orbital time due to this of just over 1 second in 6 years. In 1993, Hulse and Taylor were awarded the Nobel Prize in Physics “for the discovery of a new type of pulsar, a discovery that has opened up new possibilities for the study of gravitation” [17] – an observation that helped to better understand the physics of the Universe and kindled interest in one day directly observing and studying the underlying gravitational phenomena.

In recent years, as will be discussed in the following sections, gravitational wave signals have now been directly measured by large scale interferometric detectors, opening a new observable window to the Universe. They allow for observation of events where the electromagnetic radiation cannot be detected (or there is none) such as those involving black holes. They also assist us in inferring the populations and composition of stars. Answers to many previous questions regarding the existence of certain astrophysical systems, and their dynamics, have been provided through their study [18, 19] and we have also gained insights into fundamental aspects of life such as the origins of heavy elements in the Universe [19, 20]. Gravitational wave detection data has also helped further investigate the expansion rate of the Universe [21], and has been able to provide further tests of general

relativity [22, 23, 24, 25, 26]. Though great scientific feats have been accomplished, many questions still remain, and indeed new questions have arisen. Through continuing to expand the field of gravitational wave astronomy, further secrets of our Universe will likely be unveiled.

1.4 Sources of gravitational waves

All asymmetrical mass acceleration in space-time will produce gravitational waves, but as exemplified through Equation 1.3, only events emitting the largest of energies have any hope of being observed here on Earth. As such, in order to detect gravitational waves producing sufficiently large strain amplitudes, we must turn to space. Gravitational radiation can be classified into four categories depending on the source of origin: those originating from compact binary coalescences, burst sources, continuous-wave sources, and sources from the stochastic background.

1.4.1 Compact binary coalescences

Compact binary coalescences (CBC) are the inspiral and merger of two compact objects. Compact objects, in this sense, are extremely dense astrophysical objects, possessing high mass relative to their radius [11]. The term encompasses black holes, neutron stars and white dwarf stars (stellar remnants). Thus far, the only observed gravitational wave signals were produced by CBCs. This type of transient signal occurs when two celestial bodies orbit each other until they coalesce. There are three distinct phases to this process: the inspiral, the merger, and the ringdown. Both bodies start out orbiting each other, and the orbit then slowly decays due to the energy transference to gravitational radiation, they then collide and merge, settling into one gravitationally bound object at the end of the inspiral phase, presenting a characteristic “chirp” signature just before it forms.

Gravitational waves were directly observed for the first time in 2015, with the discovery of the first stellar-mass binary black hole system GW150914 [18]. The final stages of inspiral were observed, with the binary black hole coalescing and emitting three solar masses of energy as gravitational waves, reaching a peak emission rate in the last few milliseconds

of the merger of approximately $\sim 4 \times 10^{49}$ W [18]. This was a level of power greater than all emitted electromagnetic radiation from stars in the observable universe combined over the same duration of time, yet produced peak strain amplitudes of only $\sim 1 \times 10^{-21}$ here on Earth. So far three flavours of CBC have been detected:

- Binary black hole mergers (BBH) [18, 27]
- Binary neutron star mergers (BNS) [19]
- Neutron star – black hole mergers (NS–BH) [28]

To date, between the first and third observing runs of advanced interferometric gravitational wave detectors, over 90 such CBC signals have been detected [29], of which two were confirmed to be BNS, three NS–BH, and the rest BBH. Mergers involving two white dwarfs (BWD) or a white dwarf - neutron star/black hole (WD–NS/BH) are yet to be observed. Despite this, the latter are in fact expected to be relatively common in the Milky Way [30], however their signatures are thought to be weaker, and the low characteristic frequencies of the gravitational waves they would emit are more favourable for future detection by space-based detectors [31].

1.4.2 Bursts

In the context of gravitational waves, a burst refers to a sudden and short-lived event that produces a significant, and sharp strain signal [32, 33]. Bursts, otherwise known as un-modelled/weakly modelled transients, have an extensive list of potential sources. Some sources of bursts include core-collapse supernovae, the coalescence of intermediate-mass black holes ($100M_{\odot} < M < 10000M_{\odot}$), magnetar flares, and pulsar glitches [32, 33]. When searching for burst signals, few assumptions are made on the shape of the gravitational waveform, and they are instead identified through the observation of excess power spikes in narrow detection bands [34]. The spikes are analysed for statistical significance above normal background levels and high levels of signal coherence between multiple detectors are required in order for one to be verified [32]. No present gravitational wave detections have been confirmed to originate from burst sources thus far.

1.4.3 Continuous sources

Continuous gravitational waves are those emitted over a long period of time at a by-and-large constant singular frequency [11, 33]. They can be generated by spinning neutron stars which possess axial asymmetry in their shape (i.e. those with bulges and localised regions of differing surface height). They can also be generated by binary black hole/star systems orbiting each other during the long time prior to merger. These sources are expected to produce comparatively weak gravitational waves when compared to transient signals such as CBCs since they evolve over much longer periods of time and usually arise from interactions with less acceleration [32]. As none have been detected to date, this puts an upper limit on the maximum strain from such signals. Hence, in the case of non-perfectly-spherical neutron stars, constraints on the surface height deformations/ellipticity have been refined [35].

1.4.4 Stochastic background

The gravitational wave background (GWB), often referred to as the stochastic background, is speculated to comprise of numerous primordial gravitational wave signals from the early universe, and myriad incoherent astrophysical signals [11, 33]. It can be visualised as a background level of static permeating the Universe, given rise to from a superposition of all these sources, identical in all parts of the sky. In this way it is akin to the cosmic micro-wave background (CMB), which is likely to be the leftover electromagnetic radiation from the Big Bang.

The Big Bang is a prime candidate for the forge in which many of these random processes required to generate stochastic gravitational waves (and the CMB), were created. As such the GWB may carry information about the history and origins of the Universe. The CMB was produced approximately 379,000 years after the Big Bang [36], when the Universe had cooled enough for electrons and protons to form atoms, and it illuminates the entire sky. Before this, the Universe was completely opaque, and so dense, that photons could not traverse through it unimpeded. Effectively, no electromagnetic signals were able to escape the dense pre-CMB Universe (otherwise known as the “epoch of recombination”), and it represents a limit to how far in the past we can observe with light. However, the timescales

for the formation of the earliest GWB sources are of the order 10^{-36} – 10^{-32} seconds post Big Bang [37]. As such, rather excitingly, through their study we could glean information about the Universe in its earliest states, inaccessible with electromagnetic radiation, a fraction of a second after its inception. One can also imagine locating previously unseen supermassive black holes, via decoupling and tracking their signals back to the origin point using the GWB. Further, mapping the structure of space-time with the GWB could also be achieved [38].

Primordial GWB sources produced in the Big Bang are not expected to be detectable by current terrestrial interferometric detectors as they would be too low in frequency. However, our ground based studies can constrain certain cosmological models that predict large amplitudes for the GWB, through all-sky, all-frequency searches [39, 40]. Recently there have also been breakthroughs in using pulsar timing arrays (PTA) to make initial probes of the stochastic background [41].

Searches for the stochastic GWB using pulsar timing arrays effectively compare the measured correlated responses between multiple pairs of Earth-pulsar radio signal arrival timing baselines with their expected correlated response to an unpolarised, isotropic stochastic gravitational-wave background (plotted as a function of the angle between the Earth-pulsar baselines) [42, 43]. This determines whether a signal from an unpolarised, isotropic background is present in the data. The key property of a PTA is that the signal from a stochastic GWB will be correlated across all the baselines, while signals from other processes will not [42]. In June 2023, four independent PTA collaborations verified, after over a decade of measurement, that the expected correlated response of pairs of Earth-pulsar timing baselines in their networks was indeed indicative of the existence of the GWB, with the measured correlations, consistent with the curve predicted by general relativity [41].

With the sources of gravitational waves defined we will now discuss interferometric detectors, of which this PhD work aims to improve the sensitivity.

1.5 Interferometric gravitational wave detectors

Over the past 9 years, more than 90 gravitational wave signals have been detected with kilometre-scale interferometric detectors. Terrestrial gravitational wave detectors (GWD) at their core are Michelson interferometers, the simplest design of which is shown in Figure 1.2. A laser beam is separated into two partial beams by a 50/50 beam splitter (BS). Afterwards the beams travel approximately equal distances, where each is reflected off a mirror to return to the beam splitter, carrying information about the distance it travelled. The beams then recombine and interfere with each other at the output port of the interferometer. The interference pattern observed will be determined by the difference in optical path lengths travelled by each laser beam. Through monitoring the laser power at the output port, the difference in path length can be continuously measured with high precision. If one then operates the interferometer near the dark fringe configuration - i.e. one return beam is almost completely out of phase with the other - then anything that causes the distances from the BS to the mirrors, known as the “arm lengths”, to change with respect to each other will be perceptible as a relatively large change in output signal. A passing gravitational wave at the right orientation will periodically act to alter the arm lengths and can in principle be observed by such an instrument if it is sufficiently sensitive. Note, the detectors are kept locked with active control systems to keep the distance and




Figure 1.2: Schematic depiction of a basic Michelson interferometer.

orientation of the optics fixed such that the output remains fixed at a point close to the dark fringe. This feedback signal is constantly monitored, with the gravitational wave signal actually observed as a change in the amount of feedback signal required to maintain detector lock.

The challenge consists of building an interferometer with the required sensitivity. To facilitate this, the basic design must be adapted. The general modifications implemented by all current terrestrial interferometric gravitational wave detectors are illustrated in Figure 1.3. The ethos of these improvements is to produce larger optical path lengths for the laser beams to travel and to generate higher power signals. Longer arms allow for smaller strains to be detected. For a given strain, the absolute arm length difference induced is proportionally larger for a longer arm via Equation 1.1, hence weaker signals can be detected. A higher total laser power in the interferometer arms also results in the sharpening of the interference fringes that appear when the return beams are superimposed. The sharper these fringes, the easier it becomes to identify passing gravitational waves.




Figure 1.3: Schematic depiction of a modified dual recycling Fabry-Pérot Michelson interferometer scheme adopted in current GWD.

To increase both the arm length and laser power in the arms, Fabry-Pérot cavities are constructed, where additional mirrors are installed near the beam splitter to facilitate multiple reflections of the laser beams. This keeps the beams circulating inside the interferometer for longer, and increases the distances they travel. For example, in the Advanced LIGO detectors this changes the path length in each arm from 4 km to around 1200 km [44] and it builds up the laser light within the interferometer, which increases detector sensitivity (more photons tracking the lengths of the arms). In addition, a “power recycling” mirror is installed at the input port of the interferometer. The laser initially travels through the mostly transparent side, and upon return from the arms, meets a highly reflective surface which directs most of the returning light back into the arms, serving to further increase the power inside the Fabry-Pérot arm cavities, without the need of a much higher initial input power. The power recycling mirror typically enhances the light power circulating in the arms by more than a thousandfold [45, 46].

Lastly, a “signal recycling” mirror is placed at the output port fulfilling the same function as the power recycling mirror, but it enhances the possible signal instead of the base power. The changing arm length of the interferometer leads to a frequency modulation of the signal, and this modulation results in signal sidebands at frequencies above and below the original laser frequency, which carry information about the gravitational wave. For GWDs, which typically operate in the dark port configuration, the gravitational wave signal sidebands are coupled to the anti-symmetric readout port of the interferometer [47]. The placement of the semi-transparent signal recycling mirror near the output port forms a coupled resonance cavity with the main interferometer. Different resonant conditions of this so-called signal recycling cavity, defined by the signal recycling mirror’s reflectivity as well as the cavity length, can improve the sensitivity of the GWD across different frequency bands [48, 49]. In a resonant sideband extraction scheme the main interferometer and the signal recycling mirror form such a cavity in which the gravitational wave signal of a particular frequency is resonant, and as such its signal is amplified [47].

The combination of all these enhancements to the standard Michelson, produces what is known as a dual recycling Fabry-Pérot Michelson interferometer - shown in Figure 1.3. To reduce the coupling of external vibrations to measurements in a GWD, all the primary optics, especially the arm cavity mirrors, are also suspended in quadruple-pendulum stages inside a vacuum tube of pressures $\sim \times 10^{-9}$ mbar [50] - one trillionth atmospheric pressure. In addition, active vibration isolation systems are also employed, all to assure the highest possible detector sensitivity is achieved.

1.6 Noise sources in current detectors

As for all waves, a key characteristic which can be used to describe a gravitational wave is its frequency. To determine different components of the signal acquired on the photodetector at the output port of the interferometer (and their physical origins), data is taken continuously, and the signal is analysed across frequency space to gauge the varying sensitivity of the detector. This generates a frequency dependent power spectral density (PSD) of the signal with units of $\frac{\text{m}^2}{\text{Hz}}$. However, for a gravitational wave it is the amplitude of the signal that is of interest (as one wants to measure the strain-amplitude it is imparting), and one ultimately wants to extract the fractional length change of the arm. As such, the frequency dependent detector sensitivity is conventionally represented as the amplitude spectral density (square-root of the PSD) divided by the detector arm length. This results in a strain noise spectrum across frequency with units $\frac{1}{\sqrt{\text{Hz}}}$ - an example of which for the current Advanced LIGO detectors is shown in Figure 1.4.

This section discusses the main sources of noise present in an interferometric GWD that gives rise to the minimal measurable strain across the range of frequencies which they are designed to observe gravitational wave signals. For the Advanced LIGO and Advanced Virgo detectors this range is from 10 Hz - 10 kHz [52, 53].




Figure 1.4: Simulated design sensitivity curve of the Advanced LIGO GWs showing the total noise (black) and the contributions of each major noise source individually – taken from [51].

1.6.1 Newtonian noise

Newtonian noise, also known as gravity gradient noise, arises as a result of the test mass mirrors directly coupling to local fluctuations in the gravitational field [54]. These fluctuations can result from changes in mass or density in the vicinity of the detector, including changes caused by seismic surface waves or localised variations, such as a vehicle or aircraft passing near the suspended mirrors.

Shielding the detector from gravitational gradient noise is impossible. However, models can be developed by deploying arrays of seismometers around the detector site, akin to those used for seismic isolation, to measure the Newtonian noise, such that it may be subtracted from the detector signal [55, 56]. This noise source is significant at frequencies ≤ 20 Hz and, since the test masses cannot be shielded from gravitational fields, it effectively imposes a lower limit on the low-frequency sensitivity attainable on the Earth [57]. The impact of this low-frequency noise source can, however, be mitigated by constructing the detector underground, farther from seismic waves travelling near the surface [54], though a more bold solution would be to place the detector in space - see Section 1.8.3.

1.6.2 Seismic noise

The Earth is a dynamic object, constantly experiencing smaller and larger movements, such as shifts in its crust and surface activities like earthquakes, varying wind speeds, ocean waves, and of course human activity. These all generate propagating seismic waves that can couple through the ground and induce vibrations in the optics of the detector [58]. The intensity of these seismic vibrations varies based on local geographic or geological factors such as the relative size of, and distance from, local population centres, and factors such as soil composition in the immediate vicinity [59]. Obviously this makes quiet and remote locations preferable. Even with such considerations of the environmental factors taken, it is crucial to significantly diminish the vibration coupling of the optics to achieve an acceptable noise level.

Overcoming the influence of the ever-present ground motion, characterized by micrometre-range amplitudes, is no small feat. To address this challenge, active and passive isolation systems are employed for the main interferometer optics. In Advanced LIGO, for instance, the active systems incorporate multiple vibration sensors and magnetic actuators to respond to coupled vibrations. Complementing the active isolation is a passive system featuring a quadruple pendulum stage, each of which contributes to further reducing the coupling of noise [60, 61]. This passive seismic damping is achieved by exploiting the principles of a simple pendulum. For any pendulum system, the transfer function from the suspension point to the suspended mass decreases as $\frac{1}{f^2}$ above the pendulum resonance (where f here is the frequency of a coupled oscillation). Consequently, suspending the mirrors on a series of four pendulums greatly reduces the impact of this noise source [62].

Seismic noise becomes problematic for detection at frequencies ≤ 10 Hz, and much like with Newtonian noise, its impact can be mitigated by building the detector underground [54], or in juxtaposition, placing it in space free from all terrestrial vibrations, allowing for interferometric gravitational wave detection well below 1 Hz - this will be discussed further in Section 1.8.3.

1.6.3 Quantum noise

Quantum noise is a direct manifestation of the Heisenberg uncertainty principle, which states that two complementary observables cannot be measured simultaneously to infinite accuracy [63]. It consists of two parts: photon shot (counting/sensing) noise and radiation pressure (back-action) noise [64], and it arises due to the quantization of photons comprising the interferometer laser beam. Its impact is the most wide-spanning of the gravitational wave detector noise sources in frequency space, which can be seen in Figure 1.4.

The emission of photons by the laser is an inherently random process and as such the number of photons detected in a given time will fluctuate following Poisson statistics. So a photodetector placed in the laser beam path to record the power levels would observe a Poisson distribution in the intensity over time. The fluctuation in the detected (photon) counts it would measure is known as the quantum shot noise (SN) [65].

Now consider that in the interferometer the laser beam is used to probe the position of/distance to the suspended mirrors at the end of its arms by reflecting from them. Each photon has both energy and momentum and will impart a ‘kicking’ force on the mirrors as it is reflected, transferring a portion of its momentum. The test masses in a gravitational wave detector are suspended such that they can be considered as quasi-free particles and as such will respond freely to forces acting on them. Hence, using a laser to probe the mirror positions will intrinsically disrupt the mirrors, causing them to move and fluctuate




Figure 1.5: Visual representation of the constituents of quantum noise. (Left) Shot noise: the fluctuations in number of photons at a given time will lead to a different power/voltage being read out on a photodetector. (Right) Radiation pressure noise: fluctuation in momentum transfer from reflected photons will lead to varying displacement of the suspended mirror Δx .

in position due to the imparted momentum they receive from the laser beam. Since the input power of the laser beam fluctuates, so too will the amount of transferred momentum to the mirror and hence quantum radiation pressure noise (RPN) arises. Illustrations of both shot and radiation pressure noise effects are shown in Figure 1.5

The strain amplitudes of the two quantum noise constituent parts can be expressed by the following equations [64]

$$h_{\text{SN}}(f) = \frac{1}{L} \sqrt{\frac{\hbar c \lambda}{2\pi P}} \quad (1.4)$$

$$h_{\text{RPN}}(f) = \frac{1}{mf^2L} \sqrt{\frac{\hbar P}{2\pi^3 c \lambda}}, \quad (1.5)$$

with f being frequency, L being the rest arm length to a mirror, m the mirror mass, P the average input optical power from the laser, λ is its central wavelength, \hbar is the reduced Planck's constant, and c is the speed of light.

Notice that altering the input power improves either shot or radiation pressure noise at the detriment of the other. Ultimately, for every observation frequency there is therefore an optimum power which leads to equivalent magnitude contributions from the shot noise and radiation pressure noise at that specific frequency. This can be found by setting $h_{\text{SN}} = h_{\text{RPN}}$ and solving for power. One can see the radiation pressure noise will drop as the square of frequency, and that it can also be reduced without increasing the shot noise by employing larger mirror masses. A larger mirror mass would lower the total quantum noise level in the \sim few Hz – few-tens-of-Hz range where it dominates detector sensitivity. However, this is easier said than done as larger mirrors would require thicker suspension fibres and these would introduce a larger suspension thermal noise effect (see the next section). Also there is a technological limit to the size high quality mirrors can be manufactured. So a trade-off in increasing the mirror mass to reduce radiation pressure noise would also introduce new challenges needing to be overcome.

In the previous consideration, at some time (t), laser light is used to measure the relative position of a mirror (x), but the momentum its photons impart perturb this position making subsequent measurements of position at some time later (t') intrinsically less accurate by the Heisenberg uncertainty principle. Similarly phase and amplitude signals are not commutable. The quantum noise of an interferometer is defined as the (uncorrelated) sum

of photon shot and photon radiation pressure noises. In the context of an interferometric measurement, shot noise can be considered a noise on the phase of the signal, and radiation pressure a noise acting directly on the amplitude of a signal [66, 67]. The previous description also provided a more simplistic view of the quantum noise. In actuality, in a GWD it arises from zero-point fluctuations of the vacuum field which enters through the dark port of the interferometer [68]. A field entering through the dark port affects the two arms differentially and these fluctuations give rise to a fluctuating force (RPN) and affect the phase sensitivity (SN) [68, 69]. As an example, imagine some light entering the interferometer from the port with the photodetector instead of the laser port (see Figures 1.2–1.3). If this light had the right phase to increase the intensity in one arm, it would decrease the intensity in the other arm. Thus, this additional light would produce a force which alters the signal. Although there is no light source which injects radiation from this port, inevitably vacuum fluctuations incident from this direction will manifest. These vacuum fluctuations, interfere with light from the laser, leading to the fluctuating force and phase noise [68, 69].

The standard quantum limit (SQL) defines the lower bound envelope of the quantum noise at all frequencies, for all optical powers circulating in an interferometer of a specific design configuration [70, 71]. Its strain amplitude is defined from the previous two equations as

$$h_{\text{SQL}}(f) = \frac{1}{\pi f L} \sqrt{\frac{\hbar}{m}}. \quad (1.6)$$

For a laser interferometer with sufficient suppression of classical noise sources, quantum shot noise will place a lower boundary on the response of the detector in the higher frequency regime, and radiation pressure will dominate at lower frequencies. For current detectors the quantum shot noise does indeed dominate at all frequencies above a few-hundred Hz. However it is conceivable the components and configuration of an interferometer may be adjusted to produce a lower quantum noise level. Due to the parameter that a standard Michelson interferometer measures (the continuous position of the mirrors), this configuration will always be fundamentally limited by quantum noise. Despite this, through novel techniques, the SQL can be surpassed: these techniques include optical squeezing [72], speedmeter interferometry [73], and optical springs [74].

1.6.4 Thermal noise

Thermal noise, also known as Brownian noise, arises as the interferometer operates at some finite temperature, T , and therefore all of the detector components (and the atoms that comprise them) have some thermal energy and thus will exhibit some thermal motion [75]. The equipartition theorem states that any mechanical system in thermal equilibrium will have a mean thermal energy value of $\frac{1}{2}k_B T$ associated with each degree of freedom (where k_B is the Boltzmann constant). This energy results in vibrations of the molecules comprising the system [76, 77]. In a GWD, this thermal energy will excite resonant mechanical modes of the mirrors, and their suspension system, producing a significant source of displacement noise in the position of the mirror faces [57] and hence the cavity length, which is translated into measured power fluctuations in the output port.

Thermal noise, arising from this thermally driven motion of the molecules in the GWD test masses and their suspensions, is one of the major limits to the sensitivity. As can be seen in Figure 1.4, the current generation of room (ambient) temperature GWD are limited in their most sensitive frequency band: 50 Hz – 100 Hz by thermally-induced vibrations arising in the detector mirror coatings [78], even more so than by quantum noise effects. Improving this ‘coating thermal noise’ is an ever present challenge for detector designs in both the near and far future.

Since the internal modes of the suspended mirror, for the most part, reside at high frequencies outwith the gravitational wave detection band, the thermal noise levels witnessed mostly arises from the tails of the resonance peaks. The thermal noise spectrum can be crafted to concentrate most of the thermal motion near the resonances, thereby minimizing off-resonance motion within the operational band of the detector [57]. This is accomplished by utilising materials with ultra-low mechanical losses, for the mirror substrate, its coatings, and suspension fibres (such as fused silica at room temperature). Mechanical loss is discussed more extensively in Section 2.3 of Chapter 2, but for now it can be thought of as the level of frictional damping in a material as a result of purely internal processes while it is oscillating.

Broadly speaking, minimising thermal noise involves reducing all sources of dissipation in the mirrors and their suspensions through the use of materials with low mechanical loss at the temperatures and frequencies of interest. Finding novel solutions to reduce the coating thermal noise levels for the next generation of gravitational wave detectors is the primary goal of this thesis work of which the challenges will be discussed in greater depth in Chapter 2.

1.7 Current gravitational wave detectors

The global interferometric gravitational wave detector network currently consists of five sites:

- the two Advanced Laser Interferometer Gravitational-wave Observatories (aLIGO), in Hanford, WA and Livingston, LA USA,
- the Advanced Virgo detector (adV), near Pisa in Italy,
- the Kamioka Gravitational Wave detector (KAGRA), near Hida in Japan,
- the Gravitational European Observatory (GEO600), near Hannover in Germany.

There are a great deal of similarities between the designs of each detector, though each has its own nuances. Each employs a large-scale dual recycling Michelson interferometer design operating at a laser wavelength of 1064 nm. The interferometer physical arm lengths range from folded 600 m (effectively 1.2 km) at GEO600, through 3 km at adV and KAGRA, to 4 km at the aLIGO sites [79]. For their main mirrors, they also utilise fused silica substrates held at room temperature with the exception of KAGRA, which operates at 20 K with sapphire substrates. Unlike the other four detectors KAGRA operates underground, and unlike the others GEO600 does not make use of Fabry-Pérot arm cavities, but instead uses only its power recycling mirror to build up power in the arms.

All these GWD utilise substrates and coatings of the highest quality to enable gravitational wave detection. Each mirror, or “test mass”, is made from large high purity ingots, and polished to specified curvatures, with minimal surface roughness to mitigate optical scattering down to the low parts-per-million (ppm) level [80]. The substrate diameters of the test masses vary across different detectors, with GEO600 employing ones of approximately 18 cm, KAGRA at 20 cm (sapphire), and both aLIGO and adV detectors at

34 cm and 35 cm, respectively [79]. These test masses are all coated with multiple alternating layers of near quarter-wavelength silica (SiO_2) and tantala (Ta_2O_5), providing the required reflectivity, and low optical losses essential for gravitational wave detection [81]. The aLIGO and adV detectors further improve upon the latter two by instead employing 27% titania (TiO_2) doped Ta_2O_5 layers, in place of pure Ta_2O_5 , to improve their coating thermal noise level at room temperature by $\sim 25\%$ [82].

To date over 90 gravitational wave signals, all resulting from compact binary coalescences, have been confirmed between the first and third observing runs from 2015–2020 of aLIGO and adV. In the third gravitational wave observing run (O3) we were able to detect signals from the farthest out in the Universe to date, with multiple BBHs estimated to reside at luminosity distances of around $\sim 5000\text{--}6000$ Mpc [83] - which is an order of magnitude further out than the first detection.

The detectors periodically undergo improvement, which leads to higher strain sensitivities and the detection of signals originating from greater depths in the Universe, as well as potentially from fainter sources located closer to Earth. One such improvement in the aLIGO detectors, as example, was the introduction of optical squeezing for O3, which was then subsequently upgraded to frequency dependant squeezing in O4. Squeezing reduces the quantum noise in the phase or amplitude quadrature of the laser at the expense of the other. One can inject squeezed vacuum states into the dark port of the detector to shape the quantum noise spectrum of the detector [84].

Observing run 4 (O4) started in mid-2023 and is expected to conclude in the mid-2025. Although no published data is currently available, with various improvements made between O3 and O4 facilitating heightened strain sensitivity, it was expected that the observing rates for BBH, BNS, and NS-BH mergers in O4 would increase to $79_{-44}^{+89} \text{ yr}^{-1}$, $10_{-10}^{+52} \text{ yr}^{-1}$, and $1_{-1}^{+91} \text{ yr}^{-1}$ respectively [85], which would be a significant increase in the number of observed events. The measured and predicted observing ranges for aLIGO, adV, and KAGRA during O1-O3 and O4-O5 are shown in Figure 1.6. Despite the already impressive capabilities of the current GWDs, to facilitate deeper and deeper probes of our Universe, further increase in their sensitivity in the coming observing runs is expected due to various ongoing research activities.




Figure 1.6: Timelines of the first three observing runs, detailing when each detector was active, as well as predicted timelines for O4 and O5. The measured/expected binary neutron star (BNS) coalescence detection range for LIGO, Virgo, and KAGRA detectors is also listed. Adapted from [86].

1.8 Future gravitational wave detectors

The aLIGO, adV, KAGRA, and GEO600 detectors discussed in the last section comprise the current global interferometric GWD network and are all known as 2nd generation detectors due to their incremental improvements over time to facilitate greater sensitivity. Due to operating at cryogenic temperatures, KAGRA is also sometimes referred to as a “generation 2.5” detector [87]. With the gap between observing run 4 (O4), concluding in mid-2025, and O5 currently set to commence in mid-2027, there are many planned further upgrades to the current generation of detectors. There are also plans to expand the global network further, with the construction of more, so-called, 2nd generation detectors, 3rd generation detectors, and even space-based observatories - these will be discussed in this section. Indeed, within approximately a decade’s time the number of interferometric detectors actively recording gravitational wave signals is planned to more than double.

1.8.1 Near-future upgrades

Before the start of O5, major upgrades are planned to the detectors currently in operation which will yield improved sensitivities across their whole frequency band, allowing for the detection of more distant and fainter sources – see Figures 1.6–1.7. These so-called aLIGO+ and adV+ improvements will be facilitated by employing heavier mirror substrates, and increased laser beam diameters, further developing the frequency-dependent




Figure 1.7: Strain sensitivities taken from [88] for the first 3 observing runs, and the projected sensitivities for O4 and O5. The latter being the aLIGO+ and adV+ design sensitivities. Note the L1/H1 in the LIGO graph for O3 denotes the minimum strain sensitivities achieved at the Livingston and Hanford sites respectively. The O4 and O5 projections are now slightly outdated, see Figure 1.6 for the latest values and estimates.

squeezing (and implementing it in adV+), and crucially for this thesis, lowering the mirror coating thermal noise (CTN) [89]. These upgrades for O5 are already in the process of being implemented and aim to improve the detector sensitivities by $\sim 50\%$ [89, 90]. Graphs showing the target O5 strain sensitivities of LIGO and Virgo are shown in Figure 1.7, note these graphs also show the O4 targets which adV is currently not able to fully meet – adV’s achieved BNS coalescence detection range of up to 80 Mpc is seen in Figure 1.6.

Further to the aLIGO+ upgrades, a set of subsequent upgrades are proposed for LIGO to be implemented post O5 where even heavier 100 kg test masses, and novel suspensions to hold them will be developed [92, 93]. This is known as the “LIGO Voyager” project, and it aims to further expand the detection range of the interferometers by a factor of 4–5 compared with aLIGO (improving the sensitivity by an additional factor of two versus aLIGO+), and half the low-frequency cutoff to 10 Hz [90, 94]. LIGO Voyager aims to eventually transform the current LIGO sites into 3rd generation detectors, through eventually converting them to facilitate operating at cryogenic temperatures (123 K), with corresponding appropriate coatings for the cryogenic regime, and increasing the optical power stored in the arms (3 MW). A graph showing the latest O4 strain sensitivities achieved, with projections for O5 (aLIGO+) upgrades, and the following LIGO Voyager upgrades, with the contribution of all the major noise sources in Voyager, is shown in Figure 1.8.




Figure 1.8: LIGO Voyager noise curve (grey) compared to aLIGO's measured noise curve during O4, and the intermediary A+ upgrade predictions for O5. The future planned Cosmic Explorer detector sensitivity in its stage 1 room temperature phase is also shown (CE1) - adapted from [91].

1.8.2 Additional terrestrial detectors

As well as the currently operating detectors, there are plans for four new ground based detectors to be constructed. These are:

- LIGO India; a 2nd generation clone of the aLIGO+ sites.
- The Neutron Star Extreme Matter Observatory (NEMO); a 2nd generation detector which will be based in Australia, operating at cryogenic temperatures (150 K/123 K) designed to be optimised to observing signals in the 2–4 kHz band outside of the most sensitive band of current detectors.
- Cosmic Explorer; a 3rd generation detector operating in the USA, first at ambient temperature, then eventually at cryogenics, but with arm lengths an order of magnitude longer than current detectors. Current plans are for two sites, one with 20 km long arms, and one containing arms 40 km long.
- The Einstein Telescope; a 3rd generation observatory operating in Europe (location to be determined). The site will be constructed underground and host three independent detectors, each with two interferometers, all operating in tandem, half at cryogenics (temperature to be finalised) and half at ambient temperature.

All of these new detectors are designed to be sensitive across different frequency bands, so we will soon be able to probe signals previously hidden to us.

LIGO India, scheduled to be in operation around 2030, is planned to be identical to the detectors in the USA. Having it join the global network will allow for better sky localisation and reduce detector downtime, leading to an improvement in the detection rates, and the completeness of gravitational wave surveys [95].

The NEMO detector, aimed to be in operation by 2035, will be optimized to study nuclear physics with merging neutron stars, and will potentially allow for the first gravitational-wave observations of isolated neutron stars, supernovae, and other exotica [96]. Neutron star mergers, such as GW170817, are expected to often produce rapidly-rotating remnant neutron stars, which will emit gravitational waves post-merger at around 2–4 kHz, outside of the range where all current detectors are most sensitive. In theory observing these should provide new insights into the extremely hot post-merger environment [96].

Cosmic Explorer (CE) is another, dual-recycled Michelson interferometer design, that will upscale the aLIGO+ technology to a much larger size. The arm length of the two proposed sites are 20 km and 40 km respectively, which would mean tunnelling underground in order to maintain a straight beam-line through the ground with the end stations still residing on the surface. In its initial stages (CE1), it is planned to employ much heavier 70 cm diameter, 320 kg fused silica test masses, with further plans to eventually operate at cryogenic temperatures (CE2), swapping to 320 kg silicon test masses, using LIGO Voyager research to inform CE2 design selection [97]. Cosmic explorer (and the Einstein Telescope) aims to extend the reach of gravitational wave astronomy to the edge of the observable universe, as shown later in Figure 1.11. Cosmic Explorer will be able to detect stellar-mass black holes at redshifts of up to $z \sim 20$ [97]. This extensive range will unveil, for the first time, the entire population of stellar-mass black holes, providing insights into their existence during the early stages of the universe when the first stars were forming. Cosmic Explorer will detect hundreds of thousands of black-hole mergers annually, acquiring data on their masses and spins. These observations will provide insights into the black-hole merger rate and the underlying stellar formation rate, as well as how both have evolved over the lifespan of the Universe, and their relationship with galaxy evolution. [98, 99, 100]

The Einstein Telescope (ET) project aims to construct an observatory which overcomes the limitations of current detectors by hosting more than a single interferometer on-site. Specifically, it will host three nested detectors, each comprised of two interferometers, each with arm lengths of 10 km, and all constructed underground to reduce seismic and Newtonian noise [102]. It is the only planned terrestrial detector to have multiple interferometers on the same site. Its three arms will form an equilateral triangle to host the six interferometers – see Figure 1.9. For each detector, one interferometer will operate at cryogenic temperature, detecting low-frequency (LF) gravitational wave signals from 3–40 Hz, while the other will operate at ambient temperature, searching for high-frequency (HF) signals above this (40 Hz–10 kHz) [103]. Having two interferometers in each detector




Figure 1.9: (Top right) The Einstein Telescope observatory optical layout, with three detectors forming an equilateral triangle with arm lengths of 10 km. (Bottom left) Diagram of one of the three identical detectors comprising of two interferometers, one optimised for low frequency (LF) and one optimised for high frequency gravitational wave searches. The core interferometer layout is still based on a dual recycling Michelson interferometer employing Fabry-Pérot cavities. This figure is adapted from [101].

allows for far less restrictive design space. The parameters of one interferometer can be more easily tailored to the regime of interest, at the expense of its sensitivity at other frequencies, whilst together providing high sensitivity over a wide frequency space from around 3 Hz–10 kHz, with order of magnitude higher sensitivity than current GWD [103].

The ET detection capabilities will be similar to that of Cosmic Explorer, but it will have higher sensitivity than Cosmic Explorer below 10 Hz, and lower sensitivity above 10 Hz [94], poising them to complement each other well. This heightened sensitivity at lower frequencies will enable ET to closely observe, with high signal-to-noise ratio, the inspiral phase of CBC mergers, as well as the onset of tidal effects during mergers [102]. This will yield unprecedented insights into the internal structure of neutron stars, and facilitate investigations into the fundamental properties of matter in previously untested




Figure 1.10: Predicted strain sensitivities of future terrestrial detectors (NEMO, Cosmic Explorer, and the Einstein Telescope), as well as upgrade plans for aLIGO (aLIGO+, and Voyager). Also displayed is the measured sensitivity of aLIGO during O3. This is taken from [104].

regimes (i.e. quantum chromodynamics in ultra-high density environments) [102]. The increase in sensitivity by over an order of magnitude, provided by both ET and CE (see Figure 1.10), will also allow them to better observe details from the merger and post-merger phases of all types of CBC events [97, 102].

All future planned terrestrial detectors and upgrades to the current sites expect to improve upon current detector sensitivities in different frequency regimes. Each will assist the others in the observation of different kinds of gravitational wave signals, and facilitate greatly increased detection rates, as well as the observation of signals from as of yet unobserved sources. Figure 1.10 shows the design sensitivity curves of each of these proposed new detectors and upgrades, perhaps serving as the best visualisation of the scale of sensitivity improvements to the global interferometric GWD network, whereas the astrophysical scope of the improvements is represented in Figure 1.11




Figure 1.11: Astrophysical horizon of current and proposed future detectors for compact binary systems. The lines indicate the maximum redshift at which a detection with signal-to-noise ratio 8 is achievable. The detectors shown are aLIGO during O3, the aLIGO+ upgrade targeted for O5 (A+), LIGO Voyager (Voy), Cosmic Explorer (CE), and the Einstein Telescope (ET). Simulated population of BBH and BNS mergers are represented by the white and yellow dots respectively. This figure is from [104].




Figure 1.12: The left image visualises the localisation of the GW170817 gravitational wave signal from just the two aLIGO sites (light-green) and with adV data included (dark-green), with corroborating estimates from the INTEGRAL and Fermi gamma ray observational satellites also shown (blue). The right side images show the manifested electromagnetic counterpart (the kilonova) as seen by the SWOPE observatory 10.9 hours after GWD detection (top image), and by the DTL40 observatory 20.5 days prior to the event showing no electromagnetic counterpart present (bottom). Figure from [105, 106].

Alongside the sensitivity improvements these future planned detectors and upgrades to current sites will yield, having more detectors operating simultaneously allows for more precise localisation of gravitational wave events in our sky. As example to illustrate this, upon the detection of BNS GW170817, the addition of the Virgo detector operating in tandem with both LIGO sites reduced the plausible region on the sky in which the event could have occurred from 190 deg^2 (from just both aLIGO sites) down to only 31 deg^2 – see Figure 1.12. Guided by these observations, collaborating astronomers were then able to very quickly (within 1 day) locate the source of the gravitational wave, and record the bright optical, infrared, and ultraviolet electromagnetic counterpart kilonova that followed this BNS merger (as well as the x-ray and radio components of the counterpart detected 9 and 16 days later). This event gave birth to the era of multi-messenger astronomy [106]. As such, the localisation benefit of having more detectors in the network is clearly apparent. A greater number of GWD observing simultaneously will enable a greater constriction of the source location in the sky, hence allowing collaborating astronomers to more quickly

locate and record potentially crucial data from any electromagnetic counterpart to the gravitational wave signal. The presence of more active gravitational wave observatories will also greatly reduce/prevent periods when three or more detectors are not online, minimising the loss of data during downtime of any individual detector.

To reach the predicted sensitivity targets in future detectors, advancements are needed in many areas. A primary one for all terrestrial detectors is the development of suitable mirror coatings which are able to produce the improved sensitivities in the cryogenic and room temperature regimes, without which the targets will be unobtainable. Assisting in this aspect for the next stages of upgrades is the driving goal of this thesis.

1.8.3 Space-based detectors

Future ground-based detectors will still be fundamentally limited in their detection capabilities below ≤ 1 Hz by seismic activity. Many interesting gravitational wave signals are speculated to exist in this regime, such as primordial gravitational wave signals originating from the Universe in its earliest stages of infancy as discussed in Section 1.4.4. Other sources include super massive black holes, and extreme mass ratio inspiral BBHs, along with more speculative sources such as cosmic strings. In order to overcome this seismic barrier and observe such sub-Hz signals, space-based detectors are essential.




Figure 1.13: (Left) The proposed orbit of the LISA space-based detector - from [107]. (Right) LISA's equilateral triangle configuration, with key components labelled - adapted from [108].

The **Laser Interferometric Space Antenna** (LISA) is a European Space Agency-led space-based interferometric gravitational wave detector, aiming to probe such low frequency gravitational wave signals in the 10^{-4} – 10^{-1} Hz range [107]. LISA consists of three identical drag-free spacecraft, oriented in an equilateral triangle configuration of arm lengths of 2.5 million km – see Figure 1.13. Each satellite will orbit around the sun in Earth’s Lagrange point 3 (L3), and will host two free-falling test masses, and two primary lasers, with the formation between the satellites maintained by affixed micronewton thrusters. Essentially, by directing each of the two primary lasers of one satellite at each vertex to one of the two test masses on both of the others, they will replicate three bound Michelson interferometers [107, 109]. The LISA mission was adopted by the European Space Agency in January 2024 and is planned for launch in 2035 [110].

The **DECI-hertz Interferometer Gravitational wave Observatory** (DECIGO) is a similar proposed Japanese space-based gravitational wave detector mission, set to launch in the 2030’s, but with its satellites orbiting three orders of magnitude closer at 1000 km apart. Whilst LISA targets the 10^{-4} – 10^{-1} Hz range, DECIGO will aim to be most sensitive in the 10^{-1} – 10^1 Hz range - thus bridging the gap between LISA and terrestrial detectors and opening yet another new observation window for gravitational waves [111]. The predicted sensitivity curves of LISA and DECIGO along with a representative terrestrial detector curve (analogous to aLIGO+) are shown in Figure 1.14.




Figure 1.14: Expected strain sensitivities of LISA and DECIGO in comparison to a representative terrestrial detector curve (LCGT). This figure is adapted from [111].

With both these space-based detectors operating in tandem, there would also be high prospects for observing a CBC signal in the earliest (low frequency) stages of inspiral with LISA, years before they would be detected in a terrestrial observatory, followed by DECIGO, then eventually by the plethora of terrestrial detectors. This would allow for a more complete study of inspirals, increasing our understanding of their dynamics. Multi-band observing such as this would assist greatly in planning terrestrial detector observing schedules and significantly lower the chance of missing signals that they could have observed.

1.9 Conclusions

In this chapter, gravitational waves, along with their astrophysical sources of origin were described. The first direct detection of gravitational waves was one of the great feats of physics, accomplished via kilometre-scale interferometric detectors. Subsequently, these same detectors begun the era of multi-messenger astronomy, and have provided further tests of general relativity, allowing humanity to better understand our Universe.

Major advances in physics and experimental techniques facilitated the ability of current gravitational wave detectors to observe such weak signals. They are, however, still limited in their capabilities. The origins of the major noise sources that currently limit detector sensitivity were discussed, and so too were the ways future upgrades and new detectors plan to overcome the hurdles they represent. These noise sources must be understood and novel solutions to mitigate them must be found, so that we can detect signals from fainter, more distant, and new kinds of sources.

This thesis focuses primarily on developing and investigating novel mirror coatings that can rise to the challenges set by future detector upgrade targets. Specifically, it aims to explore coating materials and designs which reduce the coating thermal noise level, the most dominant source of noise in the most sensitive frequency band of current detectors, whilst still remaining optically suitable, and it aims to find viable solutions for improvement for the most immediate round of upgrades to current detectors. This is crucial to advance the field by facilitating more sensitive detectors which can detect signals they otherwise could not, allowing us to better hear the cosmic symphony.

Chapter 2

The mirrors of gravitational wave detectors and models of their thermal noise

The mirrors of interferometric gravitational wave detectors (GWDs), often referred to as ‘test masses’, consist of large cylindrical masses, coated with multi-layer coatings to provide high reflectivity. Each arm cavity of a GWD has two such mirrors suspended at either end – these are known as the input test mass (ITM) and end test mass (ETM).

In this chapter these mirrors and their highly reflective (HR) coatings shall be discussed. The coating design and ethos will be briefly summarised, and an extensive description of the major detector noise source associated with the coating, the so-called ‘coating thermal noise’, shall be presented. Finally, a brief list of replacement candidate materials for future HR coating designs and their potential benefits will be discussed.




Figure 2.1: Images taken by the author, during an optical microscope survey, of the two ETMs previously installed at LIGO Livingston for aLIGO’s first to third observing runs.

2.1 The mirror design

Interferometric gravitational wave detectors require highly reflective mirrors ($R > 99.999\%$) at the laser wavelength used. This is achieved through coating their front surfaces with different high-quality materials. Apart from the coated surface, the test masses can generally be considered to be homogeneous glass/crystalline cylinders of the highest purity, polished to extremely high precision, only typically varying by, at most, a few nanometres across their faces [80] - images of two LIGO end test masses are shown in Figure 2.1.

Metallic coatings are not suitable for GWD mirrors due to the high levels of optical absorption they exhibit, which could potentially lead to too high levels of thermal deformation in the test mass, and damage to the coating. Instead, dielectric coatings, which typically have much lower absorption, and better durability are used [112].

HR coatings are typically produced by depositing alternating layers of two materials of different refractive index. When light is normally incident on the boundary between two media of different refractive indices n_1 and n_2 the fraction of light power reflected R is

$$R = \left(\frac{n_1 - n_2}{n_1 + n_2} \right)^2. \quad (2.1)$$

Dielectric HR coatings often incorporate multiple quarter wave layer (QWL) optical thickness layers to ensure the optical path length through the coating is such that the light reflected from every interface is all in phase and interferes constructively. An illustration of a QWL HR coating is shown in Figure 2.2. The physical layer thickness d_{layer} corresponding to a given optical thickness δ_{layer} is given by

$$d_{\text{layer}} = \frac{\delta_{\text{layer}}}{n} \quad (2.2)$$

$$\therefore d_{\text{QWL}} = \frac{\lambda}{4n}. \quad (2.3)$$




Figure 2.2: A schematic diagram of a multi-layer reflective coating stack, with alternating layers of a high refractive index (n_H) and a low refractive index (n_L) material.

For such a coating stack containing a number N of alternating high refractive index n_H and low refractive index n_L layers, deposited on top of a substrate with refractive index n_S , and situated in a medium of refractive index n_0 (for GWDs $n_0 \approx 1$) we can define a variable f such that [113]

$$f = \left(\frac{n_H}{n_L} \right)^{2N}, \quad (2.4)$$

and the total reflectivity of a stack may be described by [113]

$$R_{2N} = \left(\frac{n_S f - n_0}{n_S f + n_0} \right)^2 \quad (2.5)$$

$$R_{2N+1} = \left(\frac{n_H^2 f - n_0 n_S}{n_H^2 f + n_0 n_S} \right)^2. \quad (2.6)$$

Here Equation 2.5 represents the reflectivity for a stack comprised of an even number of total layers ($2N$), and Equation 2.6 for when the total number of layers is odd ($2N+1$). The reflectivity of the coating can therefore be increased through either increasing the number of bi-layers of high- and low-index material present, or through increasing the difference between the refractive indices of the layers. It follows, for a given reflectivity that the number of layers, and hence the total physical thickness of the coating, can be reduced by using materials with a larger difference in refractive index - this will be seen later, first in Equation 2.31, to be an important consideration for the level of coating thermal noise.

Current GWD coatings use silica (SiO_2) as the n_L material and tantala doped with 27% titania by cation concentration (27% Ti, $\text{TiO}_2:\text{Ta}_2\text{O}_5$) as the n_H material, and the interferometers operate with a primary laser wavelength of 1064 nm [82]. However, the HR stacks are not simple QWL topologies and, instead, contain slightly thicker than QWLs of SiO_2 and slightly thinner than QWLs of $\text{TiO}_2:\text{Ta}_2\text{O}_5$, with each layer pair still summing close to a $\lambda/2$ optical thickness [82, 114, 115]. This is done primarily to provide additional suitable coating reflectivity for the arm length stabilisation system lasers of current detectors which operate with auxiliary 532 nm laser light [114], though it also helps slightly improve the coating thermal noise.

2.1.1 Coating production

The multi-layer dielectric coatings of current GWDs are deposited via ion-beam sputtering (IBS) in the ‘Grand Coater’ system [116] at the Laboratoire des Matériaux Avancés [117]. IBS involves generating ions from an ion source (typically argon), the ions are accelerated to high energies, focused into a beam, and directed at a high-purity target material. On impact, the ions will eject atoms from the surface of the target, causing them to disperse throughout the chamber. An optic to be coated is held inside the chamber, and some of the ejected atoms will condense onto its surface. With precisely controlled movement of the optic, eventually a uniform coating layer will be formed – a simple schematic of an IBS system can be found in Figure 2.3.




Figure 2.3: A simple schematic diagram of an ion-beam sputtering (IBS) chamber.

IBS is a well established technique which has many advantages over other thin film deposition methods. Studies have shown that IBS coatings tend to produce significantly lower optical absorption than coatings deposited via other techniques such as electron-beam deposition [118]. It is also one of only a few methods that can produce the high uniformity required for low optical scatter [119]. IBS also tends to produce films of denser structure, which adhere more strongly to the substrates on which they are deposited, as the sputtering ions possess relatively higher energy (several eV) than in other methods [120].

The novel coatings presented in this thesis were all deposited via IBS, using argon ions as the sputtering species.

2.2 Optical loss requirements of GWD mirrors

Optical losses from the scattering and absorption of light of the primary mirrors of a GWD are two factors which need to be exceptionally low compared to the requirements of most other optical systems, in order for it to perform as intended. Scattering is the deflection of light from its intended path caused primarily by the irregularities of the mirror surface. It directly reduces the power circulating in the Fabry-Pérot cavities, leading to a reduced signal-to-noise ratio [121]. Scattered light can also couple back into the measured signal after reflection from non-seismically isolated components, imparting phase noise [121]. Optical absorption is an important aspect connected with the HR coating properties where some of the light power is converted into heat in the test mass which causes it to deform - this effect becomes even more detrimental when operating at cryogenic temperatures. It also induces a thermal lens in the mirror, which causes distortion to wavefronts passing through the optic, leading to issues in controlling the cavity length [122]. Ultimately, both of these effects limit the amount of laser power that can be stored in the interferometer and, thus, the sensitivity of the detectors at frequencies beyond a few hundred Hz where they are shot noise limited [122] - see Equation 1.4.

The current aLIGO/adV detector main mirrors have scatter of ~ 9.5 ppm [123], and absorption of ~ 0.22 ppm for the ITM coatings, and ~ 0.27 ppm for the ETMs at the main laser wavelength of 1064 nm [124], with total arm cavity loss in aLIGO of ~ 45 ppm [123]. For next generation upgrades to the detectors, scatter must be maintained below 10 ppm and absorption below 0.5 ppm in order for the detectors to meet their sensitivity goals [124]. Both effects shall be discussed and explored in greater detail in Chapter 4.

2.3 Coating thermal noise

As discussed in Chapter 1, coating thermal noise (CTN) is a major limit to the sensitivity of current detectors. This section will discuss its origin as well as historical and current models applied to gravitational wave detectors. Finding solutions for improving it is one of the primary goals for the planned upgrades for the fifth observing run of aLIGO+/adV+ [90, 125], and indeed it is the main goal of this thesis.

2.3.1 Brownian motion

Observations of thermally induced motion in a mechanical system were first reported by the Scottish botanist Robert Brown in 1828 [75]. Brown observed irregular ‘jittering’ motion of pollen grains and dust suspended in water, and despite many subsequent investigations by a number of scientists, the origins of this so-called ‘Brownian motion’ remained unclear for several decades. Sixty years after Brown’s observation, Gouy [126], followed by Exner [127], reported that the motion of the particulate would slow in more viscous liquids, and that it would increase when the temperature of the liquid was raised. However, a comprehensive mathematical description of Brownian motion did not emerge until 1905 when Einstein presented his seminal paper on the topic [128]. Einstein inferred that the motion stemmed from fluctuations in the collision rate of the pollen grains with water molecules, and that through these collisions the pollen grains were constantly gaining and losing kinetic energy. Einstein’s treatment of this observed motion served as an indirect verification of the existence of atoms and molecules. By treating the fluid as an ensemble of imperceptible atoms, he was able to deduce the cumulative impact the fluid molecules would have on the pollen grains over time, allowing for further calculation from

the amount of observed movement of the pollen, the mass of the particles in the fluid, and hence the size of the atoms that comprised it. This was also the first time that stochastic fluctuations, in this case the random displacement of the pollen and dust grains, had been linked to a dissipative process: the loss of energy due to the fluid viscosity.

2.3.2 Fluctuation-Dissipation theorem

Einstein's interpretation of Brownian motion showed a relationship between fluctuations and dissipation. More generally, all the constituent parts which comprise a linear system in thermal equilibrium will undergo spontaneous thermally driven fluctuations, the magnitudes of which are linked to the dissipative (i.e. real) part of the admittance of the system. This was later formalised by Callen et al. as the so-called Fluctuation-Dissipation theorem [129, 130, 131], providing a general relationship for excitation (fluctuation) in any linear system in thermal equilibrium, and the subsequent damping (dissipation). The theorem states that the power spectral density $S_x(f)$ of the fluctuating mechanical driving force in a mechanical system is related to the dissipative part of the mechanical admittance of the system $\Re[Y(f)]$ via:

$$S_x(f) = \frac{k_B T}{\pi^2 f^2} \Re[Y(f)] \quad (2.7)$$

where k_B is the Boltzmann constant, T is the temperature of the system, and f is the frequency. The Fluctuation-Dissipation theorem enables predictions of the thermal noise level present in a given measurement of power/amplitude spectral density of a system. For a real multi-mode system such as the suspended mirrors of gravitational wave detectors, further factors must be considered which will be discussed in the following sections.

2.3.2.1 Sources of internal dissipation

The Fluctuation-Dissipation theorem shows that the magnitude of the thermal noise in a system is directly proportional to the real (dissipative) part of the systems mechanical admittance. The latter is intrinsically related to the system's damping coefficient, and so the thermal noise level is also related to the level of damping present. When all external

sources of damping are suitably reduced, it follows that the damping is dominated by internal sources. For the following we shall assume a system where all external damping sources are repressed to a level where they are insignificant when compared to the internal damping/dissipation.

Internal damping arises when a material/object/system responds anelastically to a force being applied. Hooke's Law tells us that when an ideal elastic material undergoes stress σ as the result of an applied force, it instantly experiences a corresponding strain ϵ , represented mathematically as [132]

$$\sigma = \epsilon Y, \quad (2.8)$$

where Y is the Young's Modulus (otherwise known as the elasticity modulus). However, most real materials are 'anelastic', meaning they experience a lag in their strain response to an applied stress, developing over a finite relaxation time. When a periodic stress is applied to a material one may write this as [132]

$$\sigma(t) = \sigma_0 e^{i\omega t} \quad (2.9)$$

where t is time, ω is the angular frequency of the oscillation ($\omega = 2\pi f$), and σ_0 is defined as the stress amplitude. From Equation 2.8 it is clear stress and strain have a linear relationship, and thus, the resulting strain for an anelastic material is given by [76, 132]

$$\epsilon(t) = \epsilon_0 e^{i\omega t - \phi}. \quad (2.10)$$

Here ϕ is the phase angle by which the strain lags behind the stress, otherwise known as the loss angle. This loss angle can also be described as the internal friction/damping factor of the material, and quantifies the fractional energy lost (dissipated) during each cycle of oscillation [132] where we can say for a given initial energy stored in the system U , decaying by U_{diss} each oscillation

$$\phi = \frac{1}{2\pi} \frac{U_{\text{diss}}}{U}. \quad (2.11)$$

To highlight its dependence on frequency we can also further define it in terms of power dissipated in one cycle of oscillation P_{diss} as

$$\phi = \frac{1}{2\pi f} \frac{P_{\text{diss}}}{U}. \quad (2.12)$$

Applied to mechanical systems it is often common to refer to ϕ simply as the mechanical loss, which shall be the convention adopted hereafter.

Another physical parameter, known as the mechanical quality factor, Q , is a dimensionless quantity related to the dissipation of energy in an oscillator when a resonant frequency, f_0 is excited

$$Q(f_0) = 2\pi \frac{U}{U_{\text{diss}}} \equiv \frac{f_0}{\Delta f}, \quad (2.13)$$

which can be clearly seen to be the inverse of Equation 2.11. Thus it can be seen that when a resonant mode of a system is excited, the mechanical loss is equal to the inverse of the material's quality factor. As such the mechanical loss of different materials can be characterised by making samples comprising them resonate and recording the decay (dissipation) of the excited resonance peaks. In doing so for multiple resonant modes of the system, one can then extract mechanical loss trends with frequency and infer the loss level in the frequency band of interest for a GWD. From Equation 2.13 one can also see that Q is inversely proportional to Δf , which is the width of the resonant peak in frequency space, measured at half its maximum power (FWHM). From this, one can understand that low loss/high Q materials under oscillation will have narrow resonant peaks, and as such most of the energy will be stored at, and very close to, these resonant frequencies.

2.3.2.2 Harmonic oscillators with internal damping

The internal damping manifests as a phase lag ϕ between the response of a system x to an applied force F . This phase lag between applied stress and resulting strain can be incorporated into Hooke's law, by modelling the system as a harmonic oscillator, with a modified spring constant as follows [76]

$$F_{\text{spring}}(f) = -k(1 + i\phi(f))x. \quad (2.14)$$

One can treat the suspended mirrors in a GWD as an anelastic system, and so when their resonant modes are excited they can be treated as harmonic oscillators with internal damping. The corresponding equation of motion which describes this is [57, 76]

$$F(f) = m\ddot{x} + k(1 + i\phi(f))x, \quad (2.15)$$

where $F(f)$ is the thermal driving force on the oscillator of mass m . Since x can be described as $x \propto e^{i\omega t}$, where $\omega = 2\pi f$, we can say $\dot{x} = i\omega x$, and $\ddot{x} = i\omega\dot{x}$, thus this driving force may be rewritten as

$$\begin{aligned} F(\omega) &= i\omega\dot{x}m - k(1 + i\phi(\omega))\frac{i}{\omega}\dot{x} \\ F(\omega) &= \frac{\dot{x}}{\omega} [k\phi(\omega) + i(\omega^2m - k)], \end{aligned} \quad (2.16)$$

and using $k = \omega_0^2m$, where ω_0 is the thermally excited (angular) resonant frequency, we can further say

$$F(\omega) = \frac{\dot{x}m}{\omega} [\omega_0^2\phi(\omega) + i(\omega^2 - \omega_0^2)]. \quad (2.17)$$

This fluctuating force is related to the mechanical admittance of the system $Y(\omega)$, which is the inverse of the impedance and is a measure of how readily a force exerted at a point on the object (at a certain frequency) induces motion at the same frequency, given by [133]

$$Y(\omega) = \frac{\dot{x}}{F(\omega)} \quad (2.18)$$

$$= \frac{\omega}{m} \frac{1}{\omega_0^2\phi(\omega) + i(\omega^2 - \omega_0^2)}. \quad (2.19)$$

By multiplying both the numerator and denominator of Equation 2.19 by the complex conjugate of the denominator, we obtain

$$Y(\omega) = \frac{\omega}{m} \frac{\omega_0^2\phi(\omega) - i(\omega^2 - \omega_0^2)}{(\omega_0^2\phi(\omega))^2 + (\omega^2 - \omega_0^2)^2}, \quad (2.20)$$

the real part of which is

$$\Re[Y(\omega)] = \frac{\omega}{m} \frac{\omega_0^2 \phi(\omega)}{(\omega_0^2 \phi(\omega))^2 + (\omega^2 - \omega_0^2)^2}, \quad (2.21)$$

and written in terms of the frequencies is

$$\Re[Y(f)] = \frac{f}{2\pi m} \frac{f_0^2 \phi(f)}{(f_0^2 \phi(f))^2 + (f^2 - f_0^2)^2}. \quad (2.22)$$

Invoking the Fluctuation-Dissipation theorem, $\Re[Y(f)]$ can be substituted into Equation 2.7 to give an expression for the power spectral density of the displacement thermal noise, of the oscillating mirror in terms of its mass and mechanical loss

$$S_x(f) = \frac{k_B T}{2\pi^3 f m} \frac{f_0^2 \phi(f)}{(f_0^2 \phi(f))^2 + (f^2 - f_0^2)^2}. \quad (2.23)$$

2.3.2.3 Thermal noise associated with a single resonant mode

Dynamic systems such as the suspended mirrors of GWDs have many resonant modes which can be thermally excited. These include the internal modes intrinsic to the mirror, along with the pendulum modes of the suspended mirror and the bending modes of the suspension fibres (more commonly referred to as the fibre “violin” modes). The latter are the only modes typically present in the frequency band of current GWDs (see Figure 1.4), whereas the internal modes of the mirror lie outside the detection region at several tens of kilohertz where the detectors are quantum-shot noise limited, and the pendulum modes at low frequency where they are seismic noise limited. As such, it is off-resonance thermal noise which predominantly defines the sensitivity of a GWD within its detection band [57].

Considering Equation 2.23 in the three different frequency regimes where $f \ll f_0$, $f = f_0$, and $f \gg f_0$, and that for a gravitational wave detector low mechanical loss materials are used (i.e. $\phi^2(f) \ll 1$), we can see the equation simplifies in each case such that

$$S_x(f) \propto \begin{cases} \frac{\phi(f)}{f} & \text{if } f \ll f_0 \\ \frac{1}{\phi(f)f} & \text{if } f = f_0 \\ \frac{\phi(f)}{f^5} & \text{if } f \gg f_0. \end{cases} \quad (2.24)$$

So, at resonance, the magnitude of the thermal noise level increases with decreasing mechanical loss. However, away from this, the power spectral density of the thermal noise is directly proportional to the mechanical loss of the materials used. As such, within the gravitational wave detection band of current detectors, finding lower mechanical loss materials for the system components, such as the mirror coatings and suspension fibres, will lead to a reduction in the overall thermal noise. Decreasing the mechanical loss effectively confines a greater portion of thermal motion to a narrower frequency range centred around the resonance, leading to a narrower and elevated thermal noise peak at the resonant frequency, accompanied by a corresponding reduction in thermal noise level away from resonance. This is shown in Figure 2.4 where example amplitude spectral density thermal displacement noise (square-root of Equation 2.23) is plotted.




Figure 2.4: Thermal displacement noise spectra (amplitude spectral density) for three mechanical oscillators each of mass 40 kg, held at 300 K, with $f_0 = 10$ kHz, but with different mechanical losses (see graph). The lowest loss material exhibits lower off-resonance thermal noise, and higher on resonance thermal noise, as more of its thermal motion is concentrated close to the resonance.

2.3.3 Thermal noise associated with multiple resonant modes

The thermal noise model outlined in the previous section, culminating in Equation 2.23, is suitable for describing the thermal fluctuations arising from a single resonant mode of a mechanical system. In a gravitational wave detector, the laser beam reflected from a mirror's front face responds to the cumulative thermal displacement arising from numerous resonant modes of the suspended optic. Early models posited that the motion of all resonant modes were uncorrelated, and thus, the total amplitude/power spectral density of thermal noise could be computed simply by summing the contribution from each mode individually [134].

However, assuming uncorrelated motion from each mode implies that each mode is subject to an independent Langevin thermal driving force. Levin highlighted that this condition only holds true when the mechanical dissipation is uniformly distributed throughout the suspended mirror [135]. In reality, the losses of a suspended mirror in a gravitational wave detector are spatially inhomogeneous; the mirror coating's mechanical losses varies at each successive layer interface, and then again at the interface with the substrate; the welding points where the suspension fibres are affixed to the optic are also points of higher loss. Consequently, this leads to correlations in the fluctuations in the motion from different resonant modes, which simple summation cannot account for. As such an alternative approach was developed [135, 136, 137] to factor in this inhomogeneous distribution of loss in the system, while also accounting more precisely for the shape of the impinging laser beam and its influence on the thermal noise spectra.

Levin postulated the laser applying a notional oscillating pressure P to the front face of the mirror of the form [135, 137]

$$P = \frac{F_0}{\pi w_0^2} e^{-\frac{w^2}{w_0^2}} \cos(2\pi ft), \quad (2.25)$$

where F_0 is the peak amplitude of the force experienced, and the field amplitude radius of the impinging laser beam w_0 defines the radius w where its intensity drops by a factor of $\frac{1}{e}$ from its maximum. The real part of the mechanical admittance is then expressed in terms of the energy transferred into the suspended mirror via this oscillating pressure and

the subsequent dissipation via thermoelastic heat flow as [135]

$$\Re[Y(f)] = \frac{2W_{\text{diss}}}{F_0^2}. \quad (2.26)$$

Substituting this into the Fluctuation-Dissipation theorem (Equation 2.7), Levin concluded the power spectral density of the thermal noise of a suspended gravitational wave detector mirror can be expressed as

$$S_x(f) = \frac{2k_B T W_{\text{diss}}}{\pi^2 f^2 F_0^2}, \quad (2.27)$$

where W_{diss} is the time-averaged power dissipated in the test mass when this notional oscillating pressure is applied, which itself is given by [135]

$$W_{\text{diss}} = 2\pi f \int_V \varepsilon(x, y, z) \phi(x, y, z, f) dV. \quad (2.28)$$

Here ε is the energy density of the elastic deformation for the peak applied pressure F_0 . Notably, when the loss is considered spatially homogeneous W_{diss} can be further expressed in terms of the total energy associated with the peak elastic deformation U_{max} [135]

$$W_{\text{diss}} = 2\pi f U_{\text{max}} \phi(f). \quad (2.29)$$

2.3.4 Inhomogeneous loss applied to coating thermal noise

Through Equations 2.27–2.29 one can see that the magnitude of the thermal noise is correlated to the power dissipated in the suspended mirror which manifests from the notional oscillating pressure applied by the laser beam. The power dissipated at any point in the mirror is proportional to both the mechanical loss at that point and the elastic energy associated with the (pressure-resulting) deformation [135, 137] – see Equation 2.12. The deformation (and therefore the energy density) is greatest near the point where pressure is applied (i.e. the front surface), and consequently sources of deformation near

the front surface of the mirror contribute more to measured thermal noise of the system read at the output port of the interferometer. Therefore, to minimize thermal noise effects, if possible, more dissipative/higher loss sources should be positioned farther from the front surface of the test mass where the laser impinges.

This realisation is of great importance for interferometric gravitational wave detectors, as multi-layer dielectric coating stacks are deposited onto the mirror front surfaces to facilitate the high level of reflectivity necessary to detect gravitational wave signals. Unfortunately the current coating stacks incorporate multiple orders-of-magnitude higher loss materials than the bulk substrate materials used in both room temperature and cryogenic temperature detectors [82, 138], resulting in it being such a dominant noise source in the most sensitive frequency band of current detectors - see Figure 1.4. Finding novel solutions to this problem has therefore been a key area of research for the past few decades to further improve the detectors, and is ever more challenging the higher the detector sensitivity goals become.

2.3.5 Calculating coating thermal noise

Using Levin's method, Nakagawa et al. [139] approximated the thermal noise of a coated mirror. To do this, they approximated the multi-layer coating stack as one single composite layer of thickness equal to that of the total stack d , with identical elastic properties (Young's modulus Y , and Poisson ratio ν) as the mirror substrate. This leads to the following equation [139]:

$$S_x^{\text{total}}(f) = \frac{k_B T (1 - \nu^2)}{\pi^{\frac{3}{2}} f w_0 Y} \left(\phi_{\text{substrate}} + \frac{2 (1 - 2\nu) d}{\pi^{\frac{1}{2}} (1 - \nu) w_0} \phi_{\text{coating}} \right), \quad (2.30)$$

where w_0 is again the field amplitude radius. This has two distinct terms; a part pertaining to the thermal noise contribution from the substrate, and the other pertaining to that of the coating. The coating thermal noise specifically can thus be simplified and written as

$$S_x^{\text{CTN}}(f) = \frac{2k_B T d (1 + \nu)(1 - 2\nu)}{\pi^2 f w_0^2 Y} \phi_{\text{coating}}. \quad (2.31)$$

Harry et al. [140] then went on to incorporate inhomogeneity of the dissipation in the coating/substrate by considering the coating (denoted with subscript C) to have different elastic properties from the substrate (S). They also proposed a model where the dissipative response of the coating was different in the directions parallel and perpendicular to the coating surface, leading to two loss angles ϕ_{\parallel} with ϕ_{\perp} , and a corresponding total mirror displacement noise of

$$S_x^{\text{total}}(f) = \frac{2k_B T (1 - \nu_S^2)}{\pi^2 f w_0 Y_S} \left\{ \phi_S + \frac{1}{\pi^{1/2}} \frac{d}{w_0} \frac{1}{(1 - \nu_S^2)(1 - \nu_C^2) Y_S Y_C} \right. \\ \times \left[(1 + \nu_S)^2 (1 - 2\nu_S)^2 Y_C^2 \phi_{\parallel} \right. \\ \left. + (1 + \nu_S)(1 - 2\nu_S) \nu_C (1 + \nu_C) Y_S Y_C (\phi_{\parallel} - \phi_{\perp}) \right. \\ \left. \left. + (1 + \nu_C)^2 (1 - 2\nu_C)^2 Y_S^2 \phi_{\perp} \right] \right\}. \quad (2.32)$$

The second term representing just the coating thermal noise can be simplified and rewritten as

$$S_x^{\text{CTN}}(f) = \frac{2k_B T}{\pi^2 f w_0^2} \left\{ \left(\frac{(1 + \nu_S)(1 - 2\nu_S)}{Y_S} \right)^2 \frac{1}{(1 - \nu_C^2)} Y_C \phi_{\parallel} \right. \\ \left. + \frac{(1 + \nu_S)(1 - 2\nu_S)}{Y_S} \frac{\nu_C}{1 - \nu_C} (\phi_{\parallel} - \phi_{\perp}) \right. \\ \left. + \frac{(1 + \nu_C)(1 - 2\nu_C)}{Y_C} \frac{1}{(1 - \nu_C)} \phi_{\perp} \right\}. \quad (2.33)$$

In the case where $\phi_{\parallel} = \phi_{\perp}$, $Y_S = Y_C$, and $\nu_S = \nu_C$, Harry notes that Equation 2.33 agrees with the result of Nakagawa et al. (Equation 2.31). Harry et al. also noted that for fused silica substrates coated with alternating layers of amorphous tantala and silica (the coatings installed in detectors of the time, and still very similar to the current coatings), Equation 2.33 can be approximated to within $\sim 30\%$ by setting $\nu_S = \nu_C = 0$. This allows for a much more simplified version of the CTN equation to be defined as [140]

$$S_x^{\text{CTN}}(f) = \frac{2k_B T}{\pi^2 f w_0^2} \left(\frac{Y_C}{Y_S} \phi_{\parallel} + \frac{Y_S}{Y_C} \phi_{\perp} \right). \quad (2.34)$$

This simplified form of the power spectral density of the CTN is useful for estimating the expected level of thermal noise, and for also easily inferring the dependence of the CTN on certain parameters. One can see the CTN is a function of the mirror substrate Young's modulus and thus the same coating will produce a different thermal noise level when deposited onto substrates comprised of different materials. One also sees that the thicker the coating, the higher the CTN, so via Equations 2.4–2.6 the larger the refractive index contrast between the two materials of interest, the less coating layers are needed to achieve comparable reflectivity, and this will also directly reduce the CTN. Another notable feature is the CTN's inverse dependency on w_0^2 . One can also see from Equation 2.32, the substrate noise arising from the homogeneous dissipation in the mirror itself, varies simply as the inverse of w_0 . Therefore one potential solution to reduce thermal noise in general, and coating thermal noise in particular, is to increase the radius of the laser beams on the mirrors.

2.3.6 Bulk and shear loss dependent coating thermal noise

To enable coating thermal noise (CTN) to be calculated ϕ_{\parallel} was measured via the resonance ring-down method (explained extensively later in Section 3.4.1). However, through this technique ϕ_{\perp} cannot be measured, so it was concluded that the accuracy of Equation 2.33 is highly dependent on how well ϕ_{\perp} is known - Harry et al. often approximated it as $\phi_{\perp} = \phi_{\parallel}$ when performing their CTN calculations [140]. However, it would later be shown by Hong et al. [141] that if the assumption $\phi_{\perp} = \phi_{\parallel}$ is not used in this model, it can generate nonphysical, negative values for the dissipated energy.

Hong et al. [141] went on to develop their own method of quantifying thermal noise of a suspended GWD mirror, in which, rather than splitting its coating's dissipative response into motion parallel and perpendicular to the coating surface, it was assumed a coating material has two separate loss values, each associated with different types of deformation/motion. Specifically the dissipative response of the coating is split into motion inducing bulk deformation (with corresponding mechanical loss of ϕ_{κ}) and that inducing shear deformation (with mechanical loss ϕ_{μ}). For any given resonant mode of the system the total associated deformation can be decomposed into a sum of bulk (volumetric) deformation and shear (shape) deformation [142]. The total loss of a mode will consequently

depend on the proportion of elastic strain energy stored in bulk and in shear. For a resonating object with total energy lost given as U_T we may deconstruct this into the parts stored in bulk and shear motion as

$$U_T = U_K + U_\mu. \quad (2.35)$$

Hong et al. first modelled the CTN of a single layer coating in terms of the coating's U_K and U_μ - i.e. notably the effect of light penetration present in a multi-layer highly reflective coating was ignored. In the case where the thickness of the coating d is much less than both the total mirror thickness and the beam spot size, Hong showed the elastic deformation of the substrate is not influenced by the presence of the coating. As consequence one may write for the whole mirror [141]

$$\begin{aligned} W_{\text{diss}} &= 2\pi f (U_{\text{sub}}\phi_{\text{sub}} + U_K\phi_K + U_\mu\phi_\mu) \\ &= 2\pi f \left(\phi_{\text{sub}} + \frac{U_K}{U_{\text{sub}}}\phi_K + \frac{U_\mu}{U_{\text{sub}}}\phi_\mu \right), \end{aligned} \quad (2.36)$$

and for the part just pertaining to the coating (C) [141]

$$W_{\text{diss,C}} = 2\pi f (U_K\phi_K + U_\mu\phi_\mu), \quad (2.37)$$

where W_{diss} as before in Equations 2.25–2.29 is the time-averaged power dissipated in the test mass, and subscript ‘sub’ in Equation 2.36 pertains to just the substrate. Substituting Equation 2.37 into Equation 2.27 obtained previously by Levin, one obtains

$$S_x^{\text{CTN}}(f) = \frac{4k_B T}{\pi f} \left(\phi_K \frac{U_K}{F_0^2} + \phi_\mu \frac{U_\mu}{F_0^2} \right). \quad (2.38)$$

This CTN power spectral density in terms of bulk and shear coating loss angles can be further expanded to be represented in terms of just the elastic properties of the coating and substrate, and the coating ϕ_K and ϕ_μ . This is achieved by substituting into Equation 2.38 for the bulk term [141]

$$\begin{aligned} \frac{U_K}{F_0^2} = \frac{1}{3} \frac{d}{\pi w_0^2} & \left[\left(\frac{(1+\nu_S)(1-2\nu_S)}{Y_S} \right)^2 \frac{(1-2\nu_C)}{(1-\nu_C)^2} Y_C \right. \\ & + 2 \frac{(1+\nu_S)(1-2\nu_S)}{Y_S} \frac{(1+\nu_C)(1-2\nu_C)}{(1-\nu_C)^2} \\ & \left. + \frac{(1+\nu_C)^2(1-2\nu_C)}{(1-\nu_C)^2} \frac{1}{Y_C} \right], \end{aligned} \quad (2.39)$$

and for the shear term [141]

$$\begin{aligned} \frac{U_\mu}{F_0^2} = \frac{2}{3} \frac{d}{\pi w_0^2} & \left[\left(\frac{(1+\nu_S)(1-2\nu_S)}{Y_S} \right)^2 \frac{(1-\nu_C+\nu_C^2)}{(1+\nu_C)} \frac{1}{(1-\nu_C)^2} Y_C \right. \\ & - \frac{(1+\nu_S)(1-2\nu_S)}{Y_S} \frac{(1+\nu_C)(1-2\nu_C)}{(1-\nu_C)^2} \\ & \left. + \frac{(1+\nu_C)(1-2\nu_C)^2}{(1-\nu_C)^2} \frac{1}{Y_C} \right]. \end{aligned} \quad (2.40)$$

2.3.7 Accounting for layer penetration

While this interpretation by Hong et al. ignores light field penetration into the stack and assumes a composite single layer coating of average properties of the stack materials, they go on to produce a complete formulation for the thermal noise in multi-layer stack coatings accounting for the adapting light field present through the successive layers of the coating. This was later simplified by Yam et al. [143] under the condition that $\phi_K = \phi_\mu \equiv \phi$ and given by

$$S_x^{\text{CTN}}(f) = \frac{2k_B T}{\pi^2 f} \frac{1}{w_0^2} \sum_j b_j d_j \phi_j. \quad (2.41)$$

Here the subscript j indicates the j^{th} successive layer of the coating stack, with corresponding thickness d and mechanical loss ϕ for a given layer. The b_j term represents the (unitless) weighting factor for each coating layer and is described (after being corrected by Steinlechner and Martin in a subsequent publication [144]) by

$$b_j = \frac{(1 + \nu_j)(1 - 2\nu_j)}{(1 - \nu_j)} \left[\left(\frac{(1 + \nu_S)(1 - 2\nu_S)}{Y_S} \right)^2 \frac{1}{(1 + \nu_j)^2(1 - 2\nu_j)} Y_j + \left(1 - n_j \frac{d\theta_{\text{coating}}}{d\theta_j} \right)^2 \frac{1}{Y_j} \right], \quad (2.42)$$

where n_j is the refractive index of the j^{th} layer. The second term in the square brackets of Equation 2.42 describes CTN arising from fluctuations in the thickness of the coating. It is composed of two effects which act in opposite directions and thus partly compensate each other. The first effect represents when a thermally induced fluctuation in a layer alters its own optical thickness, and the second represents when this fluctuation in a given layer alters the position of the front surface of the mirror. The first effect results in fluctuations in the round-trip phase within each layer θ_j . The $\frac{d\theta_{\text{coating}}}{d\theta_j}$ term describes the sensitivity of the total coating phase, θ_{coating} , to a fluctuation in the j^{th} layer - it is explored at greater depth in [143]. The magnitude of $\frac{d\theta_{\text{coating}}}{d\theta_j}$ is proportional to the peak electric field intensity (EFI) in the j^{th} layer.

The EFI decreases with each successive double layer the laser interacts with, making the first effect smallest for layers furthest from the front face (i.e. decreasing towards the substrate). The second effect is independent of layer position, thus in the lower layers, this second effect is less compensated by the first effect. As a consequence, the lower down layers, closest to the substrate, contribute slightly more to b_j and hence the overall CTN.

2.3.8 Accounting for effective applied stresses and strains

In an isotropic medium, the elastic properties are the same for stresses applied in all directions. This ceases to be true in a multi-layer coating comprising of isotropic media, because the layering makes the direction perpendicular to the layers differ from the in-plane directions, which remain essentially isotropic.

In 2021, Fejer proposed an alternative approach to characterising a suspended coated GWD mirror [145], like Hong et al., incorporating differing bulk and shear coating losses, but with no layer penetration. His approach was to define an “effective-medium” for the coating where a local-variable averaging procedure is applied to compute an effective stiffness tensor which accurately approximates response of the whole multi-layered medium to applied stresses and strains (averaged over a period of the layer) [145, 146]. The essence of this approach is to arrange the components of the constitutive relations in forms that do not contain products of quantities discontinuous at the interface between layers. By applying this analysis to material properties measured from samples coated with just single layers of the individual materials which comprise the final multi-layer stack, one should be able to accurately estimate the effective properties relevant to CTN calculation of the HR stack.

Assuming a structure comprising of two alternating layers of materials A and B, with corresponding thicknesses d_A and d_B , we can define an averaging operator \bar{g} as [145, 147]

$$\bar{g} \equiv \langle g \rangle \equiv \frac{d_A}{d_A + d_B} g_A + \frac{d_B}{d_A + d_B} g_B, \quad (2.43)$$

where g_A/g_B represent any quantity or combination of quantities in layer A/B. Noting the form of the thermal noise previously given by Levin (Equation 2.27), Fejer goes onto define a W_{diss} of the coating in terms of the effective properties of the different materials which comprise it [145]

$$W_{\text{diss}} = \frac{d}{w_0^2} f F_0^2 \left\{ \left(\frac{(1 + \nu_S)(1 - 2\nu_S)}{Y_S} \right)^2 \left[\frac{1}{3} \left\langle \frac{(1 - 2\nu_j)}{(1 - \nu_j)^2} Y_j \phi_{K,j} \right\rangle + \frac{2}{3} \left\langle \frac{1 - \nu_j + \nu_j^2}{(1 - \nu_j)^2 (1 + \nu_j)} Y_j \phi_{\mu,j} \right\rangle \right] \right. \\ \left. + \frac{2(1 + \nu_S)(1 - 2\nu_S)}{3 Y_S} \left\langle \frac{(1 + \nu_j)(1 - 2\nu_j)}{3(1 - \nu_j)^2} (\phi_{K,j} - \phi_{\mu,j}) \right\rangle \right. \\ \left. + \left[\frac{1}{3} \left\langle \frac{(1 + \nu_j)^2 (1 - 2\nu_j)}{(1 - \nu_j)^2} \frac{1}{Y_j} \phi_{K,j} \right\rangle + \frac{2}{3} \left\langle \frac{(1 + \nu_j)(1 - 2\nu_j)^2}{(1 - \nu_j)^2} \frac{1}{Y_j} \phi_{\mu,j} \right\rangle \right] \right\}, \quad (2.44)$$

and substituting this into Equation 2.27 yields a power spectral density for the CTN of

$$\begin{aligned}
S_x^{\text{CTN}}(f) = & \frac{2k_B T}{\pi^2 f} \frac{d}{w_0^2} \times \\
& \left\{ \left(\frac{(1 + \nu_S)(1 - 2\nu_S)}{Y_S} \right)^2 \left[\frac{1}{3} \left\langle \frac{(1 - 2\nu_j)}{(1 - \nu_j)^2} Y_j \phi_{K,j} \right\rangle + \frac{2}{3} \left\langle \frac{1 - \nu_j + \nu_j^2}{(1 - \nu_j)^2 (1 + \nu_j)} Y_j \phi_{\mu,j} \right\rangle \right] \right. \\
& + \frac{2(1 + \nu_S)(1 - 2\nu_S)}{3 Y_S} \left\langle \frac{(1 + \nu_j)(1 - 2\nu_j)}{3(1 - \nu_j)^2} (\phi_{K,j} - \phi_{\mu,j}) \right\rangle \\
& \left. + \left[\frac{1}{3} \left\langle \frac{(1 + \nu_j)^2 (1 - 2\nu_j)}{(1 - \nu_j)^2} \frac{1}{Y_j} \phi_{K,j} \right\rangle + \frac{2}{3} \left\langle \frac{(1 + \nu_j)(1 - 2\nu_j)^2}{(1 - \nu_j)^2} \frac{1}{Y_j} \phi_{\mu,j} \right\rangle \right] \right\}. \quad (2.45)
\end{aligned}$$

2.3.9 How CTN calculations shall be performed in this thesis

In the simplified case of equal bulk and shear coating loss angles, and neglecting the effect of layer penetration into the stack (i.e. modelling a single layer coating), both Yam's (Equation 2.41) and Fejer's (Equation 2.45) interpretation of CTN are consistent with each other. Further to this, under the additional simplifications of $Y_C = Y_S$ and $\nu_C = \nu_S$ they also both yield the results of Harry (Equation 2.33).

No strong evidence was found for significantly different bulk and shear mechanical losses of the titania-silica coatings investigated in Chapter 3. Additionally, in Chapter 6, samples coated with stacks of more complex structure than of simple quarter wave optical thickness bi-layers were investigated. Specifically, some of the highly reflective coatings measured had progressively changing layer thicknesses. In order to draw direct comparisons between these various independent studies, Yam's approach, which accounts for the adapting light field present throughout the successive layers of a coating stack (Equation 2.41) was adopted for all following CTN analysis, except in instances where it was measured directly.

2.4 Some notable coating material candidates

Current GWDs rely on highly reflecting, low thermal noise coatings of their mirrors in order to operate at the sensitivity level necessary for gravitational wave detection. The next coating upgrades for aLIGO+/adV+ are required to maintain reflectivity $R > 99.999\%$, and whilst maintaining < 0.5 ppm optical absorption and < 10 ppm scatter [124] whilst

also reducing the coating thermal noise to half the level of the current coatings [90]. As the author began their PhD studies, no clear material candidate existed that could meet all these requirements, and as such the main goal of this project was to find and further develop suitable materials to be implemented in the next detector upgrades.

The ETMs of a GWD arm cavity require a greater number of layers deposited than the ITMs, as it is necessary for the ETMs to be as reflective as possible at the intended laser wavelength - as such it is the ETMs which have the higher mechanical and optical losses of the two, and it is their numbers which will be discussed here. Current GWD ETM coatings consist of 18 bi-layers of silica (SiO_2) and titania doped tantala ($\text{TiO}_2:\text{Ta}_2\text{O}_5$) deposited via ion-beam sputtering (see Section 2.1.1). These coatings produce absorption of ~ 0.27 ppm [124], scatter of ~ 9.5 ppm [123], and a coating thermal noise amplitude spectral density at 100 Hz (where the detector is most sensitive) of $\sim 6.6 \times 10^{-21} \frac{\text{m}}{\sqrt{\text{Hz}}}$ [148]. The coating thermal noise (CTN) is dominated by the mechanical loss of the high-index $\text{TiO}_2:\text{Ta}_2\text{O}_5$, with mechanical loss of $\sim 2.4 \times 10^{-4}$ at 100 Hz, which is an order of magnitude higher than the SiO_2 layer mechanical losses of $\sim 4\text{--}5 \times 10^{-5}$ [149]. As such, a focus in the global effort to produce lower CTN coatings for next-generation room temperature GWD upgrades is to find a suitable replacement material for $\text{TiO}_2:\text{Ta}_2\text{O}_5$, which has lower mechanical loss [148]. Some of the promising coating candidate materials for both room temperature and cryogenic GWDs will now be discussed.

2.4.1 Titania mixed with germania ($\text{TiO}_2:\text{GeO}_2$)

Recent works into modelling of amorphous materials indicated that materials exhibiting higher proportions of corner-sharing to edge-sharing polyhedra may yield lower mechanical losses than those which do not [150, 151]. As such, GeO_2 (along with SiO_2) was identified as a potential low loss amorphous oxide, and a strong potential future GWD coating material candidate. It was thought that investigating doping GeO_2 with TiO_2 would lead to multiple benefits. Previous studies showed that doping Ta_2O_5 with TiO_2 led to significantly improved mechanical losses versus pure Ta_2O_5 coatings of around 25% [115,

152]. TiO_2 also has a higher refractive index ($n_{\text{TiO}_2} \approx 2.3$ [82]) than GeO_2 ($n_{\text{GeO}_2} \approx 1.5$ [153]), so a mixture of the two would have increased refractive index than compared with pure- GeO_2 , thus, resulting in both fewer and thinner coating layers required for a full reflectivity stack, both of which are also beneficial for reducing the CTN [115, 148].

In 2021 Vajente et al. reported results from a promising coating mixture of $\text{TiO}_2:\text{GeO}_2$ with 44% TiO_2 by cation concentration ($n \approx 1.88$), which yielded promising optimum mechanical losses of $(0.96 \pm 0.18) \times 10^{-4}$ - i.e. around a 60% reduction versus the current $\text{TiO}_2:\text{Ta}_2\text{O}_5$ recipe [148]. Further, from these single layer coatings, a reduction in the CTN level of around 56% of that in current detectors was predicted [148], meeting the next generation upgrade requirements. However, the optical loss of this material at the time was higher than upgrade tolerances [148], and when stacks incorporating the material with the low-index partner SiO_2 were made, they exhibited significant defects, post heat-treatment, in the form of bubbles which would make them non-viable for use in GWDs [115].

It is still currently the favoured material to replace $\text{TiO}_2:\text{Ta}_2\text{O}_5$ for next generation upgrades [115], and research in the last few years has been in mitigating defect formation, and in lowering its optical losses. In Chapter 6, studies of IBS $\text{TiO}_2:\text{GeO}_2$ coatings, including both single layers and stacks with SiO_2 as a low-index partner, produced via a recipe designed to mitigate defect formation are presented. In particular, the mechanical losses of the first stack designs which meet the GWD ETM reflectivity requirement were characterised. One of these also featured a novel design to mitigate optical losses.

2.4.2 Titania mixed with silica ($\text{TiO}_2:\text{SiO}_2$)

For the same reasons $\text{TiO}_2:\text{GeO}_2$ is of interest as a future GWD mirror coating material, $\text{TiO}_2:\text{SiO}_2$ is also deemed a promising material candidate [150, 154]. SiO_2 is the amorphous oxide which appears to exhibit the lowest known room-temperature mechanical loss across the acoustic frequency range [82, 155, 156], and doping it with TiO_2 has potential to lead to a low mechanical loss, low optical loss, high refractive index partner mix material, to pair with pure- SiO_2 in an HR stack. Though as the refractive index of SiO_2 ($n_{\text{SiO}_2} \approx 1.44$ [82]), is lower than GeO_2 , potentially greater TiO_2 doping and/or thicker coating stacks of it would be needed to make it viable. In 2021, collaborators in the LIGO and Virgo

scientific collaborations launched investigations into both $\text{TiO}_2:\text{GeO}_2$ and $\text{TiO}_2:\text{SiO}_2$ [154, 157], though general research focuses shifted more towards the former material in their subsequent investigations. However, $\text{TiO}_2:\text{SiO}_2$ is a mix of materials known to produce low optical loss [158] so was deemed of great interest to study in conjunction with the ongoing developments in $\text{TiO}_2:\text{GeO}_2$. As such, this thesis provides the most comprehensive (and highly promising) investigations on mixes of this material as a GWD candidate to date, which are explored through Chapters 3, 4, and 5.

2.4.3 Amorphous silicon (aSi)

A promising high refractive index material for cryogenic detectors is amorphous silicon (aSi). It is the only considered candidate material comprising of a single atomic species, and as such the stoichiometry during deposition is not as likely to drift, eliminating a potential point of failure during production [79]. The refractive index of aSi at wavelengths of interest for GWDs (between 1–2 micron) is the highest of all materials considered, meaning it could lead to significantly thinner coatings, though this value greatly varies with deposition parameters. The refractive index of aSi ranges from around 2.7–4.0 in the literature [79, 159, 160]. Its mechanical losses are also very low at cryogenic temperatures, while at room temperature its losses are moderately higher but still around the levels of $\text{TiO}_2:\text{Ta}_2\text{O}_5$ in current detectors [161, 162]. An HR coating using aSi as a high refractive index material at cryogenics could therefore be much thinner, and exhibit lower mechanical loss compared to other materials, thus resulting in significantly reduced CTN.

One of the major challenges with amorphous silicon is that its optical absorption is too high at present for gravitational wave detection [115, 159]. Thus, work is being undertaken to reduce its absorption through adjusting factors such as deposition parameters [161], and subjecting it to hydrogenation [163]. Though perhaps one of the most promising avenues of research in aSi resides in incorporating it as part of multimaterial coating designs (i.e. more-than-two-material HR stacks), wherein burying it in the lowest layers of the stack, after most of the light is already reflected by layers of lower optical loss materials, its influence on the absorption drops significantly - this concept was experimentally verified in 2020, with the author contributing significantly to the optical absorption and mechanical loss characterisation effort [164].

2.4.4 Silicon nitride ($\text{SiN}_X \text{H}_Y$)

Silicon nitride is another promising coating material candidate with mechanical losses at around $\times 10^{-5}$ [165], both as a high refractive index material replacement to $\text{TiO}_2:\text{Ta}_2\text{O}_5$ in future room temperature and cryogenic detectors [115, 166], and also at cryogenics as a potential low index partner to aSi [167]. However, like $\text{TiO}_2:\text{GeO}_2$ and aSi, at present its absorption is still too high for implementation in future GWDs. To tackle this, research is currently ongoing to lower this material's absorption, and also to reduce the impact of its potentially higher absorption through incorporating it into multimaterial designs [168].

2.4.5 Aluminium gallium arsenide (AlGaAs)

Aluminium gallium arsenide (AlGaAs) is a crystalline material being considered for next generation gravitational wave detectors. Crystalline based HR mirrors incorporating AlGaAs as the high index material and GaAs as the low index material, produced via molecular beam epitaxy, have been shown to produce total mirror stack mechanical losses at the extremely low $\times 10^{-5}$ level or below [169]. However, these crystalline coatings must be grown on GaAs substrates, and then transferred onto the final substrates, which poses great challenges for large scale applications, such as for future GWD mirrors (which are also planned to greatly increase in size in the future [97]). Incorporating crystalline coatings such as AlGaAs/GaAs [169], which exhibit extremely low mechanical loss, could potentially reduce coating thermal noise to the level where the mechanical losses linked with the SiO_2 mirror substrates may then significantly contribute to the overall detector noise [170]. Current research challenges of this material are primarily in up-scaling the process from which they are currently grown, to produce homogeneous coatings of the diameter needed for current detector test mass sizes ~ 30 cm [115]. GaAs substrates of the necessary size to facilitate this do not currently exist.

A selection of the author's contributions to the global effort to find and develop future optical coating solutions with extremely low thermal noise, and optical loss, meeting the next upgrade requirements of GWDs are described throughout the following chapters.

Chapter 3

Studies of titania-silica mixtures as a low thermal noise coating material

3.1 Introduction

This chapter presents investigations into characterising the mechanical loss and the coating thermal noise (CTN) for ion beam sputtered titania-silica mix ($\text{TiO}_2\text{:SiO}_2$) based coatings for application in future gravitational wave detectors. Specifically, the mechanical losses of two different highly reflective (HR) coating stack designs incorporating 1064 nm quarter wave layers of both $\text{TiO}_2\text{:SiO}_2$ and SiO_2 were measured, with CTN values then calculated. Independent to this, for one of these stacks the CTN was also directly measured by collaborators at the Massachusetts Institute of Technology, and comparisons drawn between the results from the two techniques. Also presented are investigations aiming to characterise key properties of these mixed-material coatings, such as their cation concentration, refractive index, density, Young's modulus, and Poisson ratio, all of which are vital for extracting the mechanical loss and CTN.

The culmination of the studies presented here, as well as those in the following two chapters, resulted in a first author Physical Review Letters publication in 2023 [171].

3.2 Why titania-silica mixes?

As discussed in Chapters 1–2, coating thermal noise is one of the dominant noise sources in current gravitational wave detectors and ultimately limits their ability to observe weaker or more distant astronomical sources at their most sensitive frequencies [78]. Improving this CTN limit is an ever present challenge for detector designs in both the near and far future. The current coatings of the aLIGO and adV mirrors are amorphous dielectric

coatings, and incorporate alternating near quarter-wave layers of 27% titania, titania-doped-tantala and silica as the high and low refractive index materials respectively [82]. Doping the tantala with titania was shown previously to reduce its mechanical loss by around 25% [115], yet after this the mechanical loss of the titania-doped-tantala is still substantially higher than the silica [82, 172]. As such the mechanical loss of the full coating stack is dominated by the mechanical loss of the titania-doped-tantala, and finding a suitable alternative high refractive index material is important for next generation detector upgrades.

The magnitude of the CTN amplitude spectral density is directly proportional to the square root of the mechanical loss of the coating and its thickness. Therefore, in order to reach the required sensitivity, and full astronomical potential, of planned detector upgrades and future detectors, a reduction in thickness via the use of materials with a higher contrast in refractive index, n , or a reduction in coating mechanical loss is essential. Numerous promising replacement high- n material candidates and their application were discussed previously in Section 2.4. The ethos of trialling titania-silica-mixes as a high- n candidate is similar to that of titania-doped-germania: whilst having a lower refractive index than the current titania-doped-tantala (meaning thicker coating stacks would be required), the mechanical loss may prove low enough to reduce the CTN overall [115]. Indeed the interest in both titania-silica-mixes and titania-doped-germania arose as a result of structural modelling studies which predicted that materials with a high portion of corner-sharing polyhedra (such as these) should exhibit low loss at room temperature [150].

Titania mixed with silica is a particularly interesting composition of two materials known to have low optical absorption, which is another important requirement for future detector upgrades [115]. As such, developing and characterising its performance could yield potential solutions for improvements on current detector designs. In recent years, Vajente, Menoni and others, presented initial promising mechanical loss investigations of different dopings of this material alongside their much wider spanning titania-doped-germania investigations [157], which, in large part, inspired this project. In this thesis, the most extensive study of this material to date, as a GWD mirror coating candidate, is presented.

3.3 Titania-silica coating deposition and composition

For these studies FiveNine Optics [173] in Colorado, USA was commissioned to produce multiple single layer coatings of $\text{TiO}_2\text{:SiO}_2$ and two HR stack designs incorporating pure SiO_2 as a low-index partner to the $\text{TiO}_2\text{:SiO}_2$ high-index layers. These coatings were deposited onto fused SiO_2 substrates via ion beam sputtering (the same technique used for the current aLIGO/adV coatings) with Ti and SiO_2 targets and argon as the sputtering ion. To investigate the effects of the composition, four single layer coatings were deposited, of which three were different in composition, while that with the highest TiO_2 concentration was deposited twice, allowing consistency checks of the deposition process.

The cation concentration of the single layers was nominally targeted at 40%, 50% and, 60% Ti, chosen due to the previous works by Vajente and Menoni [157], and aiming to target the seemingly most interesting region of the parameter space. Coated samples were sent to collaborators at the University of Montreal, who used Rutherford backscattering spectrometry (RBS) to determine the elemental composition of the single layers [174] and elastic recoil detection with time-of-flight measurements (ERD-TOF) to determine the composition of the HR stack layers. The author performed key analysis towards the latter measurements. The same samples were also sent to a collaborator at Hamburg University to record the transmission spectrum and measure the refractive index of each single layer coating.

From transmission spectra of the single layers, measured using an Agilent Cary 5000 spectrophotometer [175], refractive index n and film thickness t were obtained using the software tool SCOUT [176]. Table 3.1 shows the results for the as-deposited single layers and for selected annealing temperatures, showing the expected increase in n with increasing Ti cation concentration, as well as slight variations in the refractive index with heat treatment. For the HR stacks, the refractive indices of the high- n layers were estimated via a linear interpolation with concentration of the single layer n results. With the thickness and refractive indices known, the samples were shipped to colleagues at Montreal to measure the composition.

Table 3.1: Nominal and measured Ti-cation concentrations, from single layers of coating thickness t and density ρ . Refractive index n at 1064 nm at selected heat-treatment temperatures T and as deposited (as dep.), with corresponding stack thickness t_{HR} required for $R \geq 99.9994\%$ (using $n_{\text{SiO}_2} = 1.45$) are also shown. Ti-cation concentrations results are also shown for the stack samples that were produced.

| Ti/(Ti+Si) [%] | | t [nm] | T [°C] | n | t_{HR} [μm] | ρ [kg/m^3] |
|----------------|----------------|--------|---------|------|-----------------------------------|-----------------------------------|
| target | measured | | | | | |
| single layers | | | | | | |
| 40 | 46.2 ± 0.4 | 298 | as dep. | 1.77 | 10.83 | 2686 ± 50 |
| | | | 450 | 1.75 | 11.56 | |
| 50 | 58.5 ± 0.5 | 274 | as dep. | 1.88 | 7.94 | 2889 ± 50 |
| | | | 450 | 1.86 | 8.63 | |
| | | | 750 | 1.88 | 7.94 | |
| 60 | 62.3 ± 0.5 | 272 | as dep. | 1.93 | 7.21 | 2967 ± 50 |
| | | | 450 | 1.92 | 7.54 | |
| | | | 500 | 1.91 | 7.56 | |
| 60 | 62.6 ± 0.5 | 256 | as dep. | 1.93 | 7.21 | 3023 ± 50 |
| | | | 450 | 1.91 | 7.56 | |
| HR stacks | | | | | | |
| 50 | 63.2 ± 1.7 | | as dep. | 1.93 | 7.21 | |
| 60 | 69.5 ± 1.3 | | as dep. | 1.97 | 6.83 | |

RBS is a technique which uses high-energy ions to bombard the surface of a specimen. When these ions collide head-on with atoms in the coating, a small portion are backscattered at high angles (close to 180°) as a result of interactions with heavy nuclei. Through measuring the number and energy of the backscattered ions, RBS can identify the elements present in the coating, determine their relative concentrations, and the depth profile of particular nuclei, and can achieve this intrinsically for a material without need of any reference standard. The measurements were conducted with a 2032 keV He beam on a Tandetron accelerator [177], with an incidence angle of 7° , and the detector placed at a scattering angle of 170° . Simulations with the ion beam analysis software SIMNRA [178], using a uniform and stoichiometric layer model, were used to estimate the measured cation ratios Ti/(Ti+Si)- again see Table 3.1. The two coatings with nominally identical target composition were found to be the same within measurement uncertainty, showing the consistency of the process. However the actual Ti contents were also found to differ significantly from the target Ti concentrations.

Impurities heavier than Si in all of the single layer coatings were also measured. All as-deposited single layers were observed to contain around 0.5% argon - this was the largest detected impurity and arises from the argon sputtering beam. 200 – 300 ppm of molybdenum was also found in each coating, likely originating from sputtering from the acceleration grid in the deposition chamber [179], and under 50 ppm of lanthanide, possibly neodymium was also detected. From these single layer RBS results, the uncertainty on the Si and Ti concentration was around 0.1%. The Ti cation concentration errors in these were calculated by adding the numerator and denominator fractional uncertainties in quadrature, yielding cation concentration errors of 0.4 – 0.5%. RBS also allows the areal density of the scattering atoms to be determined, which combined with the measured film thickness was used to find the coating density ρ . Indeed the results obtained for the density align broadly with mixtures of TiO_2 and SiO_2 , which one would expect for relatively even concentration splits to be around $\sim 3000 \text{ kg/m}^3$, with increasing density with Ti content, which is what was observed.

Two HR multi-layer stacks were also deposited, using SiO_2 as the low- n partner to the high- n TiO_2 : SiO_2 layers – the first (denoted HR-1) had a target high- n Ti concentration of 50%, with 17.5 layer pairs (starting and ending with the high- n material), and the second stack (HR-2) had a target high- n Ti content of 60%, with 21.5 layer pairs. For the first attempt FiveNine did not want to deposit too many layer pairs in order to mitigate potential problems arising in the deposition process - once they had experience with the materials, they were happy to go further and make a thicker stack with full ETM reflectivity such as HR-2. The stacks were designed for 1064 nm wavelength and each layer was nominally $\lambda/4n$ in thickness to give high reflectivity. The Ti cation concentrations were verified using ERD-TOF analysis to probe the elemental composition of the top four layers of each stack. The main advantage of ERD-TOF over RBS is that it provides a good sensitivity for light elements, such as hydrogen, carbon, nitrogen and oxygen, even in matrices containing heavy elements [180, 181]. Since, in ERD-TOF, both the velocity and the energy for each detected atom are measured, it enables each individual element from the spectra to be distinguished. This eliminates issues of overlapping of the different elements in RBS, and allows for more complex stacked structures of different materials to be better characterised. The depth of probe however, is limited compared to RBS [181].




Figure 3.1: ERD-TOF atomic content data as a function of depth (expressed in atoms per square centimetre) for the as-deposited HR stacks. HR-1 is plotted on the left and HR-2 on the right. The argon and carbon is multiplied by a factor of 10.

Chicoine et al. [181] provide details of the specific ERD-TOF setup used. Collaborators in Montreal conducted the measurements with the author then analysing the data and extracting the cation ratio $\text{Ti}/(\text{Ti}+\text{Si})$ of the mixed material layers.

Figure 3.1 shows the measured elemental compositions of the first four layers of both HR stacks. Note that the most abundant element, oxygen, has been excluded from the plots, and the carbon and argon impurity spectra are multiplied by $\times 10$ for visibility. The plateau regions of the Ti and Si spectra indicate the four layers probed, with a $\text{TiO}_2:\text{SiO}_2$ layer at the top of the stack $\sim(0-90) \times 10^{16} \text{ at./cm}^2$, then a SiO_2 layer $\sim(100-210) \times 10^{16} \text{ at./cm}^2$, then another $\text{TiO}_2:\text{SiO}_2$ layer $\sim(220-300) \times 10^{16} \text{ at./cm}^2$, then another SiO_2 layer beyond $\sim 310 \times 10^{16} \text{ at./cm}^2$. Note that the trapezoid shape arises from imperfect depth resolution of the measurement [181]. In reality the layer boundaries are expected to effectively be completely sharp. It can be seen that HR-1 has a lower Ti content in both its mixed material layers than HR-2, and that the two observed mixed material layers of each HR stack have generally consistent Ti contents to within a few percent. As with the single layers, the argon content in each HR coating is around 0.5% as-deposited, but note the oscillating trend with slightly more argon in the mixed layers, and the carbon content that is approximately 2–3 times larger than this. This is true for both the mixed material and ‘pure’ SiO_2 layers. The slightly higher carbon content on the top (lowest depth) layer of HR-2 compared with the deeper layers likely arises from surface contaminants resting atop the optic during measurement, rather than indicating a real structural trend.




Figure 3.2: ERD-TOF silicon and titanium cation concentrations for the as-deposited HR stacks as a function of depth (expressed in atoms per square centimetre). HR-1 is plotted on the left and HR-2 on the right.

The Ti and Si cation concentrations of the two stacks are plotted in Figure 3.2. Layer regions were identified as the ‘flat’ parts of these spectra (i.e. in-between where the point-to-point gradients cross through zero), with values for the high- n layers being extracted from the two plateau regions of high Ti content. For HR-1 this corresponded to depth regions of $(24-54)$, and $(240-273) \times 10^{16}$ at./cm². For HR-2 this corresponded to depth regions of $(27-48)$, and $(252-282) \times 10^{16}$ at./cm². HR-2 having slightly thinner layers than HR-1 is consistent with the trend of increasing refractive index with Ti content, and hence reduction in required quarter wavelength thickness. With the Ti and Si atomic concentrations in the plateau regions defined, the Ti cation content could then be calculated, and the mean and standard deviation of the measurements acquired. This yielded a high- n layer Ti cation concentration of $(63.2 \pm 1.7)\%$ for HR-1 and $(69.5 \pm 1.3)\%$ for HR-2. The measurement uncertainty in these cation concentrations is known to be around 1% for this system, and the observed spread of the data here in each case is larger than this, and is what is represented in the stated uncertainty - highlighting as well the slight, but, measurable variability in cation content of the successive layers in each stack.

The cation concentration and thickness values obtained for each of the single layer and HR stack coatings are essential for extracting accurate mechanical loss values from the coated substrates. From the transmission spectra, RBS and ERD-TOF studies, all these vital parameters were extracted, with each being used in subsequent coating material analysis as described in Section 3.4.3.

3.4 Titania-silica mechanical loss studies

Investigations of the mechanical losses of the $\text{TiO}_2\text{:SiO}_2$ coatings described were carried out in collaboration with a fellow PhD student at Maastricht University, Viola Spagnuolo. For efficiency, some samples were measured in both labs, with analysis being carried out jointly. At least 50% of the measurements presented in this chapter were carried out in Glasgow by the author, over 90% of the four HR-stack measurements were performed in Glasgow, with the single layer disk mostly measured at Maastricht, but regularly shipped to Glasgow for post-deposition heat treatment and system crosschecks. It is well known that heat treatment can significantly reduce the mechanical loss of bulk fused SiO_2 [182, 183], and so all uncoated samples in this study were heat treated at 950°C for 5-10 hrs prior to any measurement, and after they were characterised in this uncoated (otherwise known as “blank”) state they were coated.

Low mechanical loss materials have narrow mechanical resonant peaks in frequency-space and therefore when resonances are excited by thermal energy, the majority of the resulting motion will be concentrated close to the resonant frequencies. Mechanical loss can be calculated from the time dependent decay of the displacement amplitude of a given excited sample - i.e. the ringdown of its vibrational motion after being excited at resonance. The mechanical loss ϕ of a sample vibrating at resonant frequency f_R is given by [57, 165]

$$\phi \approx \frac{1}{Q} = \frac{1}{\pi f_R \tau}, \quad (3.1)$$

where Q is the quality factor of a mode of the sample, and τ is the characteristic ringdown time of the decay in the oscillation brought about by the internal dissipation of energy. Specifically τ represents the time taken for such an oscillation to decay to $\frac{1}{e}$ of its starting level. For a complex resonator such as a glass disk on which a single or multilayer coating is deposited, this measured loss will depend on both the mechanical losses of the coating and substrate materials [140], and in real measurements some external damping terms that will ideally be minimised by the measurement setup used. Comparing the measured loss of the same sample before and after coating deposition allows for extraction of the

coating loss [140]. In order to extract this coating loss at a particular frequency, the strain energies stored in the coating and substrate for that particular mode shape must also be calculated. Essentially a five step process must take place for extracting coating mechanical losses from a sample:

- The ringdown data for a given excited sample is acquired pre and post coating.
- The uncoated and coated sample mechanical losses are extracted from the ringdown data via a fitting algorithm.
- The strain energy distributions between coating and substrate are calculated for each resonant frequency.
- The coating mechanical loss is then extracted from the previous steps.
- In performing these steps for multiple modes of resonance, and following multiple rounds of post-deposition heat treatment, the evolution of coating loss as a function of frequency, and as a function of heat treatment temperature are acquired.

3.4.1 Measuring loss via gentle nodal suspension

The mechanical losses of a sample can be extracted via exciting its resonant frequencies, using a Gentle Nodal Suspension (GeNS) to hold the sample whilst minimising external sources of damping. In 2009 Cesarini et al. [184] demonstrated the first GeNS system and its ability to extract precisely the mechanical losses of thin cylinder (disk) specimens. In a GeNS system, disk samples are balanced on the top of a sphere, which touches one of the nodal points of vibration. The mechanical modes of the disk are typically excited using electrostatic comb actuators placed close to (but not touching) the sample. An advantage of GeNS versus other mechanical loss measurement techniques such as clamped cantilever excitation, or disk wire suspension, is that the contact surface is minimised, and there is an absence of applied forces from the setup to the sample [184]. The sample has one relatively small point of contact with the sphere, which compared to other loss measurement techniques with larger contact surfaces, or multiple points of contact, potentially reduces the influence of friction damping. Additionally to this, it is a nodal support so there is little motion at the point of contact (for the excited modes) that could experience friction damping in the first place. GeNS was coined ‘gentle’ as the vibrating disk can freely roll over the sphere with the only applied force present being the samples own weight [184].




Figure 3.3: Diagram of a disk of thickness t balanced on a hemisphere lens of diameter D . Geometrical parameters useful to calculate the equilibrium condition in GeNS are shown.

The GeNS systems at the University of Glasgow used to characterise disk samples employ 1-inch wide, 60.4 mm radius of curvature (i.e. equivalent spherical diameter $D = 120.8$ mm) crystalline silicon plano-convex lenses, upon which disks are suspended. Figure 3.3 illustrates a suspended disk on a spherical lens and the geometrical parameters useful to define the equilibrium condition. The angular position of the disk with respect to the horizontal plane is defined as θ . Considering the vertical position h of the centre of mass as a function of θ , one can say [185]:

$$h(\theta) = \frac{D}{2} \cos \theta + \frac{D\theta}{2} \sin \theta + \frac{t}{2} \cos \theta. \quad (3.2)$$

where D is the diameter of the sphere corresponding to the curvature of the lens, and t is the thickness of the disk. Assuming that the static friction is sufficient to prevent a sample slipping during oscillation about the balance point, it follows that a stable equilibrium is achieved whenever the diameter of the sphere is greater than the thickness of the disk, i.e. $D > t$. In this regime, as the disk rolls over the sphere the centre of mass is raised for a wide range of tilting angles. If instead $D < t$, then conversely h is always decreasing away from the centre point with angular roll and the disk would fall. These two cases are illustrated in Figure 3.4.

As an example, for a sphere diameter 10% greater than the disk thickness, it follows that the angular range where stable oscillations are possible is therefore $\sim 30^\circ$, which corresponds to disk oscillation far greater than typically induced with a comb exciter, alleviating the concern of knocking a sample off the sphere. Further, for a setup with a fixed lens curvature/sphere diameter, this stability range rapidly approaches 90° with decreasing disk thickness. For all samples presented in this thesis the stability range




Figure 3.4: Vertical position of the centre of mass h as a function of angular position of the disk θ (in arbitrary units), when $D < t$ (left) and when $D > t$ (right) - the latter being the equilibrium condition. Specific values input in the second case are $D = 120.8$ mm and $t = 2.7$ mm as representative of most samples in this study.

can be calculated to be $\sim 89^\circ$. The thickest disks studied were still 42 times smaller than the sphere diameter yielding this $\sim 89^\circ$ stability. However, as a convex lens represents a fraction of a spherical hemisphere, the stability range will shrink, but it is still far greater than any motion which could be induced by the comb exciter.

The Glasgow system consists of two GeNS platforms held within the same vacuum tank, such that two samples can be measured simultaneously. In each, an electrostatic drive plate (comb actuator) is placed a few millimetres above the sample and driven with an oscillating high-voltage AC signal at the frequency of the mode to be excited, with DC offset such that the AC excitation is always positive, and a 633 nm laser beam of a few mm in diameter is used to probe the displacement of the disk. The first system makes use of a custom-built split photodiode, upon which the reflected laser beam from the stationary sample is centred, and once a mode of the sample is excited, the sensor is used to monitor the displacement of the disk. The second GeNS instead uses a commercial SIOS GmbH vibrometer [186] to record the disk displacement. Both of these configurations form optical levers and allow for small motions of the probed part of the excited disk to be recorded as a function of time and frequency during oscillation. An illustration and image of the setup are shown in Figure 3.5.

Before any sample is suspended, the horizontal plane must be found, which is achieved by placing an opaque dish filled with water where a sample would reside while suspended, with the laser beam reflecting from the water at the centre of the dish, to avoid any optical distortions associated with the meniscus at the edge of the dish. The water surface orientation defines the horizontal plane as it is independent of any tank/floor tilt. The laser




Figure 3.5: (Left) Top-down illustration of the GeNS setup placed inside a vacuum tank, with vibrometer readout - adapted from [182]. (Right) Annotated image of a disk balanced in the setup. Note, in both, the electrostatic comb actuators (placed 1 mm above the disk during excitation) have been omitted).

beam is then set to enter the tank in the horizontal plane, and elevated a few centimetres above where a sample would rest on the spherical lens, and is then directed to and returns from the suspended sample/water dish via reflection from a 45° mirror inside the vacuum tank - see Figure 3.5. In the system with the split photodiode, the laser enters the tank with a small horizontal tilt (i.e. normal in the vertical plane, but with a small angular offset in the horizontal plane). This small horizontal offset is introduced in the split photodiode system to avoid the laser beam clipping its own casing on return from the tank, allowing a second steering mirror to be placed beyond this at a large distance (~ 2.5 metres) from the sample before directing the return laser spot to the split photodiode. This creates a large optical path length of in total $\sim 4\text{-}5$ m for the return spot which aids measurement sensitivity. The vibrometer system does not have this tilt introduced during measurement because, for a reading to be taken, the beam must return to its point of origin. However, a similar horizontal angular offset must still be implemented in this system to determine and align the disk to the water surface/horizontal plane, before steering the return beam back to its origin point (i.e. back to normal incidence).

During alignment, the location of the return laser-spot from the water surface is marked with a target at a distance at least two metres from the system and provides a reference point for suspending a level sample in equilibrium on the sphere. With this reference defined, a sample can be suspended and its tilt adjusted until it sits level in the horizontal plane as defined by this reference point. After these steps have been taken, the vacuum tank is closed and the system is pumped down to a pressure of at most 5×10^{-5} mbar,

and the alignment target/s are used to check for any significant path drift before loss measurements commence. In the event that this evacuation of air from the system led to a large permanent change in return spot position, it meant the disk had slipped on the sphere, and the system would be brought back up to atmosphere, with the disk re-suspended and the process repeated until it remained stable during the vacuum tank evacuation.

3.4.1.1 External factor to consider: gas damping

For mechanical loss measurement all sources of external damping must be suitably minimised, as when they are, the resulting damping of the motion arises solely from internal processes within the sample itself. As discussed in the last section, GeNS is a promising technique for reducing suspension damping by suspending a sample on a vibrational node point. In work by Tait [182], it was shown that the same sample measured on a wire suspension (with two contact points) had typically around 20% higher losses and also had far less consistent losses between degenerate modes than when measured on GeNS (with one contact point). By having the sample levelled to the horizontal plane in each successive GeNS measurement, and the only applied force present being the sample's own weight over one small point of contact in the centre, corresponding to points of nodal motion for particular resonances, the friction can be considered well minimised.

The effects of gas damping must also be considered. As an object moves through a gas, the gas molecules resist this motion by creating drag forces. It is for this reason the GeNS measurements are carried out under high vacuum conditions. The level of loss of an oscillator due to gas damping ϕ_{gas} can be expressed as [187]

$$\phi_{gas}(f) \approx \frac{AP}{2\pi mf} \sqrt{\frac{M}{RT}}, \quad (3.3)$$

where A is the surface area of the oscillator, m is its mass, P is the gas pressure, f is the resonant frequency, M is the mass of one molar mass of the gas (for our purpose air), R is the gas constant, and T is the temperature of the system. We can see the gas damping should have greater influence for lower mass samples at lower frequencies, and with increasing pressure. Inserting typical values for air at room temperature, and the

smallest mass and dimensions of a fused SiO_2 disk used in this study (~ 50 mm diameter by ~ 1 mm thickness), and its lowest resonant frequency of around 2.5 kHz, one can calculate for a pressure of 5×10^{-5} mbar the gas damping loss should be $\sim 1 \times 10^{-9}$. This pressure value is around the highest at which measurements would typically begin, and indeed, given the inverse relationship with frequency and mass, is representative of the maximum level of gas damping expected in this experiment. It will be seen later in Figure 3.9 that this level is at least two-orders of magnitude lower than the lowest measured losses recorded on excited samples in these investigations, and thus can be considered negligible.

3.4.2 Fitting ringdown data to extract sample loss

The resonant modes of a disk are located for the first time through a combination of a low resolution broadband frequency excitation search (fast frequency sweep), and the construction of simple analytical and/or finite element analysis (FEA) models of the disks to inform the search. Once the resonant frequencies were found each would be excited individually and allowed to freely decay back to rest, with the (evolving) amplitude of the decaying oscillation being recorded. To excite a mode, an oscillating sinusoidal signal would be passed through the electrostatic drive plate with a frequency a few Hz lower/higher than the located resonance frequency, then the input frequency would be gradually increased/decreased until the peak in the amplitude signal was located, at which point the drive signal would be terminated. Sweeping forward or backward in




Figure 3.6: Ring up (left) and ringdown (right) signal amplitude of the fundamental resonance of one of the disks used in this study. The amplitude of the peak is shown to be near three-orders of magnitude higher than the background.

frequency was found to achieve the same mechanical loss results for any mode within error, though sweeping backward in frequency would also lead to observed peaks a few tenths of Hz lower than when sweeping forward. Tests were also undertaken which verified that the excitation voltage level (when still producing peak amplitudes significantly above background) did not influence decay results either. Figure 3.6 shows representative data for the excitation process, with the left graph showing the excitation of a disk until cutoff at a resonant frequency (where the amplitude peaks), and the right graph showing the decaying amplitude with time after excitation had been terminated. The level of the peak signal can be seen to be nearly three-orders of magnitude higher than the background level. One can also observe a characteristic beating in the oscillatory decay signal.

The beating pattern in the decay signal arises due to imperfections/inhomogeneities in the sample, its geometry and/or orientation with respect to the GeNS lens and electrostatic drive plate. Each measurable resonance mode will have a degenerate mode partner, which will fundamentally have the same displacement profile but rotated close to 45° with respect to the other. Examples of these mode shapes and their pairs can be viewed in Figure 3.10. These should arise at the same exact frequency but due to the aforementioned imperfections will actually be separated by $\lesssim 1$ Hz [182] (and often the level of separation is $\lesssim 0.1$ Hz). As such, a beat emerges in the decay which is observed as a periodic variation in amplitude with rate or beat frequency Δf given by [188]

$$\Delta f = f_A - f_B, \quad (3.4)$$

where A and B represent the two (slightly) different frequencies of the degenerate mode pair. The position of the beam spot on the disk will also affect the amplitude of the signal, as different regions of the disk will be deformed by different amounts during oscillation. The measured level of beating can be affected by this, and it is quite possible to observe next to no beating (i.e. something more akin to a simple exponential decay) if the spot is placed in a region of high motion for one of the twin modes but not the other. The decay, when external damping factors are minimised, depends on the mechanical losses of the resonator material/s, and the beating, decaying signal has to be modelled in order to extract the mechanical loss value of the sample for this mode.

3.4.2.1 The modelled amplitude decay

To extract mechanical loss when two degenerate mode twins, very close in frequency, have been excited, the decay is fitted as the superposition of two decaying waves of the form

$$A(t) = A_0 e^{i(\omega_A t + \varphi_A)} e^{-\frac{t}{\tau_A}}, \quad B(t) = B_0 e^{i(\omega_B t + \varphi_B)} e^{-\frac{t}{\tau_B}}. \quad (3.5)$$

Each defines amplitude evolution with time, with both a periodic oscillatory and a decay component. The amplitude of each waveform at a given time t has initial starting amplitude components A_0/B_0 , angular frequencies ω_A/ω_B which are simply the measured frequencies of each twin mode in Hz multiplied by a factor of 2π , as well as the initial phases of each wave φ_A/φ_B in radians at the start of data acquisition. Each also has a decay component given by the decay constants $\frac{1}{\tau_A}/\frac{1}{\tau_B}$ where

$$\tau_{A/B} = \frac{2}{\omega_{A/B} \phi_{A/B}}, \quad (3.6)$$

with the $\phi_{A/B}$ being taken as the ‘measured’ losses of a given excitation. Note that this is a rearranged form of Equation 3.1. The superposition of two such waveforms yields a decaying sinusoidal wave akin to that measured on GeNS and is given by

$$C(t) = A(t) + B(t) \quad (3.7)$$

$$\begin{aligned} &= A_0 e^{i(\omega_A t + \varphi_A)} e^{-\frac{t}{\tau_A}} + B_0 e^{i(\omega_B t + \varphi_B)} e^{-\frac{t}{\tau_B}} \\ &= A_0 e^{-\frac{t}{\tau_A}} \left(e^{i(\omega_A t + \varphi_A)} + \frac{B_0}{A_0} e^{i(\omega_B t + \varphi_B)} e^{-\left(\frac{1}{\tau_B} - \frac{1}{\tau_A}\right)t} \right) \\ C(t) &= A_0 e^{-\frac{t}{\tau_A}} \left((1) e^{i(\omega_A t + \varphi_A)} + \left(\frac{B_0}{A_0} e^{-\left(\frac{1}{\tau_B} - \frac{1}{\tau_A}\right)t} \right) e^{i(\omega_B t + \varphi_B)} \right). \end{aligned} \quad (3.8)$$

For more simplistic waveforms with no phase shifts or decay components undergoing beating the superposition can be expressed as

$$\begin{aligned} C(t) &= A(t) + B(t) \\ &= A_0 e^{i(\omega_A t)} + B_0 e^{i(\omega_B t)}. \end{aligned} \quad (3.9)$$

The intensity of such a wave can be represented by taking the absolute square of this value [189]:

$$I(t) = |C(t)|^2 \equiv C(t)C^*(t) \quad (3.10)$$

$$\begin{aligned} &= A_0^2(\cos^2 \omega_A t + \sin^2 \omega_A t) + B_0^2(\cos^2 \omega_B t + \sin^2 \omega_B t) \\ &\quad + 2A_0B_0(\cos \omega_A t \cos \omega_B t + \sin \omega_A t \sin \omega_B t) \\ &= A_0^2 + B_0^2 + 2A_0B_0 \cos(\omega_A - \omega_B)t \end{aligned}$$

$$I(t) = A_0^2 + B_0^2 + 2A_0B_0 \cos(\Delta\omega t). \quad (3.11)$$

Similarly, Equation 3.11 can be adapted to express the more complex beating amplitude temporal evolution, incorporating the decay factor and phase difference of each component wave as outlined in Equation 3.8, where the prefactor terms are not simply just A_0 and B_0 , but use the prefactor components in Equation 3.8 to yield:

$$\begin{aligned} I(t) &= \left(A_0 e^{-\frac{t}{|\tau_A|}} \right)^2 \left(1^2 + \left(\frac{B_0}{A_0} e^{-\left(\frac{1}{|\tau_B|} - \frac{1}{|\tau_A|}\right)t} \right)^2 + 2(1) \left(\frac{B_0}{A_0} e^{-\left(\frac{1}{|\tau_B|} - \frac{1}{|\tau_A|}\right)t} \right) \cos(\Delta\omega t + \Delta\phi) \right) \\ &= A_0^2 e^{-2\frac{t}{|\tau_A|}} \left(1 + \left(\frac{B_0}{A_0} \right)^2 e^{-2\left(\frac{1}{|\tau_B|} - \frac{1}{|\tau_A|}\right)t} + 2\frac{B_0}{A_0} e^{-\left(\frac{1}{|\tau_B|} - \frac{1}{|\tau_A|}\right)t} \cos(\Delta\omega t + \Delta\phi) \right) \end{aligned} \quad (3.12)$$

$$I(t) = \left(A_0 e^{-\frac{t}{|\tau_A|}} \sqrt{1 + \left(\frac{B_0}{A_0} \right)^2 e^{-2\left(\frac{1}{|\tau_B|} - \frac{1}{|\tau_A|}\right)t} + 2\frac{B_0}{A_0} e^{-\left(\frac{1}{|\tau_B|} - \frac{1}{|\tau_A|}\right)t} \cos(\Delta\omega t + \Delta\phi)} \right)^2. \quad (3.13)$$

Here A_0 , B_0 and $\Delta\omega$ are defined as real positive numbers. The form of the intensity evolution of the beating signal as shown in Equation 3.13 is useful for modelling and interpretation. Specifically if a model where the two degenerate modes have different decay constants is fit, the equation can be used as is, whereas if the decay constants are equal it is easy to see how the equation reduces. The final fit for a beating and decaying amplitude signal takes in the intensity in Equation 3.13 as input, as well as an estimated (small) background level from the data, yielding:

$$C(t) = \sqrt{I(t) + C_{\text{background}}^2}. \quad (3.14)$$

The individual parameters need to be estimated and refined in order to fit the model to the data via Equations 3.13 - 3.14. The flow diagram for this fitting process is shown in Figure 3.7. First, all ringdown data is screened and cleaned to ensure e.g. none captures a large portion of background signal after detectable oscillation has ceased, and no signal spikes arising from electronic interference remain, both of which can lead to poor fitting. A simple exponential decay is fit to all the passed data which allows for initial estimates of A_0 , τ_A and an upper limit on the value of B_0 . Next, the exponential fit is subtracted




Figure 3.7: Flow diagram showing the major processes used to perform the optimised fit of Equations 3.13–3.14 to the beating ringdown data.

from the measured beating decay signal in order to produce a flattened sinusoidal trend, to which the beat frequency is fitted. With these results, estimates for the phase shift are trialled iteratively until the best fitting phase shift of the beating decay signal is found. Lastly, with all of these initial parameter estimates, a fit of the form of Equation 3.14 is made, by first setting $\tau_B \neq \tau_A$ and trialling multiple iterations of all the parameters, until an optimised fit with minimised residuals of the difference between the modelled and measured vibrational amplitude time dependence is found. These minimised residual τ values are taken and via Equation 3.6 yield ϕ_A and ϕ_B i.e. the best estimate of the mechanical losses of the two degenerate modes for a given measurement.

In Figure 3.8 the measured ringdown data and final best-fit for a particular excitation of a blank silica disk are shown, as well as the residuals between the data and fit. The fit of the losses of the two excited degenerate modes close in frequency are displayed in the figure, along with other fit variables. Typically in practice at least six repeat measurements are




Figure 3.8: Comparison between the measured and best-fit ringdown data of the 16.7 kHz mode on a blank 75 mm diameter, 2.7 mm thick silica disk. The lower plot shows the residual data from the fit (the direct difference between best fit trend line and the measured data). Best-fit parameters are also displayed.

carried out of each mode pair per balance of the disk on the sphere, and at least three re-balances/re-seating of the disk on the sphere, hereafter referred to as ‘suspensions’, are made to verify measurement repeatability. Once this repeatability is verified, the lowest mean loss value obtained from the different suspensions is carried forward as a mode’s ‘best’ loss estimate. In repeating this measurement procedure for multiple modes, multiple times per mode, across multiple suspensions of the sample, a distribution in the measured losses can be built, with trends across frequency observed. Then in further repeating this whole process after different temperature, post-deposition heat treatments of the sample, its optimum lower loss state can be found.

3.4.2.2 Factors to consider before accepting the fitted losses

Fundamentally for a perfect homogeneous cylinder suspended precisely on its geometric centre, the losses of degenerate modes should be the same. Any difference in measured loss likely arises due to imperfections in the suspension or in the geometry of the disk leading to one mode’s motion being more damped by external sources than the other. For example, as the sample oscillates, geometry imperfections and suspension off-centring can result in greater sphere surface area contact with a nodal line which defines the motion in a particular mode shape across the disk, and not its twin (with motion $\sim 45^\circ$ rotated across the disk’s face) thus damping its motion more. Contamination and surface inhomogeneities can also potentially lead to damping of one or both twin modes and as such great care must be taken to ensure all suspended samples are well cleaned before measurement, and that the substrates themselves are well polished. As a result of these factors, for a given measurement the lower of ϕ_A and ϕ_B is typically taken forward as the best approximation of the true measured loss.

However, artificially low loss fits can sometimes be generated for one of the two loss values in a beating decay and care must be taken to not choose these values. This typically occurs if very little beating is observed in a particular measured decay, and it appears like a simple exponential. Fitting Equation 3.14 to a decay such as this results in only one of the two fitted loss values being accurate, and both can be compared with the initial

simple exponential fit value to determine which. Artificially low values can also arise if the data to be fitted includes a significant portion of the background noise following the ringdown of the excited motion, so again care is taken to crop the data to encompass as much of the signal as possible where decay is clearly still ongoing.

In fringe cases, unlike the example shown in Figure 3.8, where the fit does not model the ringdown data well, either one, or often both, loss values for the fit must be totally discounted and excluded from subsequent analysis. There are many factors that enter into this decision, such as comparison of both fit loss values with the simple exponential decay modelled loss, the magnitude and shape of the residuals, and comparison with data from the same resonance measured multiple times on the same and across other suspensions. Finally, for a given sample suspension, care must also be taken to verify that there exists no clear evolving trends in the measured losses solely with time. As an example, if the loss of one mode is getting significantly worse with repeat measurement, then the disk is likely not in equilibrium and gradually slipping, in which case all values should be discounted and the disk re-suspended. These checks were applied to every measurement in this and all GeNS studies undertaken.

3.4.2.3 Measured mechanical losses

The procedure and considerations outlined in Section 3.4.2.1 and Section 3.4.2.2 are used to measure the loss of a number of modes of a sample, allowing possible frequency-dependence to be investigated. For each mode, the mean and standard deviation of the loss values obtained from the repeated measurements of a given mode during a particular sample suspension are calculated and plotted. Figure 3.9 shows the four suspensions taken for the first uncoated fused SiO₂ disk measured solely by the author on GeNS; two on the vibrometer setup, two on the split photodiode setup. In general one can see some variability between the different suspensions, but a relatively consistent lower bound level of the loss is observed amongst the four suspensions. These ‘best’ level results are plotted on the right-hand side of Figure 3.9 for each setup, and represent the lowest mean loss found for each mode, across all suspensions. The good agreement between the vibrometer and split photodiode setup is apparent.

For some modes there is relatively high variability between the four suspensions. This highlights the importance of performing multiple suspensions until the ‘true’ loss level for all modes of the sample is found, which is approximated by the lowest loss results with observed repeatability. It should be noted as well that the data from the authors’ first measured disk is presented here, and it has some of the largest suspension variability of any sample measured throughout the work in this thesis. A likely explanation for this is the simple fact that for these measurements the balancing, aligning, and cleaning techniques were still being developed and practiced. Due to the observed equivalence in the best losses between the two setups, going forward, measuring on either system was determined to be equivalent, and best losses acquired irrespective of the measurement setup.

One final consistency observation of note, true to most samples, and illustrated in the representative data in Figure 3.9 is that the last mode around 30 kHz here has generally the highest variability between different suspensions. It also typically exhibits the lowest signal-to-noise of all modes, and decays the quickest due to both its higher loss and relatively low excitation level above the background. As such it is often harder to measure and fit this mode well, which the higher suspension-to-suspension variability reflects. The




Figure 3.9: (Left) Mechanical loss values measured for eight modes of a blank SiO_2 disk between 2–32 kHz, after re-balancing twice on the two GeNS systems at Glasgow; one with a vibrometer read-out, one with a split photodiode. (Right) The lowest mean loss values obtained between the multiple repeated measurements of each mode acquired between the two suspensions on either system.

lowest losses of this mode from suspensions VB2 and PD1 shown here had higher SNR than the other two suspensions, likely resulting from the laser spot reflecting from a region of significantly higher motion, resulting in more consistent results than for the other two suspensions.

The graph of the final ‘best’ losses in Figure 3.9 (and indeed for all the disks measured in this study) appears to show two distinct loss trends with frequency on these first eight measured resonances; a lower loss trend shared by the higher of the two 16 kHz modes and the ~ 25 kHz mode, and a high loss trend with steeper gradient between the other six modes. Though the mode frequencies between the two geometries of disk measured in this study differed, the same trend with mode shape was witnessed. This loss variation with frequency and/or mode shape, is likely to partially originate from a different split into bulk and shear motion dependent on the mode shape [141] or indeed excess losses due to disk face/barrel polishing discrepancies [190]. To better understand this mode shape dependence and eventually pull out coating losses the disk motions were modelled with finite element analysis.

3.4.3 Modelling vibrating disks

Finite element analysis (FEA) was performed on all disk samples undergoing oscillation in GeNS in order to estimate their resonant frequencies, understand the pattern of motion induced at those frequencies, and the elastic strain energy density stored in the coating and substrate for coating loss extraction. The resonant frequencies are found by modelling the disk both coated and uncoated in COMSOL Multiphysics [191] and running its built-in eigenfrequency model to produce resonant frequency estimates for the modelled disk. These estimates are used to assist in experimentally finding and exciting the mode frequencies. Both geometric and material properties of the disk are required to perform accurate analysis. The diameter and thickness of each disk were measured with a set of digital calipers with error of ± 0.1 mm at different points across the sample, which were averaged and input into the model to build the geometry. For amorphous materials, the Young’s modulus, Poisson ratio and density are also required of both the substrate and any coating modelled on top. For fused SiO₂ disks, these substrate properties were taken from generally accepted values [192], with all geometric parameters later being refined for

all uncoated disks, once the real measured frequencies were known, so that the models better reflected these. This was achieved by inputting a range of values for each geometric property, constrained by known uncertainties in each geometric measurement, and parameter optimising until the set of values yielding the closest matching frequencies to the real data were found.

FEA is a numerical modelling method which can be used to investigate an object's physical properties and its response to various applied forces. After defining an object's geometry and its material properties, FEA software divides the geometry into a mesh of individual elements. Each element in the mesh is allocated specific partial differential equations, which the model will use to characterise the element's behaviour in response to some process [193]. An FEA model will then solve each of these differential equations and combine all the results to generate a solution for the entire object. The accuracy of a finite element model generally increases with increasing number of elements. Therefore to ensure reliable results, the number of elements is increased incrementally until the desired output converges - i.e. there is no further significant change with further increase in element density.

For the FEA models used in this thesis work, the number of elements was changed until a model consistently reproducing eigenfrequency estimates and elastic strain energies to within $<1\%$ of the previous number of elements trialled was achieved for all of the first forty predicted modes of a disk - a quantity greater than the number measured on any sample. After these converged models were produced, the resonant mode frequency estimates were recorded. From this, information of their mode shape and percentage of elastic strain energy stored in bulk and shear motion can also be extracted. Indeed, one learns that the two modes with a lower loss trend at 16 kHz and 25 kHz, shown on the right side of Figure 3.9, exhibited a greater fraction of bulk motion and are both from a separate family of mode than the other six studied modes.




Figure 3.10: Finite element analysis representation of the first twenty-two resonant mode shapes of a vibrating disk, grouped by family, with each twin mode paired. For a given mode shape m/n represent the number of present nodal circles/lines.

Figure 3.10 represents the first twenty two modes predicted with FEA, grouped by ‘family’. Mode shapes are defined by the notation (m,n) where m is the number of nodal circles and n is the number of nodal lines representing the regions/shape of displacement. Modes with zero nodal circles ($m = 0$) represent the ‘clover’ modes, and ones where both $m,n \neq 0$ represent the ‘mixed’ modes. The first sixteen modes present between these two families on a thin disk are represented in Figure 3.10, and it is these which are measured across all the disks in this study. Observed beating in the ringdowns result from the interference between twin modes close in frequency, with identical shape that is rotated $\sim 45^\circ$ with respect to each other.

The other two families of modes represented in Figure 3.10 are the ‘drumhead’ modes with no nodal lines ($n = 0$), and the ‘rocking’ modes with only one nodal line ($n = 1$). Neither of these are shown in the following results but are worth mentioning. The drumhead modes cannot be measured on GeNS as the vast majority of their motion is concentrated at the centre of the geometry, at/close to the suspension point and as such are totally friction damped by interaction with the sphere/lens on which the disk is balanced. This is not the case for the clover or mixed modes presented in these studies, which have nodes at the centre of the disk and most motion occurs close to the edges of the disk, far from the balance point. Similarly the rocking modes, so named as when excited the disk begins to ‘see-saw’ back and forth in one axis, are often not detected due to both significant motion near the centre and mostly minimal displacement along the vast majority of the edge region. However, they can occasionally be measured in GeNS if the laser spot happens to coincide with an edge region of higher motion, but when this happens their measured losses are typically much higher than clover or mixed modes due to the damping experienced near the centre as the sample rolls back and forth over the sphere, so are discounted as true sample mechanical loss measurements with GeNS.

Figure 3.11 shows, as a visual aid, the mode shape of each measured mode plotted alongside the measured losses from Figure 3.9. Though the exact frequencies and magnitude of the losses differed amongst the different disks, this same general trend in lower measured loss of the mixed modes versus the clover modes, with the clover losses increasing more rapidly with frequency, was seen on all samples.




Figure 3.11: Finite element analysis representation of the first twenty two resonant mode shapes of a vibrating disk, grouped by family, with each twin mode paired. For a given mode shape m/n represent the number of present nodal circles/lines.

3.4.3.1 Modelling a coated HR stack topology

Modelling multiple connected thin interfaces in FEA comes with a host of associated complexities. Due to this the $\text{TiO}_2:\text{SiO}_2/\text{SiO}_2$ coating stacks in this study are approximated in the FEA models as a composite mono layer of thickness equal to the total thickness of the stack and of composite material properties given by suitable thickness weighted averaging of each individual material's properties. Specifically the composite Young's modulus can be expressed as a weighted sum [76, 182]

$$Y_{\text{composite}} = \frac{\sum_{i=1}^n Y_i t_i}{\sum_{i=1}^n t_i} \quad (3.15)$$

where Y represents Young's modulus, t the thickness, and subscript i denotes the different materials to be averaged. The composite density was calculated following this same form. Given that there are two materials to be averaged for the coating stacks in this study, and each layer of the same material in this stack are nominally the same thickness Equation 3.15 becomes

$$Y_{\text{composite}} = \frac{Y_1 t_1 + Y_2 t_2}{t_1 + t_2}, \quad (3.16)$$

The composite Poisson ratio can be calculated as follows [76, 182]

$$\nu_{\text{composite}} = \frac{Y_1 t_1 \nu_1 (1 - \nu_1^2) + Y_2 t_2 \nu_2 (1 - \nu_2^2)}{Y_1 t_1 \nu_1 (1 - \nu_1^2) + Y_2 t_2 \nu_2 (1 - \nu_2^2)}. \quad (3.17)$$

The composite Young's modulus, density, and Poisson ratio can be used in combination with the stack thickness to produce a model of a coated disk to be solved with FEA. The two HR stacks in this study each had a specified number of quarter-wavelength optical thickness layers (QWL) at 1064 nm. This means the total thickness of the stack t_{stack} can be defined as

$$t_{\text{stack}} = \sum_{i=1}^n \frac{1064[\text{nm}]}{4n_i} N_i = \sum_{i=1}^n t_i N_i, \quad (3.18)$$

where n represents the refractive index of material i , N the total number of layers of that material, and t is its physical thickness. HR-1 had 18 TiO₂:SiO₂ layers (63.2% Ti composition) and 17 SiO₂ layers, each $\lambda/4n$ thickness, making its total thickness ~ 5.7 microns. HR-2 was specified for full gravitational wave detector ETM reflectivity with 22 TiO₂:SiO₂ layers (69.5% Ti composition) and 21 SiO₂ layers, each $\lambda/4n$ thickness, making its total thickness ~ 6.8 microns.

3.4.3.2 Inferring doped material elastic properties

To perform the single composite layer HR stack FEA approximation, as outlined in the previous subsection, first the material properties of both the mixed TiO₂:SiO₂ and the SiO₂ comprising the stack had to be estimated. Nominal amorphous IBS SiO₂ Young's modulus, density and Poisson values were taken from the literature [165], and significant deviation from these is not expected. For the mixed material, density values were acquired for the single layer coatings from Rutherford backscattering spectrometry measurements as shown previously in Table 3.1, and these values were broadly consistent with being a mix of the densities of pure IBS TiO₂ and SiO₂. The densities of the mixed layers in the HR coating stack (which differed in cation content from the single layer counterparts) were estimated via volumetric weighted averaging of values for pure TiO₂ and SiO₂, following a similar form as Equation 3.16.

To estimate the Young's modulus and Poisson ratio of the mixed material, the author built an analytic model in MATLAB following relations derived by S. Barta [194]. These relations describe the effective Young's modulus and effective Poisson ratio of a multi-component particulate mixed material in the framework of the average field approximation. The average field approximation method follows from the notion that a randomly chosen isotropic granule of globular form with characteristic Young's modulus Y and Poisson ratio ν is submerged into an (unlimited) effective medium, with characteristic effective Young's modulus Y_{EFF} and effective Poisson's ratio ν_{EFF} [194]. The parameters Y and ν are stochastic quantities, which are determined through n-point correlation functions. The average field approximation method only requires knowledge of the local distribution function and not the n-point correlation functions. On the one hand, this is advantageous as only local distribution information is needed, but on the other hand, being an approximation it uses incomplete information of the structure and statistics of the mixed material on a sub-macroscopic level.

Barta discusses these limitations of this average field approximation and indeed after a long derivation defines two equations required to be solved simultaneously to yield the relations for the effective material parameters - Equations 3.19–3.20. Despite the simplifications, it was determined that this average field approximation approach was suitable for calculating effective mechanical-elastic properties in this study, as has been successfully done prior in others [195, 196, 197]. The equations are obtained by means of the average stress and average linear strain tensors and through defining an effective material tensor. Their solution gives the relations for effective parameters and are defined as follows [194]

$$\sum_{i=1}^n c_i \frac{A - A_i}{2A + \frac{A_i}{B_i}(\mu_i + 1)} = 0 \quad (3.19)$$

$$\sum_{i=1}^n c_i \frac{\frac{A}{B} - \frac{A_i}{B_i}}{2A + \frac{A_i}{B_i}(\mu_i + 1)} = 0, \quad (3.20)$$

where c_i is the volumetric fraction of each constituent material i in the mix, μ_i is defined as the Poisson constant (inverse of the Poisson ratio) such that

$$\mu_{\text{EFF}} = \frac{1}{\nu_{\text{EFF}}}, \quad \mu_i = \frac{1}{\nu_i}, \quad (3.21)$$

and the variables A and B are defined as follows

$$A = \frac{\mu_{\text{EFF}} Y_{\text{EFF}}}{\mu_{\text{EFF}} + 1}, \quad A_i = \frac{\mu_i Y_i}{\mu_i + 1} \quad (3.22)$$

$$B = \mu_{\text{EFF}} - 2, \quad B_i = \mu_i - 2. \quad (3.23)$$

From these relations one can see how the effective Young's modulus and Poisson ratio depend on the volume fraction and the properties of the individual components. Once the simultaneous equations are solved to yield values for A and B , Equation 3.21 and Equation 3.23 can be used together to solve for the effective Poisson ratio of the mix, and then this can be input into Equation 3.22 to solve for its effective Young's modulus. It should be noted the solution of Equations 3.19–3.20 in fact yields two possible values for B and hence for the effective Poisson ratio. One of these solutions lies between the Poisson ratios of the two pure materials and one does not and as such the latter is treated as nonphysical and discarded. Notably the solutions of Equations 3.19–3.20 can generally be obtained only by numeric methods, such as those incorporated with the model the author developed using MATLAB's Symbolic Math Toolbox [198].

In lack of directly measured material properties of pure TiO_2 and SiO_2 made in the same IBS system under the same conditions, values for the Young's modulus and Poisson ratio of the pure IBS materials taken from the literature were used in this analysis and outlined in Table 3.2, along with the effective material properties calculated for the mixed material layers of both stacks. In order to achieve this, the cation concentrations of both TiO_2 and SiO_2 acquired via ERD-TOF for the mixed material layers of each stack had to be converted into a volumetric fraction, by considering the atomic mass and density of each. The mixed material was modelled as a binary system where only TiO_2 and SiO_2 were considered, and the known impurities such as argon and carbon comprising at maximum of around 3.5% of the real mixed materials were not. Figure 3.12 shows how the calculated mixed material properties varied with cation content. Once the mixed material properties for the desired Ti cation concentration were found, the FEA composite mono layer HR stack material properties were calculated using Equations 3.16-3.18 and are also detailed in Table 3.2.




Figure 3.12: Effective Young's moduli and Poisson ratios calculated for mixed $\text{TiO}_2:\text{SiO}_2$ materials from Ti cation concentrations of 1% – 99% (in 1% steps).

Table 3.2: Material properties of pure and mixed IBS SiO_2 and TiO_2 , as well as a composite stack property estimates.

| Material | ρ [kg/m ³] | Y [GPa] | ν |
|--|-----------------------------|---------|-------|
| SiO_2 | 2202 | 72.0 | 0.170 |
| TiO_2 | 4230 | 151.0 | 0.280 |
| $\text{TiO}_2:\text{SiO}_2$ (63.2% Ti) | 3303 | 103.8 | 0.241 |
| HR-1 composite (63.2% Ti) | 2695 | 86.2 | 0.202 |
| $\text{TiO}_2:\text{SiO}_2$ (69.5% Ti) | 3443 | 109.5 | 0.249 |
| HR-2 composite (69.5% Ti) | 2746 | 88.5 | 0.204 |

3.4.4 Extracting coating loss from measured losses

Measurements of the loss of a sample before and after being coated can be used to find the loss of the coating itself. To do this, it is necessary to know what fraction of the total elastic energy of the sample, vibrating with some mode shape, is stored in the coating - as it is only this energy which 'experiences' the loss of the coating material. With the mode shapes of the measured resonances identified, and converged well-matching FEA models produced for the same sample both coated and uncoated, the elastic strain energy residing in the coating and substrate for all the excited modes can be extracted. These can then be used to infer the coating mechanical loss.

The measured loss ϕ_M of a coated sample can be expressed as

$$\phi_M = \frac{1}{E_T} (E_C \phi_C + E_S \phi_S), \quad (3.24)$$

where E_T is the total elastic strain energy stored in the sample undergoing some particular motion, and E_C and E_S are the portions of this energy stored in the coating and substrate respectively ($E_T = E_C + E_S$), with ϕ_C and ϕ_S being the corresponding intrinsic coating and substrate losses. Both ϕ_M and ϕ_S are measured using the GeNS, and the elastic strain distribution in a given mode shape is calculated through FEA. This equation can be rearranged to find the coating mechanical loss [140]

$$\phi_C = \frac{E_T}{E_C} (\phi_M - \phi_S) + \phi_S. \quad (3.25)$$

Of note, the prefactor term in Equation 3.25 is the inverse of a term sometimes defined in the literature as the dilution factor D – the ratio of elastic energy stored in the coating to that in the full sample defined as [165, 199]

$$D = \frac{E_C}{E_T} = \frac{E_C}{E_C + E_S}. \quad (3.26)$$

Following an approach highlighted by Reid and Martin [200], Equation 3.25 can be simplified when the coating volume is much smaller than that of the substrate (e.g. the coating thickness is much thinner than the substrate thickness). In this case, the strain energy stored inside the substrate is much larger than that stored in the coating i.e. $E_C \ll E_S$ and so $E_S \approx E_T$. Therefore for a given resonant frequency f_0 , the coating mechanical loss can be approximated as [200]

$$\phi_C(f_0) = \frac{E_S}{E_C} (\phi_M(f_0) - \phi_S(f_0)). \quad (3.27)$$

This equation frame work allows for extraction of the intrinsic coating loss under the assumption that thermoelastic loss of the coating material is not a significant contributor to the overall dissipation. If it was significant then Equations 3.25 and 3.27 would instead represent the sum of the intrinsic coating loss, and coating thermoelastic loss. This $\frac{E_S}{E_C}$ term defined as the ‘energy ratio’ is approximately equal to the inverse of the dilution

factor defined in Equation 3.26 when $E_C \ll E_S$ and can be extracted directly from the FEA model. The energy ratios for each of the four HR stack coated samples in this study are shown in Figure 3.13 and Table 3.3. Note that lower values for the mixed modes compared with the clover modes can be observed.

The different energy ratio values for different samples arise from variation in their physical dimensions. Both HR-2 samples were of similar geometries and as such have very similar energy ratios to each other. HR-1 A was a sample with coating deposited onto a substrate with similar dimensions to both the HR-2 samples, its energy ratios are higher, reflecting the thinner HR-1 coating with a lower value for E_C . With these ratios computed, the analysis framework is complete and all measured ringdown data measured using the GeNS can be converted into coating mechanical loss values, with trends through frequency, mode shape and post-deposition heat treatment able to be extracted.

One aside to this analysis is that E_C can be further split into the contributions of strain energy in the coating stored in bulk ($E_{C,\kappa}$) and shear $E_{C,\mu}$ (where $E_C = E_{C,\kappa} + E_{C,\mu}$). One finds, while both families of mode have a majority of their elastic strain energy stored in shear, the mixed modes generally possess a significantly higher proportion in bulk. This, however, was not necessary for these $\text{TiO}_2:\text{SiO}_2$ coating studies for converting the coating mechanical loss into CTN values.

Table 3.3: Energy ratios calculated in COMSOL Multiphysics for the four coated $\text{TiO}_2:\text{SiO}_2/\text{SiO}_2$ HR stack samples approximating the stacks as a composite mono layer.

| Stack energy ratios sorted by mode shape | | | | |
|--|--------|--------|--------|--------|
| Mode | HR-1 A | HR-1 B | HR-2 A | HR-2 B |
| (0,2) | 134.6 | 52.3 | 109.5 | 109.5 |
| (0,3) | 135.9 | 52.2 | 110.6 | 110.6 |
| (0,4) | 137.6 | 52.3 | 111.9 | 111.9 |
| (0,5) | 139.5 | 52.4 | 113.5 | 113.5 |
| (1,2) | 131.4 | 51.4 | 106.9 | 106.9 |
| (0,6) | 141.7 | 52.6 | 115.3 | 115.3 |
| (1,3) | 137.8 | 51.8 | 108.9 | 108.9 |
| (0,7) | 144.1 | 52.8 | 117.2 | 117.3 |




Figure 3.13: The elastic strain energy ratios calculated for the resonant modes of the four samples coated with the two $\text{TiO}_2\text{:SiO}_2/\text{SiO}_2$ HR stacks. The values for HR-1 B coated on a much thinner, and smaller diameter disk than the other three samples is displayed on the inlay with axes representing the same terms as the main plot.

3.4.5 Titania-silica coating mechanical loss results

Following the measurement procedure outlined in Section 3.4.1 to the measure the mechanical losses of fused SiO_2 disks before and after coating via gentle nodal suspension, and applying the analysis framework described from Sections 3.4.2 - 3.4.4 the mechanical loss of the coating can be extracted. Samples coated with the two highly reflecting coating stack designs incorporating QWLs of $\text{TiO}_2\text{:SiO}_2$ and SiO_2 as the high and low index layers with topology described in Section 3.3 were studied. HR-1 was determined to have nominally 63.2% Ti in its mixed layers with HR-2 having nominally 69.5% Ti, as shown in Figure 3.2.

Each coating was deposited onto two disks; four samples in total, denoted HR-1 A, HR-1 B, HR-2 A, HR-2 B. Sample HR-1 B had approximate dimensions of ~ 50 mm diameter and thickness ~ 1 mm, while all the others had approximate dimensions ~ 75 mm diameter and thickness ~ 2.7 mm. Their individual dimensions were verified from multiple point measurements per disk with a set of digital calipers. Figure 3.14 shows the as-deposited coating mechanical results for all four of these samples with frequency. Each data point

represents the ‘best’ coating loss values acquired after multiple GeNS suspensions, where for each coated disk suspension, the mean of the losses from multiple measurements of each resonant mode was calculated, and the lowest of these mean values between the different suspensions is taken as the best value. The corresponding best uncoated (blank disk) values are subtracted from these, and the difference is scaled by the elastic strain energy ratio between coating and substrate for the relevant mode as defined in Equation 3.27. The plotted coating loss error bars represent the corresponding standard deviation of the coated and uncoated best results added in quadrature and also scaled by the energy ratio.

With an average of 6 repeated measurements per mode, per suspension, at least 3 suspensions per state of the disk, and both the uncoated and coated disk data considered to extract the coating losses, this corresponds to typically between 10-15 measurements being represented by each data point on the coating loss versus frequency plots presented. For all the individual measurements a fit similar to that represented in Figure 3.8 is made towards performing this analysis.

From Figure 3.14 a clear agreement between the measurements of both samples of each coating can be seen. Particularly the HR-1 data, where two different geometries of disk were coated, highlights the success of isolating coating loss independent of sample dimensions. The coating mechanical loss is an intrinsic material property and should not be influenced by the geometry of the sample onto which it is deposited, and indeed for all samples very good agreement in coating loss is observed. There are 13 measurements from which direct comparisons can be drawn for both coatings, with all but one of these the




Figure 3.14: As-deposited coating mechanical losses as a function of frequency for both samples coated with HR-1 (left) and both coated with HR-2 (right).

values, for one mode of one sample, lying within measurement error of the other sample deposited with the same coating. The agreement between samples is apparent. A few less modes were measured for sample HR-2 B as-deposited than for the other three disks in the as-deposited state, but all modes were measured for all samples in the future runs after heat treatment.

As well as agreement between the samples with the same coating, one can also see very similar levels of loss between the two different coatings, as may have been expected for coating stacks comprised of thicker pure-SiO₂ layers and of thinner layers of mixed-material differing in Ti content by only $\sim 5\%$. Each exhibits a broadly constant loss trend with frequency at a level of $\sim 4\text{-}5 \times 10^{-4}$, suggesting that both high-index materials have very similar losses as-deposited. This flat loss trend with frequency suggests no large difference between the bulk and shear loss of the coatings as well, as this would be reflected with the mixed modes (with a significantly greater portion of energy stored in bulk motion) having different losses than the clover modes.

It is well known that heat treatment can significantly reduce the mechanical loss of bulk fused SiO₂. This reduction in loss is believed to be a result of a reduction of internal stresses in the SiO₂ [201]. It is also believed for many amorphous solids above 10 K, the internal dissipation largely results from thermally activated re-orientations of bonds over potential barriers [202] and that heat treatment possibly alters the distribution of these potential barrier heights [203] reducing the loss. Ion-beam sputtered coatings are often heat treated after coating deposition to reduce the stress in the coating and to reduce the optical absorption [204] and this is indeed true for current GWD coatings.

Each sample was put through various rounds of progressive post-deposition heat treatment, with the losses being recorded in order to find which temperature yielded the lowest losses of the coatings, and therefore potentially would lead to the largest improvement in GWD coating thermal noise levels. Each sample was heat treated to the target temperature in air, similar to the procedure undertaken for current GWD coatings. Specifically the samples, in all bar one instance described later, were heat treated to target temperature with a controlled ramp up rate of 100°C/hr, and held at the desired temperature for 10 hr

before being allowed to freely cool back down to room temperature, and removed from the oven for measurements. At least one common temperature point was reached between the two samples deposited with the same coating to further check consistency. The results of these studies will now be discussed.

3.4.5.1 Coating loss of the 63.2% Ti HR stack through heat treatment

Figure 3.15 shows the coating loss of the samples coated with HR-1 versus heat treatment temperature, excluding the 700°C heat treatment of HR-1 B, which is shown in Figure 3.16. The coating losses from both HR-1 disks results maintain good agreement with each other after heat treatment. The results from the common heat treatment of 550°C yielded near identical coating losses from each disk - taking the average of all modes for comparison (which exhibit no trend with frequency), the loss of the two samples agreed to within 2%, which is also less than the variance exhibited between modes of the same sample of 3% and 8% respectively. Neither sample had the same heat treatment prior to this step, suggesting in this case that heat treatment temperature was more dominant to the overall coating loss than heat treatment history. A general decreasing trend in loss




Figure 3.15: Coating mechanical losses as a function of frequency for the 63.2% Ti $\text{TiO}_2\text{:SiO}_2/\text{SiO}_2$ HR stack, after different stages of post-deposition heat treatment. The results for sample HR-1 A are plotted on the left, and for HR-1 B on the right.

with heat treatment temperature, with the losses again being by and large constant across the frequencies measured was observed until 600°C. At 650°C (on HR-1 B, shown by the red pentagons) a loss split between the mode families was witnessed, with both the mixed modes ($>0,n$) now exhibiting higher coating losses than the clover modes (0,n).

This loss increase at 650°C was accompanied by macroscopic physical changes to the samples, with large blisters of a few millimetres in size noticed after heat treatment. Blister formation in optical coatings results from issues with the adhesion/cohesion [205]. Specifically, blisters can arise during post-heat treatment cooling if the contraction of the coating differs significantly from that of the substrate, giving rise to compressive stress. As such, blisters result from so called compressive stress failure [206]. This stress can lead to the formation of blisters, as the coating tries to shrink but is constrained by the substrate. As the mixed modes have a larger proportion of their motion across the face of the disk (and larger bulk strain energy storage) than the clover modes do (see Figure 3.10), it is conceivable that these inhomogeneities across the face led to higher measured losses in these modes. It is also possible that the blister formation led to micro-cracking around the defect regions which could also have increased the losses, though unfortunately this data is not available for this particular sample as the disk was subsequently heat treated to 700°C before it could be examined under an optical darkfield microscope. Though the effect on the clover modes is also likely not zero, the excess loss of the mixed modes of the sample at 650°C is clearly apparent for HR-1, and as such the losses for the measured (1,2) and (1,3) modes will be discounted from subsequent analysis and averaging of values at this temperature step.

The lowest loss results for this coating arose after heat treatment at 600°C before any noticeable defects manifested. This was found on sample HR-1 A, with loss of $(1.32 \pm 0.16) \times 10^{-4}$, though the losses of the clover modes after 650°C on HR-1 B were quite similar, and it was decided to heat treat this sample which had already blistered further to observe whether the clover modes continued the downward trend in loss, and to study if more macroscopic defects arose. HR-1 B was therefore heat treated next to 700°C with coating loss results shown in Figure 3.16. After this heat treatment the sample now also exhibited macroscopic cracking in the coating across the whole diameter of the sample with more blisters forming. Cracking, in contrast to blister formation (with too great compressive

stress), results from tensile stress failure [206]. A material experiences tensile stress when it is elongated or pulled and compressive stress when it is compacted. Too high levels of tensile stress has been well documented to induce cracks in optical coatings [206, 207]. It is apparent that the sample experienced critically high levels of (various kinds of) stress as its material properties (of both the coating layers and substrate) were altered through the heat treatment process to these temperatures. Details of the different samples for which defects arose and their corresponding heat treatments procedures are shown in Table 3.4. At 700°C the loss of HR-1 no longer appeared to be flat across frequency, and nor did it seem to have a significant loss split between mode families as at 650°C for the same sample. Instead it exhibited a general increasing trend in loss with frequency, with all modes now exhibiting higher losses than at the previous heat treatment step.

The losses of this sample at 700°C were witnessed to become progressively worse with each re-suspension of the sample (three were attempted), and indeed the losses also became progressively worse with each repeat measurement on the same suspension. Disk slipping for these was ruled out by monitoring the stable return laser spot position, and another disk was simultaneously measured on the other GeNS system inside the same vacuum




Figure 3.16: Coating mechanical losses as a function of frequency for sample HR-1 B, after different stages of post-deposition heat treatment, including the cracked coating 700°C measurement.

tank, and did not exhibit this worsening trend with time (ruling out contamination) which meant the increasing losses could only have come from sample change. The number of cracks was in fact witnessed to increase from the first to the third suspension, and the accompanying increasing losses with these seemed to strongly suggest correlation between the cracking and measured loss of the disk resonator. It is therefore believed that the numerous cracks were the primary factor which led to this increase in loss, and the sudden divergence from the previously witnessed flat loss trend with frequency. As such the values reported in Figure 3.16 for 700°C come only from the first suspension.

The mean loss acquired from all measured modes of these HR-1 samples through heat treatment is presented later in Figure 3.19 (excluding this 700°C suspension). Their results are discussed alongside the results from the samples coated with HR-2, the latter of which shall now be discussed.

3.4.5.2 Coating loss of the 69.5% Ti HR stack through heat treatment

A similar investigation into the effects of post-deposition heat treatment on HR-2 with 69.5% Ti high index $\text{TiO}_2\text{:SiO}_2$ layers was conducted. This coating was manufactured at a much later date and during this time it was discovered that the oven used to heat treat the samples in Glasgow had developed a fault which meant as the samples were set to cool at a specific rate from their dwell temperature, they instead cooled naturally back to room temperature following a broadly exponential cooling decay. As such for these samples it was decided to test in another oven if the cool rate was controlled, if defects could be avoided at the higher temperatures. Therefore when both samples HR-2 A and HR-2 B were heat treated to 600°C, HR-2 B was placed in an oven with 50°C/hr controlled cooling, whereas HR-2 A (and indeed the previous temperature steps of HR-2 B) followed the heating procedure outlined in Section 3.4.5.

Figure 3.17 shows the post-deposition heat treatment coating loss evolution for each of the HR-2 coated samples. The as-deposited data is not plotted here to provide a more zoomed in view of the fine temperature step changes of HR-2 B but can be viewed in Figure 3.14. The coating losses for HR-2 were, similarly to HR-1, found to be lowest at 600°C. The




Figure 3.17: Coating mechanical losses as a function of frequency for the 69.5% Ti $\text{TiO}_2\text{:SiO}_2\text{/SiO}_2$ HR stack, after different stages of post-deposition heat treatment. The results for sample HR-2 A are plotted on the left, and for HR-2 B on the right.

loss levels at this common point heat treatment between the two HR-2 samples agree very well with each other, despite their different heat treatment histories and slightly different heat treatment procedures. This again suggests, as with HR-1, that the actual dwell temperature, rather than the sample history or heating/cooling rate dominates the resulting coating loss. In the next section the mean coating loss acquired from all the measured modes is also shown in Figure 3.19.

Sample HR-2 B was heat treated using a much finer range of temperature steps compared to the other three samples. In doing so it was noticed that its losses generally began to increase slightly from 500°C to 525°C, before subsequently reducing at 550°C and again at 600°C - this is perhaps best viewed later in Figure 3.19. This small peak in the loss ultimately may simply arise because any coating loss improvement between the 25°C temperature steps is minuscule compared to the variability of the measurement technique itself. After the 600°C heat treatment there is clear improvement witnessed for the coating loss on both samples.

Both HR-2 samples, upon reaching 600°C, irrespective of heat treatment procedure/history, showed signs of compressive stress failure, with numerous macroscopic blisters of a few millimetres in size forming in the coating. Most of these were localised to the outer 20 mm of the coating face where the quality of polishing would be expected to be worse, which may indicate some correlation between the two. The slower and controlled ramp rate did not have a noticeable effect on the defect formation nor the loss level: see Table 3.4.

Sample HR-2 A was also heat treated to 650°C where its coating was witnessed have cracked (showing tensile stress failure) after removal from the oven, with multiple large fractures across the coating face. Upon measuring its losses, they were found to have increased an order of magnitude compared to 600°C, and are shown in Figure 3.18.

Notably, unlike HR-1, no frequency dependent loss was seen to arise in the HR-2 B sample after the onset of cracking, and indeed much higher levels of loss were noted post damage in HR-2 than in HR-1 on the whole. The HR-2 coating stack is ~20% thicker than HR-1, and assuming similar material properties, as exemplified by the similar loss levels and close cation concentrations of the high refractive index layers in the stacks, HR-2 likely experiences greater levels of stress on the whole and this perhaps goes some way to explain the dichotomy in observed coating loss levels/trends at 650°C in both coatings. Of course if the defects arising in HR-1 happened to line up with the regions of high strain-energy storage for the mixed modes but not the clover modes, that may also explain differences in the measured trends.




Figure 3.18: A selection of coating mechanical losses as a function of frequency for HR-2 B, after different stages of post-deposition heat treatment. Specifically previously omitted 650°C data is now shown.

3.4.5.3 Titania-silica HR stack losses discussion

Both HR stack designs presented in this study achieve promising coating mechanical loss levels slightly lower than the current aLIGO/adV ETM coating stack design of $\sim 1.50 \times 10^{-4}$ after post-deposition heat treatment at 600°C . HR-1 achieves loss of $(1.32 \pm 0.16) \times 10^{-4}$, and HR-2 achieves loss of $(1.46 \pm 0.14) \times 10^{-4}$. These represent reductions from the as-deposited states of the coatings through heat treatment of around 70%, and show the importance of this post-deposition process for IBS coatings. The current GWD coatings, unlike these investigated stacks, are not simple QWL topologies, they in fact have slightly thicker layers than QWL of SiO_2 and slightly thinner than QWL of $\text{TiO}_2:\text{Ta}_2\text{O}_5$, which has the effect of decreasing their stack loss overall versus the QWL design. This was done primarily to provide additional suitable coating reflectivity for the arm length stabilisation system lasers of current detectors which operate with auxiliary 532 nm laser light [114], but provides coating thermal noise improvements as a secondary boon. As such, one might expect a slightly larger improvement of coating stacks incorporating this new material versus current GWD coatings, if they were also manufactured with similar layer thicknesses optimisation.

Figure 3.19 shows the average coating losses calculated from taking the mean of the losses from all measured resonant modes at a given heat treatment, with the error bars representing the standard deviation. This approach was adopted because of the observed by-and-large constant coating loss trend in frequency. The two exceptions to this are both for sample HR-1 B at 700°C where the losses increased and were not constant with frequency and at 650°C where just the mixed mode losses were observed to increase with respect to the clovers (see Figure 3.16). At both of these heat treatments the sample gained new defects. For 700°C an average is not plotted at all, and for 650°C only the clover mode losses were considered for the averaging shown in Figure 3.19.

Of course it is not just the mechanical loss that matters for coating thermal noise reduction, and when the improvements in loss are only slight, the refractive index of the stack layers will play a large role, as this determines the total thickness of the coating for a given reflectivity, directly impacting the thermal noise. In these HR topologies the aim was to test a replacement candidate for $\text{TiO}_2:\text{Ta}_2\text{O}_5$ in current room temperature detector




Figure 3.19: Coating mechanical losses for both TiO₂:SiO₂/SiO₂ stack designs as functions of post-deposition heat treatment temperature. Each point represents the average loss of 6–8 measured modes for a particular heat treatment step, with the standard deviation between them used as the error bars. Note that the points at 100 $^{\circ}\text{C}$ represent the as deposited average loss as the deposition temperature was near this.

designs. Though the losses of the TiO₂:SiO₂ based stacks were found to be slightly lower than the current TiO₂:Ta₂O₅ based stacks, the refractive index of TiO₂:Ta₂O₅ is larger than the TiO₂:SiO₂ investigated here. This means that the individual layers of TiO₂:SiO₂ will be thicker by Equation 3.18, and due to a lower refractive index contrast with SiO₂ additional bi-layers would be needed to achieve the same reflectivity of current detector coatings. Both of these factors increase the thickness of a stack incorporating the new high refractive index material, and would counter some of the improvement in the coating thermal noise from the lower mechanical loss. In the following section all these factors are considered and used to calculate the coating thermal noise level (shown on Figure 3.20).

Conversely, both compositions of TiO₂:SiO₂ investigated here have higher refractive index than another proposed high- n candidate TiO₂:GeO₂ with 44% Ti concentration, and therefore would yield a thinner overall stack than this other candidate material. As will be seen in Chapter 6, TiO₂:GeO₂ seems to have significantly lower loss at (its own optimal heat treatment of) 600 $^{\circ}\text{C}$, but current recipes yield worse optical losses than TiO₂:SiO₂. Thus, a favoured candidate is not yet clear.

The defects that emerged in these $\text{TiO}_2:\text{SiO}_2$ HR coatings following the heat treatments above 600°C appear to have very likely impacted the measured losses. There may be scope to alter the substrate, deposition, and post-deposition heat treatment parameters to suppress their manifestation, as has been achieved in the past for other IBS oxide coatings [208], which could result in higher achievable heat treatments and even lower losses. A subsequent study based on the promising mechanical losses of these $\text{TiO}_2:\text{SiO}_2$ based coatings was launched, into directly measuring their coating thermal noise on smaller scale super polished substrates in collaboration with colleagues at the Massachusetts Institute of Technology (MIT).

3.5 Directly measured coating thermal noise of the 69.5% Ti titania-silica HR stack

For the second HR coating run of these $\text{TiO}_2:\text{SiO}_2$ based coatings, with calculated $\sim 69.5\%$ Ti cation concentration in its high- n layers from the ERD-TOF measurements (see Figure 3.2), three 25 mm diameter, 6 mm thick super polished Corning 7980 fused SiO_2 disks were included. These samples were sent to colleagues at MIT to perform a complementary study in quantifying the coating thermal noise (CTN) these coatings would produce in the most sensitive frequency range of current GWDs at ~ 100 Hz, where it dominates the detector noise. The highly-polished substrate surface was essential for reducing optical scatter down to acceptable levels to enable the direct CTN measurement. The samples were only coated within an 8 mm diameter circular region in the centre and also had an anti-reflective coating of approximately 5 mm in diameter on the opposite face to mitigate backside reflections. A similar post-deposition heat treatment study to the one performed by the author for the mechanical loss samples was undertaken, with each sample being brought to and from target temperature with a ramp rate of $50^\circ\text{C}/\text{hr}$. In all cases, bar one, the samples were left to dwell at temperature for 10 hours. The CTN measured in this experiment could then be compared to calculations of the CTN made by inputting the measured coating losses shown in Figure 3.19, into Equation 2.41.

3.5.1 The directly measured CTN setup briefly explained

The CTN measurement setup at MIT [209] employs a Fabry-Pérot folded cavity, in which the folding mirror is the sample to be measured. To mitigate external noise coupling, the cavity is mounted on a vibrationally isolated platform and placed inside a vacuum chamber. Three optical modes co-resonate in the cavity - the TEM00, TEM02 and TEM20 modes. Each mode, due to their differences in shape, probes different areas of the sample's face, with the TEM02 and TEM20 overlapping only in a small region at the centre. Three laser beams are used: one locked to the TEM00 cavity mode, used to suppress cavity length noise and laser frequency noise which is common to all three modes. Two pick-off beams are frequency shifted and locked to the TEM02 and TEM20 resonant frequencies of the cavity, which emerge at slightly different frequencies due to differences in the horizontal and vertical radii of curvature of a real (imperfect) optic being measured.

The CTN, which is spatially independent between the modes, results in a fluctuation in this frequency difference between the TEM02 and TEM20 resonances. The fluctuations in the frequency difference between these two beams, emerging due to CTN (as well as other typically weaker noise sources that, unlike the CTN, mostly cancel on subtraction), is measured by interfering the beams and monitoring the resulting beat signal. This beat frequency, can then be converted into an equivalent cavity length change, and thus, the amplitude spectral density of the CTN N_{CTN} extracted. The operation of this setup is outlined in much more detail in a 2018 article by Gras and Evans [209].

Though this experiment quantifies the thermal noise sensed by TEM02 and TEM20 modes in a folded cavity across frequency space, more typically of interest for GWD is the thermal noise of the fundamental mode (TEM00) of a linear cavity N_{CTN}^{00} . Thus correction factors must be implemented to scale for the cavity geometry, mode shape, and for difference in laser beam size, which are discussed at length in [210]. Ultimately the final conversion from the measured CTN amplitude spectral density N_{CTN} in this setup with beam size w_S (typically ~ 50 microns) to CTN experienced in a linear cavity where a beam of size w_L locked to the TEM00 mode is used is [209]

$$N_{\text{CTN}}^{00}(f) = 0.616 \times \left(\frac{w_S}{w_L} \right) N_{\text{CTN}}(f). \quad (3.28)$$

Substituting in the beam size values typical to this experiment and for aLIGO (~ 6.2 cm), the prefactor terms simplify to $\sim 5 \times 10^{-4}$. With Equation 3.28, one can use the measurements from this folded cavity setup to estimate the level of CTN experienced in a GWD over a range of frequencies and fit trends in frequency to the measured data. The sensitivity band well encompasses the most sensitive frequency range of current room temperature GWDs where CTN is dominant, and through these measurements new HR coatings can be rapidly tested with their GWD CTN improvement prospects characterised. Samples with exact coatings used in aLIGO/adV currently were also measured in this setup, so direct comparisons can be made against these as a reference standard without this conversion to its expected level in a GWD cavity.

3.5.2 Direct CTN measurement results of 69.5% Ti HR stack through heat treatment

Of the three 25 mm diameter super polished samples coated with the HR-2 design across their central 8 mm, two were characterised extensively by collaborators at MIT in the experimental setup described in the last section. These two samples were both alternately heat treated in 50°C steps (100°C steps individually) from 250–700°C with fixed heating/cool rates of 50°C/hr, where each sample was held at temperature for 10 hours. Thereafter, influenced by the previous mechanical loss results, sample 1 was solely heat treated in increasing 25°C steps, until a CTN minimum point, followed by a defect formation point were found. After the optimum CTN post-deposition heat treatment temperature was found, sample 2 was heat treated once more to this temperature. The third sample was held in reserve, and all were eventually shipped back to the author for further optical and structural characterisation studies outlined in Chapters 4 and 5.

Figure 3.20 shows the directly measured CTN values for the HR-2 stack from the high-finesse folded cavity setup at MIT, alongside the calculated CTN from the average coating mechanical losses measured by the author and colleagues in Section 3.4.5, and shown in Figure 3.19. The loss-to-CTN conversions were made using the values acquired for the average HR-2 stack coating losses in Equation 2.41. The results from this calculation were used to compare results between the two techniques, and show very good agreement between the six common temperatures measured on both setups. The calculated CTN




Figure 3.20: Directly measured CTN at 100 Hz of the 69.5% Ti HR coating (HR-2), scaled to the aLIGO end test mass beam size (represented as diamonds and squares for samples 1 and 2), and CTN calculated from the HR coating mechanical loss shown by the red triangles, corresponding to the losses in Figure 3.19. The black dashed line represents the current aLIGO CTN at 100 Hz, with the blue dotted dashed line representing the upgrade goal.

values from the mechanical loss measurements, agree to within at worst 4% of the values measured directly at MIT. Of note, the slight upturn witnessed in the CTN calculated from the loss at 525°C is not captured in the direct CTN measurements, as 25°C steps were only adopted post 700°C, with the step size being 50°C prior to this. However the common temperature values measured either side agree well, and the gradient through temperature in the direct measurements significantly plateaus in this region when compared with, say, below 400°C and above 600°C. This supports the previously discussed hypothesis in Section 3.4.5.2 that the material's loss/quality factor is not improving greatly at these intermediate temperatures.

The major difference between the results from the two setups is that on the direct CTN setup the samples continued to improve after 600°C, and were in fact found to produce optimum CTN levels on samples free of defects after heat treatment at 850°C. In contrast the mechanical loss samples, with much larger geometry and coating area, on potentially not as well polished substrates, manifested blisters and cracks along the coating at tem-




Figure 3.21: Images of the two super polished $\text{TiO}_2:\text{SiO}_2$ HR coated samples used in the direct CTN measurements, after the final heat treatment of each. A crack emerged on sample 1 after 950°C , 10 hour heat treatment, whereas sample 2 remained immaculate.

peratures beyond 600°C , with accompanying increases in the measured loss. This could imply the stress experienced in the coating was different for both sets of samples, leading to the associated compressive/tensile stress-linked defects (see Section 3.4.5.1) occurring at lower temperatures in the loss samples.

Sample 1 measured in the high finesse cavity was heat treated up to 950°C where it experienced one macroscopic crack across the coating, and the coating appeared to visually become slightly cloudy, perhaps as a result of advanced crystallisation. Prior to this, the lowest CTN measured was at 850°C after 10 hours, and as such sample 2 was heat treated to this temperature but for a longer 100 hour dwell duration. Images of both samples after these final heat treatments can be seen in Figure 3.21.

The direct measurement of the CTN is a much more localised measurement (at the level of the beam spot size) than mechanical loss measurements, which sample the loss from the entire vibrating sample. As such any defects, imperfections, as well as the macroscopic cracks, are less likely to be impactful on the direct CTN measurement when they are far from the area of the coating sampled by the cavity laser beam. Indeed similar low levels of CTN have been seen with the folded cavity setup for similar samples with and without cracks present [211], whereas these clearly had an impact on the measured loss.

Sample 1 and sample 2 yield equivalent CTN values as-deposited of $1.2 \times 10^{-20} \frac{\text{m}}{\sqrt{\text{Hz}}}$, serving as a consistency check of the deposition process. After the 850°C, 10 hour heat treatment, the CTN levels were recorded at $5.0 \times 10^{-21} \frac{\text{m}}{\sqrt{\text{Hz}}}$, and $4.8 \times 10^{-21} \frac{\text{m}}{\sqrt{\text{Hz}}}$ after 100 hours. These are a reduction in the CTN levels of the as-deposited coating of 57% and 59% respectively as a result of heat treatment, and represent a reduction of the aLIGO CTN to 79% and 75% of its current value, respectively. This was achieved from a coating which dark field microscopy showed was free from defects, clearly highlighting the high promise of titania-silica mixes as a next generation GWD upgrade coating material candidate. The CTN calculated from the loss measurements, at 600°C just before the onset of defects (which is in good agreement with the directly measured CTN, as discussed), is equivalent to the aLIGO CTN level.

Despite the significant CTN improvement upon current ground-based GWD mirrors at 850°C, this coating by itself does not yet meet the aLIGO+/adV+ goal of a factor of two reduction in CTN [90], with it only reaching half way towards this. It is worth noting that this stack is four layers, or $\sim 10\%$, too thick as the high- n refractive index was slightly larger than predicted. One can see via Equation 2.41, that the amplitude spectral density (square-root of the power spectral density) of the CTN scales directly with the square-root of coating thickness. Having more layers should not greatly impact the extracted average coating mechanical loss, but will influence the overall CTN, meaning that $\sim 5\%$ lower CTN results could likely be achieved.

The 69.5% HR stack optimum heat treatment of this coating is close to (only 50°C below) the optimum heat treatment temperature of another promising very low mechanical loss, but higher optical absorption coating candidate - silicon nitride [115]. Should the optical loss (absorption and scatter) prove low for $\text{TiO}_2:\text{SiO}_2$, then pairing these materials with SiO_2 in a multimaterial design with a few bilayers of $\text{TiO}_2:\text{SiO}_2/\text{SiO}_2$ on top could be very promising for future GWDs. A similar argument could apply as well for pairing it in a multimaterial design with $\text{TiO}_2:\text{GeO}_2$ (discussed in Chapter 6) for heat treatment at around 600°C.

Material properties can vary significantly depending on the deposition process. As such, for a preliminary study, these coatings, which were not made in the same system that deposits the current state-of-the-art GWD optical coatings, yield very impressive results. Due to this deposition variability, it is quite possible as well that the stack mechanical loss/CTN is significantly influenced by the loss of the pure-SiO₂ layers. The loss of SiO₂ single layers produced by FiveNine optics has not been studied separately. The possibility of SiO₂ being the source of excess losses in certain coating designs has been suggested before, though is not thought to be the case for current aLIGO/adV coatings [82, 212]. Attempts to determine this through comparison of the stack mechanical losses with mechanical losses of single layer coatings of TiO₂:SiO₂ deposited onto thin disk substrates will now be discussed.

3.5.3 Estimating the loss of the pure silica layers and inferring an optimistic CTN

Similar to the loss samples measured for both HR stacks, a ~ 50 mm diameter, by ~ 1 mm thick disk was deposited with the 62.3% Ti TiO₂:SiO₂ single layer coating of 272 nm thickness referred to in Table 3.1, and characterised on GeNS following the same analysis procedure. The bulk of the measurements on this disk were undertaken by collaborators at Maastricht University, with the author carrying out all its heat treatments using the same oven at the University of Glasgow to keep procedural consistency with the stacks. A few repeat measurements in Glasgow to serve as cross checks between the two systems, or to simply acquire the data and avoid shipping delays, were also performed by the author, with all cross check measurements between the two systems yielding the same losses within error. The average coating loss results through heat treatment for this single layer are shown in Figure 3.22 alongside the previously characterised HR samples from Figure 3.19.

To within error (see Table 3.1) HR-1 and this single layer have the same cation concentration $\sim 63\%$. Interestingly, both of these, despite one being a single layer and one being an HR stack including pure-SiO₂ layers, have identical coating losses. A lower HR loss would be expected due to the loss of the SiO₂ layers in the stack generally being low [82, 213].




Figure 3.22: Coating mechanical losses for both $\text{TiO}_2:\text{SiO}_2/\text{SiO}_2$ stack designs as functions of post-deposition heat treatment temperature, as well as the 62.3% Ti $\text{TiO}_2:\text{SiO}_2$ single layer (black crosses). Each point represents the average loss of 6–8 measured modes for a particular heat treatment step, with the standard deviation between them used as the error bars.

Attempts were made to calculate the SiO_2 layer losses from the total HR stack losses using the single layer mixed material results. Following a method outlined by Penn et al. [172] the loss ϕ_2 of one of two materials comprising a stack can be inferred from the composite coating stack loss ϕ_C and the loss of the other material ϕ_1 using

$$\phi_C = \frac{Y_1 t_1}{Y_{\text{composite}} t_{\text{stack}}} \phi_1 + \frac{Y_2 t_2}{Y_{\text{composite}} t_{\text{stack}}} \phi_2 \quad (3.29)$$

$$\therefore \phi_2 = \left(\phi_C - \frac{Y_1 t_1}{Y_{\text{composite}} t_{\text{stack}}} \phi_1 \right) \left(\frac{Y_{\text{composite}} t_{\text{stack}}}{Y_2 t_2} \right). \quad (3.30)$$

For this case, the subscripts 1/2 represent the $\text{TiO}_2:\text{SiO}_2/\text{SiO}_2$ respectively. The single layer losses and HR-1 losses (with same high-n layer Ti content as the single layer) were compared at the common temperature points. Then, using Equation 3.30, estimates of what losses the pure SiO_2 layers in the stack would need to have to result together with the $\text{TiO}_2:\text{SiO}_2$ single layer losses, to produce the measured HR-1 losses, were made. This SiO_2 loss is shown by the blue triangles in Figure 3.23. The green and black points show values of IBS- SiO_2 coating mechanical losses from the literature. The grey, dashed lines




Figure 3.23: Mechanical loss estimates of the SiO₂ layers of the HR stack, obtained from the mixed-single layer and HR losses [214], and SiO₂ loss from the literature A [213] and B [165].

represent exponential least-square fits through all three data sets, indicating their trends. These trends were used to infer the improvement factors at 850°C for these IBS-SiO₂, as both were not actually measured at this temperature, but do have data at higher and lower heat treatment temperatures. The one hollow blue triangle assumes further improvement of our SiO₂ after heat treatment at 850°C, by the same factor observed for the green circle literature data A. [213], which was lower than the improvement factor of the black squares [165], as a conservative estimate of potential improvement. This is an improvement of a factor of $\times 1.85$ at 850°C, versus the SiO₂ loss estimate at 450°C.

This suggested the loss of the SiO₂ coating layers within the HR stacks was quite high compared to other IBS-SiO₂ coatings - observable in Figure 3.23. Excess loss in the HR coating, e.g. due to layer interface effects, may be an alternative explanation for the discrepancy of single layer and HR loss [149]. Other effects, like thickness related stress relaxation may also cause excess loss particularly on thin substrates [149], but this seemed unlikely to be the dominant source in this study with relatively thick samples showing consistent loss for different thicknesses as shown in Figures 3.15 - 3.19.

In order to infer the effect of potentially high SiO₂ loss, or other excess loss in the stack, on the CTN, the fit of the SiO₂ losses (blue triangles in Figure 3.23) was used to extract the corresponding high-*n* TiO₂:SiO₂ losses from the directly measured CTN at each heat treatment temperature via rearrangement of Equations 2.41–2.42. These high-*n* losses were then combined with the best SiO₂ loss numbers from the literature (the fit of the black squares in Figure 3.23) in order to produce an ‘optimistic’ potential CTN of a stack




Figure 3.24: Directly measured CTN at 100 Hz of the 69.5% Ti HR coating (HR-2), scaled to the aLIGO end test mass beam size (represented as diamonds and squares for samples 1 and 2), and CTN calculated from the HR coating mechanical loss shown by the red triangles, corresponding to the losses in Figure 3.19. The black dashed line represents the current aLIGO CTN at 100 Hz, with the blue dotted dashed line representing the upgrade goal. The green circles show projected CTN values calculated from individual $\text{TiO}_2:\text{SiO}_2$ and SiO_2 losses, assuming lower loss SiO_2 layers than those present are incorporated.

employing the best known loss SiO_2 , the results of which are shown in Figure 3.24. These literature values come from the same IBS chamber where the current GWD coatings are deposited [165]. The other parameters used for SiO_2 in this calculation were a Young's modulus of 73.2 GPa, a Poisson ratio of 0.11, and $n = 1.45$ which are parameters identical to those used by Granata et al. [165]. This analysis suggests that a maximum CTN reduction to $\sim 45\%$ of current aLIGO/AdV could be achieved, going beyond the aLIGO+ targeted improvement. This is the maximum possible improvement, if we could eliminate all excess loss by using the best possible silica instead of that used in these samples. However, improved SiO_2 could also possibly reduce the loss of the $\text{TiO}_2:\text{SiO}_2$ layers in the coating, improving performance even further.

3.6 Table of defect formation

Table 3.4: Overview: Crystallisation and damage of samples after heat treatment. The heat treatment temperature after which the defect was observed is shown, with the highest temperature after which the sample appeared undamaged given in brackets.

| Ti [%] | type | substrate | diam. | polish | T [°C] | observations |
|--------|------|----------------|-----------------|--------------------------------|------------------------|---|
| 58.5 | SL | thin disk | 2" | $\lambda/10, \leq 5\text{\AA}$ | 600 (550) | randomly spread damage observed under microscope |
| 63.2 | HR-1 | thick disk (A) | 3" | super pol. ^d | 600 (550) | ring shaped bubble-like defects in outer region of sample |
| 63.2 | HR-1 | thin disk (B) | 2" | $\lambda/10, \leq 5\text{\AA}$ | 700 (650) | cracks across the coating |
| 62.3 | SL | thin disk | 2" | $\lambda/10, \leq 5\text{\AA}$ | 550 (525) | shows some small cracks in the centre (from GeNS) |
| 69.5 | HR-2 | thick disk (A) | 3" | super pol. ^d | 600 ^c (550) | bubbles and cracks |
| 69.5 | HR-2 | thick disk (B) | 3" | super pol. ^d | 600 (550) | ring shaped bubble-like defects in outer region of sample |
| 69.5 | HR-2 | CTN (1) | 1" ^a | super pol. | 950 (925) ^b | cracks across the sample |
| 69.5 | HR-2 | CTN (2) | 1" ^a | super pol. | 850 | no defects |

^aCoated only in centre with diameter 8mm, while all other sample where coated up to the edge.

^bHeat treated with a controlled ramp down (all steps), while all other samples were allowed to cool down naturally (in a different oven).

^cLast heat treatment step with same procedure/oven as the CTN sample.

^d The 3" mechanical loss samples were super polished on one side, and many of them were suspended with that side contacting the GeNS lens prior to coating, which could have also lead to scratches of the super polished surface.

3.7 Summary of the titania-silica loss and coating thermal noise investigations

These studies into characterising the mechanical loss and CTN performance of $\text{TiO}_2:\text{SiO}_2$ based coatings for future GWDs yielded very promising results. The evolution of two different coating stacks with high index layers which consist of $(63.2\pm 1.7)\%$ and $(69.5\pm 1.3)\%$ TiO_2 was tracked through various stages of heat treatment. Optimum amplitude spectral density CTN levels of $4.8\times 10^{-21} \frac{\text{m}}{\sqrt{\text{Hz}}}$ were found for the latter composition after 850°C , 100 hr post-deposition heat treatment, which is around 75% of current aLIGO/adV levels, and was achieved on a sample free of defects.

Mechanical loss results for each stack, measured via GeNS, were found to be independent of frequency/mode shape for all samples free of defects, with each stack producing similar levels of loss (within uncertainties) through heat treatment. For the loss samples the optimum heat treatment seemed to be around 600°C , before the onset of macroscopic cracking in the coatings. This was not seen on the samples produced for direct CTN measurement, which had a much smaller coating area and were deposited on super polished substrates until after a 300°C higher anneal at 950°C . From the measured mechanical losses, the CTN was able to be calculated. All CTN values obtained at the same heat treatment temperatures between the mechanical loss study, where CTN was inferred, and the directly measured CTN study (prior to cracking in the loss samples) agreed to within no greater deviation than 4% of each other, showing great consistency between both techniques.

In the absence of single layers of SiO_2 made via the same process as the HR stacks, analysis was performed to estimate the losses of the SiO_2 layers of the stack. These results suggest that the SiO_2 layers may have significantly higher losses than expected, and therefore would increase the CTN. Applying the losses of SiO_2 single layer coatings made in the same ion beam sputtering system which deposits the current GWD coatings, suggests the potential of further reductions in CTN to up to 45% of the current GWD levels for these stacks, which would reach the targets for next generation upgrades.

Ultimately these investigations yielded highly promising results for a new material candidate for GWD mirror coatings, and could be highly impactful for the field of gravitational wave astronomy. These results will potentially shape the direction of research in the coming months and years for facilitating sensitivity improvements of detectors. Although with some (but not all) of the samples manifesting blisters and/or cracks at higher heat treatments there may yet be issues left to solve. This work shows the potential of $\text{TiO}_2:\text{SiO}_2$ as a new material which could deliver on the CTN goals for future upgrades, with significant improvements versus current detector coatings. Indeed with the mechanical loss and CTN of these materials characterised, subsequent investigations were launched by the author into characterising their optical losses and their structure, in order to further verify their validity - explored in the next two chapters. The culmination of all these studies recently led to a first author Physical Review Letters publication in 2023 [171].

Chapter 4

Optical loss studies of titania-silica HR coatings

4.1 Introduction

In Chapter 3, the mechanical losses of ion beam sputtered (IBS) titania-silica mix coatings and their potential to meet the coating thermal noise (CTN) requirements of next generation gravitational wave detectors was presented. As well as the thermal noise, there are very stringent optical loss tolerances for future detectors. These tolerances are at the parts per million (ppm) level, with no greater than 0.5 ppm optical absorption, and no greater than 10 ppm optical scatter permitted [124]. This chapter describes investigations of the optical absorption and scattering of these coatings, with these values quantified through varying key stages of post deposition heat treatment.

The culmination of the studies presented here, as well as those in Chapters 3 and 5, resulted in a first author Physical Review Letters publication in 2023 [171].

4.2 Optical absorption studies of the titania-silica HR coatings

The optical absorption investigations were carried out primarily via photothermal common-path interferometry. These studies were conducted by the author and fellow researchers at the University of Glasgow, and by the author and Dr Liyuan Zhang while undertaking a research placement at the California Institute of Technology (Caltech).

4.2.1 Photothermal common-path interferometry

The Photothermal common-path interferometry (PCI) technique [215] involves crossing the waists of two laser beams of different wavelengths on the surface of a sample. The first beam, dubbed the “pump” beam, is at the wavelength of interest for characterising the sample’s optical absorption. This pump beam is of significantly higher-power, and with comparatively smaller waist size (typically $\times 2\text{--}4$ smaller [216]), than the second laser, dubbed the “probe” beam - making the pump beam much higher intensity than the probe. The pump beam is optically chopped at a specific frequency to provide periodic illumination, and therefore, heating of the sample, thus producing a periodic thermal lensing effect [215, 217]. The thermal lens is generated by the refractive index of the sample changing due to heating from absorbed laser light. The portion of the much larger probe beam that passes through the (pump-induced) thermal lens, will interfere with the part of the probe beam which is unaffected by the thermal lens. The resulting interference pattern causes a signal S proportional to the absorbed pump power. By scaling this signal to the signal produced from another sample of known absorption, dubbed the “calibration sample” S_{cal} , the signal from the sample of interest can be converted into an absorption value α . This can be expressed as [79]:

$$\alpha = \gamma \frac{S}{S_{\text{cal}}} \frac{P_{\text{cal}} \Omega_{\text{cal}}}{P \Omega} \alpha_{\text{cal}}, \quad (4.1)$$

where α_{cal} is the known optical absorption of the calibration sample, P and P_{cal} are the powers of the pump beam during the acquisition of the signal from the sample of interest and the calibration sample respectively. Similarly Ω and Ω_{cal} are the corresponding probe beam powers, and γ is known as the substrate-correction factor.

The properties of the thermal lens depend on the properties of the substrate upon which it is induced. Therefore, when scaling an absorption signal from a sample of a different substrate material to the calibration sample, this correction factor needs to be applied [79]. All PCI absorption measurements presented in this thesis are for coatings deposited on fused-SiO₂ substrates, scaled to signals from fused-SiO₂ calibration samples, so for these studies $\gamma \approx 1$.




Figure 4.1: Schematic diagram of the PCI system at the University of Glasgow. A sample to be measured is mounted on an x-y-z translation stage. The photodiode (PD) resides on an imaging stage that can be moved in the z-direction.

A schematic diagram of the PCI setup at the University of Glasgow is shown in Figure 4.1. The probe signal is monitored with a photodiode, and modulation of the probe signal is monitored with the lock-in amplifier. To measure the absorption of the coating, the waists of the probe and pump beams must be crossed on the surface of the sample. In practice it is easiest to keep the waist crossing point fixed, mount a sample on an x-y-z translation stage and pass it through the crossing point (in the z-direction) while monitoring the signals on the photodiode and power meter. Once the coated surface is at the correct z-position, if desired, the sample can then be moved in x and y (adjusting the z position if necessary to account for tilt) to acquire a map of the absorption across the sample surface.

The time delay in the response of the sample to the heating from the pump laser, is linked to the thermal diffusivity of the sample [217]. This delay is measured as a phase for a given absorption signal. In the case of a thin coating, heat conducts rapidly into the substrate material and, as such, the majority of the signal is produced in the substrate. The phase of the absorption signals, for all samples measured in this project, was therefore expected to correspond to the thermal diffusivity of fused-SiO₂ (with peak signal at the coating/substrate interface).

The photodiode monitoring the probe signal is placed on an imaging stage (a fine-control z translation stage). The imaging stage operates on the principle of a Galilean telescope. Its function is to maximise the area of measurement for the thermal lens-induced perturbed probe beam and its distance to the beam crossing point, so that a common photodiode

can be used to monitor the probe signal for any sample installed in the setup [79]. The imaging stage also enables compensation for measuring samples of different thicknesses - by adjusting the stage in the z-direction one can compensate to keep the probe signal maximised, and thus maintain high measurement sensitivity.

4.2.2 Absorption of the 69.5% titania-silica HR coating obtained on two PCI systems

The optical absorption of the HR-2 coating stack with $\text{SiO}_2 / \text{TiO}_2:\text{SiO}_2$ (69.5% Ti composition) was studied extensively by the author and fellow collaborators using two independent PCI systems. This coating was previously found to show lower thermal noise than current GWD coatings – see Figure 3.24.

The PCI setup at the University of Glasgow utilised a $70\ \mu\text{m}$ waist diameter 1064 nm pump beam, and $230\ \mu\text{m}$ waist diameter 1620 nm probe beam configuration, chopped at 408 Hz with no more than 2.4 W of laser power incident on a sample being measured. The Caltech PCI setup utilised a $60\ \mu\text{m}$ waist diameter 1064 nm pump beam, and $450\ \mu\text{m}$ waist diameter 633 nm probe beam configuration, chopped at 1000 Hz with no more than 9 W of laser power incident on a sample being measured. In measuring samples with identical histories, and in one case the exact same sample on both setups, the author hoped to confirm the optical absorption values of the samples, the absorption uniformity across distinct but nominally identical samples coated at the same time, and also verify the consistency between the two independent PCI systems which were used.

4.2.2.1 Discussions of the equivalence of different sampling techniques and samples

A variety of coatings deposited on Corning 7979 and 7980 fused silica substrates [218] were measured in this absorption study. The samples were of varying geometry and substrate surface quality. When investigating effects of heat treatment, most samples were directly heat treated at the desired temperature, whilst others had undergone multiple stages of heat treatment at successively increasing temperatures, as they were the same samples used in the mechanical loss and CTN studies in Chapter 3. Details of the heat treatment histories and absorption values obtained for each sample are outlined in Table 4.1.

When dealing with low absorption samples, typically below a few ppm absorption, measurements can become quite time consuming and challenging as it can be difficult to resolve such low signals. Therefore great care must be taken to align, calibrate and shield the setup from factors such as ambient noise.

As the measurements were carried out with 1064 nm pump beams, and the coating stacks permitted only ~ 10 ppm transmission at 1064 nm, the substrates only received $\sim 0.001\%$ of the total pump light power incident on the coating top surface. The substrates, though different in terms of geometry and polishing, were all made of low absorbing fused silica, and as virtually no light would be reaching them during measurement in order to be absorbed, the influence of any substrate absorption in these measurements can be considered negligible. Indeed the absorption values are dominated for the most part by only the top few bi-layers of the coating stack [219].

Three sampling techniques were employed to characterize the various coated sample surfaces via PCI: multiple point measurements, line scans, and surface maps, with the technique chosen dictated primarily by the time constraints for a given measurement. The sampling technique variations should have no great impact on any particular absorption values obtained for a coating of high homogeneity. In point measurement sampling, a certain number of randomly distributed absorption measurements across the sample are acquired, with the mean and standard deviation of these taken as the final absorption value and spread. Once a sufficient quantity of point measurements are acquired, the mean

value from these should nominally converge on the values of a large area surface map. The minimum number of point measurements for a given sample in this study was five, though if the absorption values acquired varied greatly, many more point measurements would be taken. For line and map scans a large region of the surface is systematically sampled at equally-spaced intervals, equivalent to many tens, hundreds, or thousands of point measurements, with mean and standard deviation of the data-set again calculated. In scan measurements, the sample tilt has to be very precisely accounted for, so as to ensure the beam waists are crossing at the coating surface for all the measurements. Map scans can take orders of magnitude longer to complete than point sampling measurements, and through these the entire surface can be characterised.

To demonstrate the equivalence of the different sampling techniques, two independent measurements for the optimum CTN 850°C, 100 hour heat treated sample are shown in Figures 4.2 and 4.3 as examples. Five point scan measurements were taken with the Glasgow PCI, and a 4 mm diameter circular absorption map about the optic's centre was acquired by the author using the Caltech PCI (containing 1257 sampled points). Taken together they represent the largest contrast in sampled points in the entire study, and both measurements were used to probe the same value on the same sample, so will serve well to illustrate the equivalence of the sampling techniques and the two setups themselves. Both measurements gave the same mean absorption value of 0.82 ppm with a standard deviation of 0.11 ppm for the point scan measurements on the Glasgow PCI, and 0.16 ppm for the Caltech PCI map.




Figure 4.2: Absorption measurements of the HR-2 850°C, 100 hour heat treated sample from the Glasgow PCI. The top figure shows z-scan measurements for five points of the sample taken at the University of Glasgow, and the bottom their corresponding phase values – the central peak in the absorption signal corresponds to the absorption measured at the coating surface. This maximum absorption signal is achieved as the sample’s coated surface passes (in the z-direction) through the cross point of the pump and probe beam waists. The incident pump beam power during these measurements was 2.4 W and the mean absorption value from these five point measurements is 0.82 ppm.




Figure 4.3: Absorption map of the central 4 mm circular diameter of the HR-2 850°C, 100 hour heat treated sample’s surface, taken by the author on the Caltech PCI – equivalent to 1257 point measurements. On the left figure the colour scale is capped at 1.5 ppm, whereas on the right the same data has it capped at 5 ppm, to highlight the cluster of largest/lowest absorption points (inside red circle), later determined to correspond to a location where surface contaminants were present. The incident pump beam power during this measurement was 9 W, with a mean absorption value obtained of 0.82 ppm.

Any surface absorption measurement (central peak value) accompanied by a surface phase value diverging greater than 15% from the nominal phase value was discarded as a failed measurement as it was assumed to result from misalignment of the samples tilt, and the acquired value was not originating from the coating surface. None of the point scans shown here are considered failed measurements. The same phase stringency is applied to each point of the map.

For the map data, the four highest absorption points, i.e. the only ones with greater than 5 ppm absorption across the entire surface, as well as the one point of lowest absorption, all appear in a cluster within a small 500 micron-wide region (red circle in Figure 4.3). Subsequent visual inspection revealed a small contaminated region on the optic after this measurement of approximately this size in this region of the sample, on an otherwise immaculate appearing surface. This was later removed with further cleaning, thus, these five points were also discarded as not true measurements of the coating absorption, and considered significantly influenced by the contamination on top of the surface. Unfortunately, there was not time to measure it again post-cleaning at Caltech, though this should not significantly affect the results. Once this was all accounted for, the remaining map points were considered valid, and the mean absorption of the map (now consisting of 1252 remaining points), as well as the mean of all five Glasgow PCI point scans could be calculated, along with the corresponding standard deviations.

Tests have been performed on the Glasgow PCI by the author and by others [164, 182], for other coatings, which showed the equivalence of performing multiple point scan absorption measurements versus a surface map measurement of samples with reasonably uniform absorption. The two independent absorption measurements of this same sample shown here, taken on two different PCI systems not only lie within the error bounds of each other, but also yield the exact same mean value of 0.82 ppm, showing strong agreement irrespective of the sampling technique employed, and irrespective of which system acquired the measurements. These results also highlight the independence of the final absorption result obtained with respect to incident pump intensity, given that the incident Caltech

PCI pump intensity was $\sim \times 5$ higher than in the Glasgow PCI measurements. With confidence in the agreement between the independently obtained values for the same sample on different setups, the absorption evolution of this coating through heat treatment will now be discussed.

4.2.2.2 Absorption of the 69.5% titania-silica HR coating through heat treatment

Figure 4.4 shows the measured optical absorption at 1064 nm for the HR-2 coating through different stages of heat treatment. A general downward trend in absorption with heat treatment can be seen until approximately 550°C where, thereafter, the absorption trend flattens approaching values close to and below 1 ppm. The error bars in all cases result from the variation of absorption measured across the coating and are the standard deviations between the multiple measurements at each heat treatment. Two results were acquired for two different as-deposited samples, and they lay within the error range of each other, highlighting the absorption homogeneity of different samples coated in the same run, residing in different parts of the coating chamber during deposition.

Between 600°C - 800°C no absorption measurements were acquired via PCI, so it is possible that the coating absorption significantly decreases further in this unmeasured region and then rises again, though the absorption estimates from cavity measurements shown later in Figure 4.6 suggest this is unlikely as they are consistent with a relatively flat absorption trend at these temperatures. All heat treatments lasted 10 hours, except for the 850°C step, which lasted for 100 hours – details of all heat treatment histories can be seen in Table 4.1. The value of 0.82 ppm, measured at 850°C using both the Glasgow and Caltech PCI is approximately an 86% reduction from its as-deposited value, and was achieved from a sample free of defects. This was the same sample that achieved CTN levels of 75% of current GWD coatings in the same heat treatment state.

Of note, at 600°C two samples were measured, yielding (0.84 ± 0.09) ppm and (1.69 ± 0.08) ppm respectively. The first result was from a 1-inch diameter, 1 mm thick sample, heat treated directly to 600°C with a 50 K/hr heating and cooling rate. The second result was from a 3-inch diameter, 2.5 mm thick sample, heat treated progressively to 600°C in a number of




Figure 4.4: Optical absorption values at 1064 nm of the 69.5% Ti high-n layer HR stack coating through heat treatment obtained via PCI. The heat treatment duration at 850°C was 100 hours, while the heat treatment duration at all other temperatures was 10 hours. The one point indicated with a red \times at 600°C are from measurements of a sample where many macroscopic defects formed during heat treatment. The markers and error bars represent the mean and standard deviation of all measurements on the surface at each heat treatment.

steps at different temperatures, with heating rates of 100 K per hour and with free (initially rapid) cooling to room temperature. Both values were obtained via multiple point measurements. The latter sample exhibited much macroscopic blistering, and cracking across the coating surface after heat treatment at 600°C, whereas the former continued to appear immaculate under visual inspection – though it was later learned after optical microscope inspection that it in fact had one $\sim 100 \mu\text{m}$ diameter defect. One can see that the damaged sample with many defects had around double the absorption value of its counterpart. Further, the immaculate-to-eye sample generally seems to follow the trend of absorption set by the samples at the prior heat treatments where again no defects were observed. For the sample with many defects, many points were taken both at/near and away from macroscopic defect regions, yet the variability witnessed for that sample is low, meaning the absorption is consistent across that sample and also distinct from the other sample at the same target temperature. Thus, it seems the heat treatment heating/cooling




Figure 4.5: Optical microscope images of the HR-2 sample annealed at 600°C exhibiting many macroscopic defects (a)-(e), and the sample also annealed at 600°C which did not (f)-(g).

(a)-(b) are respectively bright-field images of two of its many blisters, and a crack.

(c)-(e) show blisters under dark-field view from different parts of the surface.

(f) an image of the sample heat treated to 600°C , which showed no visible-to-eye defects.

(g) a bright field microscope image of (f)'s one small and dim $\sim 100\ \mu\text{m}$ diameter defect.

rate did not just impact the coating structural integrity (manifestation of defects), but it also impacted the absorption globally, even away from the regions where defects were observed to manifest. An assortment of microscope images for these two different 600°C heat treated samples are displayed in Figure 4.5.

Bright- and dark- field images of the samples were taken to observe the degree of scattering from the different defect regions. On the sample which manifested many blisters and cracks at 600°C heat treatment, the perimeters of each blister exhibited high light scatter when observed under dark-field (Figure 4.5(c)–(e)), as did the cracks. Figure 4.5(e) shows a dark-field image of a blister which formed at the very edge of the coating, accompanied by cracks which span across the length of the sample. The centres of some blisters appeared to contain within them a much smaller bright circular region (see Figure 4.5(c)). These appeared at different depths from the blister, as to resolve many of them clearly, the microscope focal height had to be adjusted, to a point where the rest of the blister was out of focus. This perhaps suggests these are points where the blister has burst, they could also be nexus points of coating/surface non-uniformity /contamination lower down in the sample that manifested the blister.

Interestingly, on the other sample, which appeared defect-free to the eye, one similarly sized feature to these small spherical regions formed on heat treatment (Figure 4.5(g)). However, this feature is much dimmer than any feature seen on the other sample under dark-field illumination. Subsequent heat treatment of this sample could yield interesting results and insights into its relation (or lack there-of) to defects that emerge.

4.2.2.3 Summary of the 69.5% titania-silica HR PCI results

This coating possessing a sub 1 ppm optical absorption value at the intended wavelength and optimal heat-treatment temperature for CTN is a significant and highly promising result, particularly from a preliminary coating run with no bespoke absorption optimisation parameters implemented for this HR stack design. A high level of agreement between each PCI system can be seen from the results. Independent of experimental factors such as incident pump light intensity or sampling technique, the lowest mean absorption value was found to occur at this optimum CTN heat treatment, with both systems confirming 0.82 ppm mean absorption with standard deviation of 0.16 ppm on the map measurement ($\sim 19.5\%$ variability). Current GWD ETM coatings have absorption values obtained from surface maps of (0.27 ± 0.07) ppm [124], i.e. a variability of $\sim 26\%$. Comparing the values acquired for HR-2 and these, we see that though the overall absorption value is higher, the variability is on-par with current high-quality GWD coatings. Though, care must be taken in this comparison as the absorption map size for the current GWD ETM in the literature was for an order of magnitude larger diameter [124].

Though not quite reaching the desired sub 0.5 ppm target for GWD upgrades [124], sub 1 ppm is still a significant achievement for this preliminary investigation, and there is likely to be significant scope for improvement of the absorption of this material. The actual system which the current GWD mirror coating manufacturer, Laboratoire des Matériaux Avancés (LMA), uses to coat the test masses has a track record of producing ultra-low absorption IBS coatings [124]. So identical coatings made in this system, might be capable of achieving the desired absorption.

4.2.3 Comparison with cavity absorption measurements

The optical absorption of the HR-2 coating was also estimated in the folded-cavity setup used for the direct CTN measurements described in Section 3.5.1. This was done by collaborators at MIT, by recording the resonance frequency of the high-finesse cavity as a function of circulating power with the sample installed. By comparing this against values for a coating of a known optical absorption, the absorption of these coatings could then be inferred. Each measurement from the folded cavity comes from a single point on the sample of diameter no greater than 50 μm near the centre of the coating surface, resulting in the absorption values shown in Figure 4.6 as the green + and blue \times markers. In this study the two samples investigated in the cavity setup were heat treated at progressively increasing temperatures until (1) an optimal CTN temperature was found, and/or (2) defects emerged on the given sample. The absorption can again be seen to reduce with increasing heat treatment temperature to a minimum of approximately (0.9 ± 0.4) ppm at 650°C , and then remains fairly constant until it starts to increase slightly at 800°C . Most of the heat treatments were carried out for 10 hours. However, at 850°C , cavity sample 1 was heat treated for 10 hours, and sample 2 for 100 hours. No significant difference in absorption between these was found. The uncertainty on the cavity measurements is quite high and generally larger than the PCI results in this study, but the method still serves to give a good approximation of the absorption trend in this coating.

One drawback with these cavity measurements was that due to the setup configuration at the time, neither the sample, nor the laser spot could be moved and as such these absorption values are based only on a single point measurement at each heat treatment. If the absorption varied significantly from point-to-point on the coating, this could give a value which is not representative of the coating on the whole. Hence one would be slightly less confident in these values compared to the previously discussed PCI measurements where the sample could be moved allowing any point on the surface to be measured, and an average of these measurements found.




Figure 4.6: Optical absorption values at 1064 nm of the 69.5% Ti high-n layer HR stack coating through heat treatment obtained via PCI and in situ CTN cavity measurement estimates. For cavity sample 2 and PCI measurements at 850°C, the heat treatment was 100 hours, while all other measurements were for 10 hour annealed samples. The one point indicated with a red \times at 950°C is done so to highlight that the sample cracked.

Figure 4.6 shows a general agreement between absorption measured with the PCI and with the cavity technique. Notably, at 850°C the PCI data comes from ‘sample 2’, i.e. the same one as measured in the cavity set up, and in the same state, and while both setups yield low absorption values, the values are not within error of each other (though the extremes of the ranges being only ~ 0.1 ppm apart). The cavity absorption is over double the 0.82 ppm measured on both PCI setups. From the PCI map, after removing the 5 ‘bad’ points from the contaminated region (see Figure 4.3) only 4 of the remaining 1252 points exhibited absorption values above 1.5 ppm ($< 1\%$). This then perhaps alludes to potential contamination affecting the cavity measurements or indeed limitations of the scaling technique, as both independent PCI systems arrived at the same consistent (around 50% lower) absorption value. This could stem from the singular point measurement in the cavity setup indeed being slightly higher absorption than the nominal representative surface mean, and/or due to the fact that the sample went through many rounds of

cleaning by the same operator before the PCI measurements in both Glasgow and Caltech (which themselves agree with each other). However, the discrepancy between the cavity measurement and PCI on the whole through heat treatment is still rather small, with the range of values for each temperature step agreeing for the most part.

The CTN cavity system at MIT has since been upgraded, such that measurements can now also be made at multiple points across the sample face instead of just one fixed point in the centre. Now the setup can probe CTN and absorption homogeneity. The former is more significant on the whole, but the latter would hopefully allow for both the values and the uncertainty on the cavity absorption measurement to be refined, which could lead to a more in depth comparative study for future samples.

4.2.4 Comparison of two titania-silica HR coating compositions

A brief study was carried out into the absorption of the HR stack with high-index layers of a lower Ti cation content (63.2%). Figure 4.7 shows the PCI absorption results of a 63.2% HR coating (previously dubbed HR-1) compared to the 69.5% HR coating (HR-2). The 63.2% HR results are generally higher than for the 69.5% HR coating. It is well known that coatings made via IBS (and other techniques) can be deposited oxygen-poor, and gradually become more stoichiometric as they are heat treated in air, taking in more oxygen from the environment and filling their oxygen deficiencies [220, 221]. The negative effects of oxygen deficiency on absorption and other material parameters, and improvements yielded from heat treatment have been previously well studied for oxide films [220, 222]. Indeed it may be the dominant absorption improvement factor at low temperatures, more-so than factors such as material relaxation [222].

In the as-deposited state, the 63.2% coating's higher absorption could result from it being deposited in a farther from stoichiometric state than the 69.5% sample and, thus, perhaps possessing greater oxygen deficiency. This would support results previously found by Bassiri et al. that greater Ti concentration may mitigate oxygen loss in as-deposited coatings [223]. We can see that the most drastic absorption improvement occurs after the first heat treatment at 450°C for 10 hours, and here both coating absorption values now lie within their errors resulting from the variation of absorption across the coatings. This




Figure 4.7: Optical absorption values at 1064 nm of the 69.5% and 63.2%Ti high-n layer HR stack coatings through heat treatment obtained via PCI. For 69.5% at 850°C, the heat treatment was 100 hours, while all other measurements were for 10 hour annealed samples. The points indicated with a red × represent that those samples had cracked and/or blistered. The markers and error bars represent the mean and standard deviation of all measurements on the surface at each heat treatment.

perhaps alludes to absorption improvements as a result of oxygenation, and that the two materials of slightly different cation concentration are very similar optically once their vacancies are filled. Looking at all heat treated and defect free 63.2% results beyond this, it can be seen that though around a similar level, the values either lie within the error range or are higher than the 69.5% results. Interestingly, for heat-treatments which resulted in defect formation (cracking / blistering), the 63.2% coating's absorption did not increase significantly, like was seen for the 69.5% coating – staying around the ~ 1 ppm level. Beyond this temperature, at 650°C the trend also seems to remain relatively flat, which is similar to the defect free 69.5% HR-2 samples.

Overall the data presented here is not conclusive to determine which of the two coatings has the lowest comparative absorption. The 69.5% coating was studied in more detail, due to its demonstrated promising CTN values. All 63.2% samples were heat treated at a faster rate in an oven with no temperature controlled ramp down, which seem to have been factors for defect suppression and average absorption values in the 69.5% coatings, as can be seen in Table 4.1. In order to draw a more definitive comparison between the two

Table 4.1: Optical absorption values obtained via PCI for the titania-silica mix high-n layer HR sack samples through heat treatment. The corresponding heat treatment histories and ramp up/cooling procedures are also shown. For all bar the 850°C sample the heat treatment dwell duration was 10 hours.

| optical absorption values through heat treatment [ppm] | | | |
|--|--|---|---|
| | 69.5% HR | | 63.2% HR |
| as-deposited. | 5.90±0.12 | | 8.20±1.80 |
| heat treatment | 50°C/hr ramp, controlled cool | 100°C/hr ramp, uncontrolled cool | 100°C/hr ramp, uncontrolled cool |
| 450°C (direct) | | 2.07±1.00 | 1.20±0.20 |
| 500°C (progressive) | | 1.36±0.04 | |
| 525°C (progressive) | | 1.08±0.03 | |
| 550°C (direct) | | 0.92±0.24 | 1.90±0.40 |
| 575°C (direct) | | 1.00±0.15 | 1.30±0.70 |
| 600°C (direct) | 0.84±0.09 | | |
| 600°C (progressive) | | 1.69±0.08* | 1.14±0.12* |
| 650°C (direct) | | | 1.15±0.58* |
| 850°C, 100 hr (progressive) | 0.82±0.11 | | |

* are used to indicate values obtained from the three samples with macroscopic defects.

stacks with different high-n layer concentrations, these results could be expanded upon by heat treating samples of the 63.2% coating to key temperatures with controlled cooling and at a slower ramp rate, mimicking the previous 69.5% results, as well as performing further measurements at higher temperatures, such as 850°C and select values in the, as of yet unmeasured, 200°C region below.

4.2.5 Absorption summary and further discussions

The coating stack incorporating the 69.5% Ti material achieved its lowest absorption value of (0.82±0.11) ppm at the heat treatment which is optimum for coating thermal noise, which was verified on three independent setups at Glasgow, Caltech, and MIT. Both the Glasgow and Caltech PCI setups were observed to agree well with each other at the low/high extremes of the absorption range. The previous limitations of the folded cavity absorption inference technique have been discussed, however after these studies the collaborators at MIT upgraded their sample holder stage, such that measurements can now be made at multiple points across the sample face instead of just one fixed point in the centre. Re-measuring some of these samples with these recent upgrades now in place in the folded cavity setup could prove valuable for obtaining a more global view of both the CTN and absorption variations across HR coated samples. The former is more

interesting on the whole, but the latter would lead to absorption surface maps from this different technique, which could be directly compared with the PCI maps, and allow for greater diagnosis of potential absorption discrepancies. The 63.2% Ti HR stack results were confirmed to be a similar level of absorption to their 69.5% counterparts, though a far less in-depth study was conducted for this coating. These sub-ppm absorption results for this initial study of HR stacks with titania-silica mix high-n quarter wave layers, and pure silica low-n quarter wave layers, are overall extremely promising as future gravitational wave detector coating candidates, given that they come close to the 0.5 ppm requirement for future detectors.

Following the initial measurements reported here, further titania-silica mix/silica HR coatings have been deposited by the current GWD coating vendor LMA, and fellow researchers continue to actively expand the development of this material [224]. The prospects of producing an HR coating with absorption losses below the 0.5 ppm target seem good.

Another potential avenue of further research for this material could be in incorporating it as part of a multi-material design with silica, and higher-absorption silicon nitride [115] in the lower layers. The optimum heat treatment temperatures of silicon nitride and the titania-silica studied here, appear to only be 50°C apart, with silicon nitride seemingly possessing even lower mechanical loss, and higher refractive index than titania-silica [165]. This means that depositing layers of both in a multi-material design could potentially be quite promising for producing an HR coating with low CTN levels, whilst also having low absorption. Having only a few layers of titania-silica on top of a stack comprising mostly of silicon nitride and silica, could also potentially help mitigate the defect formation witnessed on some of the samples in this study.

4.3 Optical scattering studies of the 69.5% titania-silica HR coating

Mitigation of the scattered light from GWD optics is another crucial factor for the detection of weak and distant astronomical signals. Noise due to scattered light has been a frequent disturbance in Advanced LIGO and Virgo, hindering the detection of gravitational waves [225]. It causes a direct reduction in the quality of quantum squeezed states of the interferometers [226]. Scattering from the test masses also results in a direct reduction of power in the Fabry-Pérot arm cavities, with the scattered light from the primary test masses potentially hitting walls and baffles and re-entering the interferometer beam, introducing additional noise [227, 228]. The average scatter from a collection of many aLIGO test masses was found to be around (9.5 ± 2) ppm [123].

This section describes the characterisation of the optical scatter of the titania-silica HR coatings carried out during a research placement at Caltech and California State University Fullerton (CSUF).

4.3.1 Factors which give rise to optical scattering

When light enters a new medium a number of interactions will occur. Some light will be reflected at the same angle to the normal as the incident ray, but on the opposite side of the normal. Some light will also be transmitted, refracting as it travels through the second medium, as described by Snell's Law [229]. Some of the light energy will be absorbed upon interaction, generating heat and thermal distortion in the second medium. Some light, however, will also be scattered at angles different from the specular reflection angle, the vast majority of which will be Rayleigh scattered as is described more in Section 5.1.1.

Light scattering from a layer of material can originate either from local variations of the properties in the interior of the layer, or from the irregularities of the physical surfaces or interfaces that form its boundaries. These two sources can be classed as volume scattering and surface scattering respectively, and it is generally accepted that the latter usually dominates in multilayer dielectric coatings [230]. Volume scattering is solely a function of the characteristics of the ideal multilayer and the observation conditions.

Surface scattering depends on the roughness of the various surfaces. For the production of high-performance low-loss dielectric multilayer systems, ion based deposition techniques such as ion plating deposition (IP), ion assisted deposition (IAD), and ion beam sputtering deposition (IBS) have been historically favoured [119]. Coatings deposited with these techniques show almost no structure of their own, and so topographic features of the interfaces strongly depend on the topographic structure of the substrate on which they were coated. Smooth well polished substrate surfaces are therefore essential for the production of multilayer systems with small surface roughness, and thus low optical scatter losses [119].

It is well known that the roughness of boundary surfaces is always responsible for at least part of the scattering of the incident light [230]. In the past half-century a growing number of applications have demanded that multilayer optics reach the theoretical ideal performance [230]. As an example, for current aLIGO and for its next iteration aLIGO+, the design requirements specify the scatter from a single test mass should be no more than 10 ppm [124].

Light scattered from interfaces within an HR stack does not emerge in the direction of specular reflection. As a result, unusual phase relationships and thereby scattered light patterns can also occur [231]. Regardless of this complexity, if one simply quantifies all light power scattered from a surface to all angles other than the incident angle of the illuminating light source, one can quantify its total optical scatter (for a given wavelength).

4.3.2 Using an integrating sphere to measure optical scatter

Total integrated scatter (TIS) is a measure of the total amount of light scattered by the sample in all directions [119]. This quantity can be recorded through the use of an integrating sphere and carefully aligned optics. An integrating sphere essentially spatially integrates radiant flux [232, 233], and has ports which can be opened to allow light to enter and escape. In essence, integrating spheres are hollow spheres with interior walls which form a near-perfect Lambertian surface [234], making any luminance on the sphere interior wall constant. In other words, the inner-wall can be considered as a perfectly diffusing and highly reflecting surface [235]. These properties as well as the spherical shape, lead




Figure 4.8: Images of an integrating sphere before (left) and after (right) its interior is illuminated with a collimated laser light source, showing the homogeneous diffuse distribution of the light once bounced off the inside walls.

to a homogeneous distribution of any light flux entering the sphere over its interior walls, via multiple diffuse reflections off the walls. As such the brightness of any spot on the wall is proportional to the reflected flux [119, 232]. In order to achieve this, integrating spheres can incorporate a highly reflective and spectrally-flat inner-wall material such as Spectralon® [236, 237, 238]. The uniform and diffuse light distribution produced by an integrating sphere is illustrated in Figure 4.8.

By illuminating a sphere with the scattered light signal from a diffuse reflectance standard of known optical properties one can calibrate the sphere [232, 235]. At a later time, the scattered light signal from a sample of interest can then be measured and scaled to that calibration so one can infer the TIS of that sample.




Figure 4.9: Surface radiation exchange diagrams. (Left) A diagram showing the general exchange between two randomly oriented surfaces. (Right) The exchange is now confined between two points belonging to the same hollow sphere.

The following from Equations 4.2-4.8 are derived and discussed much more extensively in the “Integrating Sphere Theory and Applications” technical guide published by Labsphere Inc. [232]. The theory of the integrating sphere is derived from the theory of radiation exchange within an enclosure of diffuse surfaces, and although the general theory can be quite complex, a sphere yields a simple, yet unique, solution. Consider first, as illustrated in Figure 4.9, two diffuse surfaces randomly orientated with respect to each other and separated by some distance S , and a radiation exchange that will take place between two differential elements on these surfaces. The fraction of energy leaving the first element on the first surface dA_1 , and reaching the second element on the other surface dA_2 , is known as the exchange factor $dF_{d_1-d_2}$, and is given by

$$dF_{d_1-d_2} = \frac{\cos \theta_1 \cos \theta_2}{\pi S^2} dA_2, \quad (4.2)$$

where θ_1 and θ_2 are measured from the surface normals. Now consider that those two differential elements dA_1 and dA_2 are in fact two points on the internal perimeter of a spherical shell of radius r . The distance S between the two points can now be written as

$$S = 2 r \cos \theta_1 = 2 r \cos \theta_2, \quad (4.3)$$

and therefore, in substituting Equation 4.3 into Equation 4.2, for the internal wall of a sphere

$$dF_{d_1-d_2} = \frac{dA_2}{4\pi r^2}. \quad (4.4)$$

This result is significant as it shows the radiative exchange factor inside a sphere is independent of distance between the two points, and of viewing angle. Therefore the fraction of flux received by dA_2 is the same for any radiating point on the sphere.

Finally, if the infinitesimal area dA_1 instead exchanges radiation with a finite area A_2 , Equation 4.4 becomes

$$dF_{d_1-d_2} = \frac{1}{4\pi r^2} \int_{A_2} dA_2 = \frac{A_2}{4\pi r^2}, \quad (4.5)$$

and since this result is also independent of dA_1

$$F_{1-2} = \frac{A_2}{4\pi r^2} = \frac{A_2}{A_S}, \quad (4.6)$$

where A_S is the surface area of the entire sphere. Therefore, the fraction of radiant flux received by a detector of active photo-surface area A_2 placed at an opening port of the integrating sphere, is simply the fractional surface area it consumes within the sphere. Another important result follows from this, that when comparing signals from two different samples (of comparable reflectivity) using the same integrating sphere, one can simply compare the ratio of the scattered light signal directly striking the photodetector in each case – no extrapolation of the total flux of all the scattered light radiated from the sample surface and covering the entire sphere wall is necessary.

The ports of the integrating sphere allow for interaction with the lab environment. These can be used to position samples of interest, photodetectors, and light sources. The reflectivity of the surface of the sphere as well as the size and location of ports, detectors, and baffles will all influence how the light bounces around inside the sphere [233]. All of these parameters will influence the light integration ability of a sphere and introduce complexity that will have to be accurately accounted for. The radiance L_S of an internally illuminated fully enclosed sphere (i.e. with no ports), from an input flux Φ_0 , when the sphere has wall reflectance ρ (where $|\rho| < 1$), after the incident light has undergone N reflections from the walls, would be given by

$$L_S = \frac{\Phi_0}{\pi A_S} \rho \sum_{n=0}^{\infty} \rho^n = \frac{\Phi_0}{\pi A_S} \frac{\rho}{1 - \rho}, \quad (4.7)$$

where n is defined as ($n = N - 1$). But when the internal wall reflectance uniformity is altered by entrance ports of total fractional sphere area f , the sphere surface radiance then becomes [232, 234]

$$L_S = \frac{\Phi_0}{\pi A_S(1-f)} \frac{\rho(1-f)}{1-\rho(1-f)} = \frac{\Phi_0}{\pi A_S} \frac{\rho}{1-\rho(1-f)} = \frac{\Phi_0}{\pi A_S} \mu, \quad (4.8)$$

where μ is referred to as the sphere multiplier.

Figure 4.10 shows the integrating sphere configuration which was used in this study to measure the TIS signal of highly reflecting samples. Measurements are made by scaling the scattered light signal from a sample of interest, to the scattered light signal from a diffuse reflectance standard calibration sample of known optical properties. Laser light enters from the north port of the sphere, striking a sample of interest at the south port, which the sphere rests atop to minimise the gap and minimise stray light entering or leaving the detector from the south port. The laser beam enters at near normal incidence striking a highly reflecting sample, so the vast majority of the light exits again from the north port and only diffusely scattered light remains to illuminate the walls of the sphere. A photodiode is placed on the west port to allow for signal acquisition - note the placement of a baffle directly between the detector and the scattering sample.

Equations 4.2-4.6 hold only for a diffusely and uniformly illuminated detector. When using integrating spheres, it is therefore important that the section of the sphere sampled with the photodiode does not contain any hot-spots directly irradiated from the sample or laser. Any light entering the photodetector, before it has been diffusely reflected off the inner walls will introduce a false response [232], poor spatial light integration [233], and consequently, give rise to artifacts in the measured signal. Conversely, the more bounces of the light off the inner-wall before entering the detector, the more uniform that distribution




Figure 4.10: The Caltech integrating sphere setup used to quantify the scatter of high quality optics. (Left) A labelled photograph of the part of the setup close to the integrating sphere. (Right) A diagram of the integrating sphere internals.

of light becomes [233]. Large errors can also arise if the angular distribution of the light reflected from the specimen of interest and entering the detector is different from that reflected by the calibration standard sample [232, 233]. To prevent all these errors, a baffle is placed inside the sphere. Baffles are struts, typically made of (or coated with) the same material as the rest of the integrating sphere inner wall, and can be considered as extensions of the sphere surface [232, 233]. They are positioned to shield the detector from first strike reflections from the sample, blocking the detector's view of this light which has not undergone diffuse reflection off the sphere surface. Its inclusion ensures the scattered light reaching the photodiode is diffuse, uniform, and free of artifacts [233].

As described by Equation 4.6, a photodiode positioned on the inner wall of the sphere detects an amount of light power which is proportional to the luminescence of the entire sphere wall, which is in turn proportional to incident light flux. This measurement is independent of the location of the detector, and (thanks to the baffle) it is also free from first strike artifacts of light coming directly from the sample [119]. Thus a perfect spatial integration is obtained.

The scattering loss from a sample is simply the ratio of the scattered power P_S to the input power from the laser beam P_I , whereas the integrated scatter (TIS) from an optic is defined as [119, 239]

$$TIS = \frac{P_S}{P_I R}, \quad (4.9)$$

where R is the sample intensity reflection coefficient.

To extract the TIS of a sample of interest (superscript S), its scatter signal must be scaled to the scatter signal from a diffuse reflectance standard (superscript REF):

$$\frac{TIS^S}{TIS^{REF}} = \frac{P_S^S}{P_I^S R^S} / \frac{P_S^{REF}}{P_I^{REF} R^{REF}}, \quad (4.10)$$

leading to:

$$TIS^S = \left(\frac{R^{REF}}{R^S} \right) \left(\frac{TIS^{REF}}{P_S^{REF}} \right) \left(\frac{P_I^{REF}}{P_I^S} \right) P_S^S. \quad (4.11)$$

Two important sources of error are electronic noise in the photodetector, and optical noise from unwanted sources coupling into the measurement [119]. The optical noise can come in the form of unwanted ambient light from the lab entering the sphere (background). It can also emerge from the laser beam in the following ways; scattering on dust particles and molecules of the atmosphere, scattering and diffraction of the laser at the entrance/exit port windows, and back-scattering from beam dumps back into the sphere [119]. All of these noise sources need to be suppressed in order to obtain a high precision measurement.

To account for the background light coupling into measurements via the open ports of the sphere, the signal on the photodiode with the laser switched off must also be recorded. We can define the signal measured on the photodiode P_{PD} as a combination of the fractional scattered light power P_S and the background signal P_B :

$$P_{PD} = \frac{A_{PD}}{4\pi r^2} P_S + P_B, \quad (4.12)$$

where A_{PD} is the surface area of the detector photo-surface, so

$$P_S = \frac{4\pi r^2}{A_{PD}} (P_{PD} - P_B), \quad (4.13)$$

and Equation 4.11 therefore becomes

$$TIS^S = \left(\frac{R^{REF}}{R^S} \right) \left(\frac{TIS^{REF}}{P_{PD}^{REF} - P_B^{REF}} \right) \left(\frac{P_I^{REF}}{P_I^S} \right) (P_{PD}^S - P_B^S). \quad (4.14)$$

With this, one can now quantify the optical scatter from multiple points across our samples and infer an average scatter for a coated surface of interest.

4.3.3 Integrating sphere scatter measurements

A full schematic of the integrating sphere setup used is shown in Figure 4.11. The light source is a 760 mW Nd:YAG 1064 nm laser. The laser is regulated by a half-wave plate and a Faraday isolator, which respectively ensure specific (linear) light polarisation and protect the laser from back reflections. An optical chopper acts to modulate the beam and suppress ambient/ electronic noise by means of lock-in detection, providing a high




Figure 4.11: Schematic of Caltech integrating sphere optical scatter set-up.

signal-to-noise ratio at the chopper frequency. Photodiodes are placed on pick-off beams before and after the optical chopper to monitor the power of the laser beam entering the integrating sphere. The beam is also passed through a quarter-wave plate and a series of convex lenses to ensure the beam is circularly polarised and remains well collimated upon entering the sphere - 0.3 mm in diameter on the sample surface. Both the sample of interest and the diffuse reflectance standard are held on a motorised linear translation stage such that a scatter map of any given sample can be made by moving the sample around the fixed position laser beam.

The integrating sphere had an internal wall diameter of 4 inches, with circular ports 1.5 inches in diameter. The polar angle collection range of the integrating sphere is defined by the sizes of the north and south ports, leading to a polar angle collection range from 1.0° to 75° . The collected scatter is measured with a photodiode coupled to the output port of the integrating sphere. The diffuse reflectance standard used is a nominally flat disk of 99% reflecting Spectralon[®], which has comparable reflectivity to the inner lining of the sphere which is made of the same material. Calibration is accomplished by installing

the diffuse reflectance standard at the south port and normalizing the measured signal to the incident power. After the calibration, Equation 4.14 can then be used to scale the scattered signal acquired from a sample of interest to the calibrated signal from the diffuse reflectance standard and infer its TIS.

4.3.3.1 Integrating sphere scatter results for the titania-silica HR coatings

The scatter from two of the 69.5% Ti “HR-2” stacks was investigated. The two samples were super polished silica disks of 1-inch in diameter where an 8 mm diameter region in the centre was coated. Scatter maps were acquired of the innermost 4.8 mm circular diameter of the samples, ensuring measurements were acquired far from the coating-substrate edge region, with a measurement taken every 250-300 microns, yielding 197 measured points. One sample was measured as-deposited, the other had been heat treated at 850°C for 100 hours - the same sample which achieved the best coating thermal noise and absorption values. Figure 4.12 shows the scatter maps taken from each sample and Figure 4.13 shows the plotted histograms of the scatter values obtained.

The mean TIS of the heat treated sample map was found to be 4.6 ppm, which is around 64% higher than for the as-deposited sample which had a mean TIS of 2.8 ppm. It should be noted that despite the increase, the scatter following heat treatment is still very low. Both maps in Figure 4.12 also show that the scatter of each sample seems to generally




Figure 4.12: Surface maps of the measured scatter from (left) the as-deposited sample and (right) the heat treated sample for optimum CTN performance. The colour bar for each is scaled from 1-6 ppm to aid visual comparison.

be slightly lower towards the edge of each mapped area. The heat treated sample has approximately double the amount of points registering greater than 10 ppm TIS than the as-deposited sample. These could be localised pockets of high scatter arising after heat treatment, but they could also arise from intrinsically worse surface roughness in the heat treated sample, or indeed differences in cleaning quality between the samples prior to measurement.

Recall, as discussed in Section 4.3.1, the surface roughness of the bare substrate makes a significant contribution to the scatter obtained. As the exact same sample was not compared prior to and after heat treatment, the differences in the amount of high scatter points could be attributed to slight differences in the substrate surface quality - overall the number of outlying high scattering points on each sample are of the same order of magnitude. Equally, cleaning quality can play a role in these measurements. The samples were cleaned with red First Contact polymer solution [240, 241, 242] and an ionised top gun source [243, 244] directly before measurements, and the measurements took place in a class 100 cleanroom, all to reduce the amount of, and opportunity for, dust accumulation on the samples. However, First Contact can occasionally leave micro-particles of residue




Figure 4.13: The data for the as-deposited sample is shown in blue (top), and the data for the sample heat treated from optimum CTN is shown in red (bottom). The bin size is 0.2 ppm, and both axes are log scaled.

on the surface when removed and of course the surface can not be completely dust-free, which statistically should not greatly impact the averaging for larger map areas with tens-of-thousands of data points, but for these 4.8 mm diameter maps with only 197 points, it could play a significant role when drawing comparisons. In fact, the 850°C sample was noted on a later microscope inspection to have a few small micro-particles of First Contact still present, and would certainly give rise to at least some of the observed higher scattering points.

The histograms of both samples in Figure 4.13 show that despite the few high-scatter points, both samples show very low overall scatter. The as-deposited sample exhibited a root-mean-squared (rms) error across the mapped area of 0.2 ppm on the mean of 2.8 ppm, and the 850°C, 100 hour heat treated, sample possessed a rms error of 0.5 ppm on its mean value of 4.6 ppm. The prediction value used for the rms error calculations was the dataset mean, and the fact that the rms error values are an order of magnitude lower than the mean in each case highlight the generally good scatter uniformity of the samples.

This same system has been used to measure all the aLIGO test masses. Specifically as part of these characterisations, one aLIGO test mass was measured multiple times under identical laboratory conditions, and the change on the mean TIS measurement of this same sample between these repeated measurements was found to be approximately 2 ppm - or 21% of the nominal mean value of 9.5 ppm. This estimate of the error arising from the repeatability of the measurements of that test mass has been applied to the scatter measurements made of all subsequent GWD ITMs and ETMs acquired in this system [123]. Accounting for this means each test mass should only need be mapped once. From this, one might expect a similar level of repeatability on any other by-and-large homogeneous HR sample measured in this system. Accounting for this 21% variance allows for better comparison of these titania-silica HR coating values with the real GWD ETM values measured on this system. This yields final values of (2.8 ± 0.6) ppm as-deposited, and (4.6 ± 0.9) ppm after 850°C, 100 hour heat treatment.

This coating at optimum CTN heat treatment has optical scatter which is near half the value of the current aLIGO test masses which was measured on the same system to be (9.5 ± 2.0) ppm [123]. Given that this coating was designed to have the equivalent reflectivity as an end test mass, and the end test masses have typically higher scatter than the input test masses [124, 245], this low scatter result is very promising. However, care of course must be taken when directly comparing TIS values from maps of different area (the aLIGO test mass maps are typically 50 mm and 160 mm in diameter) and from different substrates having undergone different face super-polishing procedures. Unfortunately, due to the demand for these samples in other measurement systems around the world, the same sample was not able to be mapped pre and post heat treatment. Doing so would have allowed for more conclusive results about the apparent increase in high scatter points after heat treatment. However a generally small increase in global scatter of around 2 ppm is clear. The low average heat treated scatter values (with the final uncertainty also accounting for known variability on repeated measurements of samples in this system) gives great confidence in the low scatter potential of this coating.

4.3.4 Angle-resolved scatterometer studies

It was postulated that perhaps any large increase in scatter with heat treatment would not be seen, if the scatter originated from the presence of crystals much smaller or much larger than the measurement beam size. The integrating sphere setup employed a 0.3 mm diameter beam, and the coating thermal noise cavity employed a 0.1 mm diameter beam. Indeed, low scatter was seen with the integrating sphere, and the CTN cavity setup would see the effects of high scatter via a significant increase in the measurement noise, for scatter values $\gtrsim 10$ ppm - which was not observed. Both of these results suggested low scatter of these samples for these beam sizes, and as a much larger diameter beam would be used in a real GWD, it was desired to verify the scatter once more, using a setup with as large as possible a probe beam. As the physical coating diameter on each sample was only one order of magnitude larger at 8 mm, this set an upper limit on the largest beam size which could be used to probe scatter from these samples.

The author launched a second scatter investigation with collaborators at California State University, Fullerton (CSUF) in which experiments were performed on the as-deposited and 850°C heat treated samples, previously measured with the integrating sphere, using an angle-resolved scatterometer (ARS) system. The light source used was a 1056 nm superluminescent diode, which was collimated such that it illuminated a circular diameter of 5.2 mm in the centre of the coated sample.

4.3.5 Angle-resolved scatterometer technique

An ARS measures the bidirectional-reflectance-distribution-function (BRDF) of a sample by keeping the angle of incident light fixed, and measuring the intensity of light reflected at many different angles. For a given surface undergoing illumination from a specific incoming direction, the BRDF describes how much light will be scattered in different directions from an illuminated point on that surface. It depends on factors such as the surface's physical properties (i.e. orientation, roughness, contamination, bulk homogeneity), surface optical properties (i.e. transmittance, reflectance, absorptance, refractive index), as well as the light power, wavelength, incidence angle, and the viewing angle of the detector [239]. A distinction between the TIS and the BRDF is that a BRDF quantifies the amount of scattered light as a function of solid angle at a particular orientation, while the TIS is a cumulative measurement of light back-scattered across all angles [246]. TIS values can therefore be inferred through integrating BRDF values measured at particular scattering angles (multiplied by cosine of the scattering angle) over the full solid angle of back-scatter (i.e. a hemisphere) [247, 248]. Note that due to measuring at discrete angles, and being limited in measurement capabilities towards high angles (close to 90°) due to signal clipping, the TIS values obtained are approximations of the “total” integrated scatter.




Figure 4.14: Angle-resolved scatterometer setup: (left) full schematic [247] and (right) image with labelled key components [249]. Both source files were edited by the author.

With the ARS, the sample surface and incoming light source orientation are kept fixed with respect to each other, and both are placed on a rotating stage and rotated in unison with respect to a fixed position detector CCD camera, which measures the scattered light at many different scattering angles. Figure 4.14 shows an image and schematic of the ARS setup.

Measuring the scattered light signal (as a function of angle) from a diffuse reflectance standard sample of known scatter, first with a power meter swapped in place of the CCD camera, and then again with the CCD, calibrates the setup, and allows for the CCD counts to be accurately converted into a power signal. The calibration samples used was a diffuse reflectance standard of 99% reflective Spectralon® [236, 237, 238], which scatters the light in a highly predictable and uniform manner [238]. This calibration enables, at a later time, the simultaneous surface image acquisition and quantification of the scattered light power from a sample of interest, placed in the system at a later time, imaged with the CCD camera alone.

The superluminescent diode of the ARS setup outputs to a 90:10 beam splitter fibre optic cable, with 10% going to a power monitoring photodiode and the remaining 90% guided to a reflective collimator. The collimated light is then fed through a linear polariser which horizontally polarizes the light, and it is then further narrowed using an iris. A superluminescent diode (coherence length $\sim 10 \mu\text{m}$) was favoured as the light source over a laser (coherence length $\sim 100 \text{ km}$) as it leads to a more constant measure of scatter because it produces less time-varying coherent effects such as the twinkling of point scatterers [246].

During the calibration measurements, the BRDF of the diffuse reflectance standard for multiple scatter angles was constructed using the following equation [239, 246, 247]:

$$BRDF = \frac{P_S}{P_I \Omega \cos \theta_S}, \quad (4.15)$$

where P_I is the incident light power, P_S is the scattered light power (detected by the power meter switched in place of the camera), which subtends a solid angle Ω , and was orientated with respect to the normal of the sample surface at a polar angle θ_S (in the plane of the light beam).

After this, with the CCD camera now in place, images of the scattering surface could be taken at the same scattering angles. As the sample is revolving and changing angle with respect to the camera, the proportion of the total image area/photo-surface that the sample's surface subtends also changes. The sample face appears widest at the lowest polar angles, but at higher angles it appears to narrow significantly. For each CCD image, a region of interest (ROI) corresponding to an area of significant measured scattered light (i.e. meeting a specific photon count threshold) can be defined. Once this is done, for each image, the counts over this entire ROI are summed and normalised by the incident light power, and the exposure time of the camera T_{exp} giving [247]

$$ARB_{\text{CCD}} = \frac{\sum_k V_k}{P_I T_{\text{exp}}}. \quad (4.16)$$

Here V_k is the value of the k^{th} pixel in the ROI, and ARB_{CCD} is the arbitrary unit count value from the CCD, normalised to the incident power and exposure time. For clarity the values pertaining to this calibration measurement specifically will now be denoted with the superscript ‘CAL’. A calibration function F^{CAL} that relates the BRDF measured previously using the power meter in place of the camera, and the corresponding CCD counts can be defined [247]

$$F^{\text{CAL}} = \frac{BRDF^{\text{CAL}} \cos \theta_s}{ARB_{\text{CCD}}^{\text{CAL}}}. \quad (4.17)$$

As the sample of interest is rotated, some diffuse light from the sample barrel and optic holder, in addition to the scatter from the beam spot, will enter from the edges of the ROI, to varying degrees at the different angles – examples of this can be seen in Figure 4.15 and later in Figure 4.19. To better estimate the optical scatter from just the sample face enclosed in the ROI, instead of one ROI, six concentric elliptical ROIs of increasing size are defined to decouple the external diffuse scatter skewing the CCD counts. The sum of the counts within each of these ROI are calculated and normalized as in Equation 4.16. These six values are then fit to a linear function versus pixel area (of each given ROI), which then defines a y-intercept value that estimates the true sample BRDF [246]. One would expect




Figure 4.15: ARS sample images taken at different polar scattering angles. The perimeters of the six concentric elliptical ROIs are defined with faint blue lines in each image - the primary (smallest) ROI has the boldest blue line and represents a height on the sample of approximately 10 mm. Note the diffuse light not originating from the sample face entering the ROIs at 17° , 53° , and 75° .

the unwanted light scattered off the barrel/mount coupling into the measurements to be much more prevalent in the larger of the concentric ROIs (closest to the barrel/mount) and have decreasing effect towards the centre of the optic in the smaller of the ROIs. For any future given sample of interest the calibration factor along with this corrected y-intercept fit value of its counts, dubbed ARB_{CCDY} , can be used to extract its corresponding BRDF as follows [246, 247]

$$BRDF = F^{CAL} ARB_{CCDY}. \quad (4.18)$$

For all images taken with the camera, a corresponding dark image of the same exposure time, but with the light source turned off, is taken and subtracted from the measurement image. This removes the camera noise and any ‘hot’ pixels. Steps are also taken to suppress all ambient room light and changing background level as much as possible during measurement: the room lights are switched off, the operator controls the apparatus remotely, and shielding is placed surrounding the setup, as shown in Figure 4.14.

The TIS can be estimated by integrating a measurement of the BRDF multiplied by $\cos \theta_S$ over the full solid angle of scatter (i.e. over a hemisphere for back-scatter). The hemispherical reflectance, called R_H , is related to the sample’s TIS and its reflectivity R (≈ 1 for HR samples), and can be stated in a rearranged and expanded form of Equation 4.9 which incorporates the BRDF [247, 248]:

$$R_H = \frac{P_S}{P_I} = R \ TIS = \int_0^{2\pi} \int_0^{\pi/2} BRDF \cos \theta_S \sin \theta_S \ d\theta_S d\phi_S, \quad (4.19)$$

where ϕ_S represents the azimuthal scattering angles. For clarity, the polar scattering angles θ_S characterize the deviation of scattered light from the incident beam’s direction in the plane perpendicular to it, while the azimuthal scattering angles describe the orientation of scattered light around the incident beam’s axis. As the angle of in-plane sample rotation with respect to the camera has to be $< 90^\circ$ so as to image light scattered from the sample face (and not have it viewed through the barrel), one is limited to $0 < \theta_S < \pi/2$ for the subtended polar angle. An assumption that the BRDF is independent of azimuthal scattering angles can be made, since the coated optics being investigated have well polished faces that should scatter isotropically at a given polar angle [247]. This allows for the integration of the accumulated ARS data over only polar angle. The solid angle integral

can therefore be approximated as a sum of the scatter in individual sectioned ring areas Ω_{ring} centred on the polar angles of measured scatter θ_S . For visualisation, imagine a segment of a sphere, sectioned from its centre to its edge by a cone with slants defined by angle θ . This conical section of the sphere, with its apex at the apex of the solid angle, and with apex angle 2θ defines the solid angle the light subtends in the field of view in the detector. This solid angle is defined as [250]

$$\Omega = 2\pi(1 - \cos \theta). \quad (4.20)$$




Figure 4.16: Diagram showing the conical section of the hemisphere of backscatter that the camera samples during ARS measurement.

To calculate the solid angle of a ringed region of this spherical based cone between two different subtended angles, one simply subtracts the subtended angle of the first region from the second, yielding

$$\begin{aligned} \Omega_{\text{ring}} &= 2\pi(1 - \cos \theta_2) - 2\pi(1 - \cos \theta_1) \\ \Omega_{\text{ring}} &= 2\pi(\cos \theta_1 - \cos \theta_2). \end{aligned} \quad (4.21)$$

Applied to the ARS experiment, Ω_{ring} is the solid angle subtended by angles between θ_1 and θ_2 defined by

$$\begin{aligned} \theta_1 &= \theta_S - d\theta/2 \\ \theta_2 &= \theta_S + d\theta/2, \end{aligned} \quad (4.22)$$

where $d\theta$ is the angular resolution (i.e. step) of the ARS measurement.

With the solid angle integral being approximated as a sum of the scatter in these individual sectioned ring areas for each of the measured scattering angles θ_S , Equation 4.19 can be simplified to yield [247]:

$$R \text{ TIS}(\theta_S) = \Omega_{\text{ring}}(\theta_1, \theta_2) \text{BRDF}(\theta_S) \cos \theta_S, \quad (4.23)$$

and finally

$$\text{TIS}(\theta_S) = \frac{2\pi(\cos \theta_1 - \cos \theta_2) \text{BRDF}(\theta_S) \cos \theta_S}{R}. \quad (4.24)$$

One can also see the similarities to Equation 4.23 as a rearranged and equated form of Equations 4.9 and 4.15 with Ω approximated as the sum of all the Ω_{ring} sections. Equation 4.24 yields the polar-angular dependant TIS, which can be sampled and summed over many polar scattering angles to achieve the integrated scatter value. It should be noted that here the measurements are not ‘total’ integrated scatter since only a given angle range can be accessed due to obstacles such as the laser and the optic holder. So for this technique the integrated scatter is extracted from integrating the BRDF over only a partial hemisphere of back-scattering (with isotropic scattering in azimuthal angles assumed) [246, 247].

4.3.6 Angle-resolved scatterometer results

The two 69.5% Ti high-n layer titania-silica/silica HR stack samples, previously measured with the integrating sphere, were remeasured using the ARS. In doing so, a direct comparison between the two techniques could be drawn, and it was hoped to learn if there was significant excess scatter witnessed in using a probe beam encompassing around $\times 300$ greater surface area in the 850°C heat treated sample than the previous integrating sphere result.

The effects of various cleaning processes on the scatter measurements were also recorded for these optics. Specifically samples were measured in the ARS after: (A) only being lightly sprayed with ionised airflow to remove surface dust after transport; (B) measured after drag wiping with isopropyl alcohol followed by ionised airflow spray; and finally (C) after cleaning with a commercial polymer cleaning solution First Contact plus ionised airflow spray. The latter was the cleaning technique the author also adopted in the integrating sphere experiment for these samples.

4.3.6.1 Analysis of images after different cleaning methods

The 850°C heat treated sample was measured in the ARS first after only being sprayed with ionising airflow to remove dust, and then it was measured a second time after being cleaned with First Contact immediately prior to ARS measurement. Upon removing it from its container at CSUF, faded streaks could be seen across the sample face upon visual inspection. These were attributed to drag wipe cleaning with contaminated chem-




Figure 4.17: ARS images for the two titania-silica 69.5% high-n layer HR stack samples before and after cleaning with First Contact ('FC cleaned'). The as-deposited sample was drag wipe cleaned with isopropyl alcohol and sprayed with ionised nitrogen gas prior to the first measurement to remove dust ('DW cleaned'), whereas the 850°C heat treated sample was not chemically cleaned and only sprayed with ionised nitrogen flow for the first measurement ('not cleaned'). The primary ROI, ~ 10 mm tall, is indicated with the boldest blue circle in each image.

icals/cloth occurring during previous measurements undertaken by colleagues for other experiments, and were subsequently easily removed by future rounds of cleaning. However, measuring it in this state after only lightly spraying with ionised airflow to remove surface dust, and comparing to the results after a subsequent round of cleaning would serve as an excellent baseline showing the impact of any surface contaminant on the scatter.

The as-deposited sample on the other-hand, looked immaculate to the naked eye. It was first measured in the ARS immediately after a standard drag wipe cleaning with isopropyl alcohol [247] and then sprayed with the ionised nitrogen gun in an attempt to remove any surface contamination not visible-to-eye. After this it was also measured a second time after cleaning with First Contact and the nitrogen gun, comparing two well established optical cleaning methods for scatter measurement preparation [246, 247].

The images presented in Figure 4.17 on the whole appear quite dark, due to the samples globally having quite low scatter, but with a few high scattering points seen as white pixels. There is a clear reduction in the number of these high scattering points after First Contact cleaning in both samples versus the previous measurements. The 850°C annealed sample in the ‘not cleaned’ state where it was only lightly sprayed with ionised air, clearly exhibits much more scatter than all the other measurements. Looking carefully at Figure 4.17 the pattern of faint streaks previously observed by eye can be clearly seen in the measurement as clustered points of high scatter, aligned near-vertically. The other three images all show a bright thin vertical line of high scatter right-of-centre. This was determined to be the edge of the 8 mm coating surface to within 0.2 mm accuracy, by scaling the feature’s pixel distance from the centre, and the total optic diameter in pixels, both into millimetres using geometry measurements from a set of digital calipers. Its presence indicates notably high scatter at the very edge of the coatings compared with the rest of the coated surface, and was particularly significant in the scatter at the lower polar angles as its presence quickly faded at angles greater than 10°. On the other-hand, for the 850°C sample before cleaning with First Contact, the coating edge scatter is clearly obscured by the much higher scatter surface contamination.

In each measurement, the roll-axis orientation of the samples is kept relatively similar with the aid of markings on the barrel. Any features then seen in both images of the same sample after separate rounds of cleaning are likely to be real features embedded in, or just below, the coating on the front face, or to originate from back surface back-scattering viewed through the front face of the optic. Any high scattering features of the front surface will rotate with the direction of travel, and any originating from the back surface would be seen to move in the opposite (yaw-axis) direction to polar angle rotation. Within the coating diameter there is no distinct pattern, clustering, or other macroscopic trend to the high scatter points present in the First Contact cleaned optics, and all of these are small compared to the total imaged surface. This is good evidence for a well cleaned optic, and clearly the First Contact cleaning compared with just drag wipe, or top gun cleaning resulted in the lowest surface contamination.

4.3.6.2 Comparison of ARS images and integrating sphere maps

The images of both titania-silica/silica HR samples which were cleaned with First Contact corroborates the integrating sphere results shown in Figures 4.12 and 4.13. Each indicate the presence of very few points on each surface with high scatter, and indeed that the 850°C sample possessed more of these high scattering points than the as-deposited sample, correlating with the higher scatter of that sample overall.

Re-orientating and overlaying the 4.8 mm diameter integrating sphere scatter maps from Figure 4.12 on the smallest angle-of-incidence scatter images for both samples yields Figure 4.18. Care needs to be taken here as the ARS images are each taken at approximately 3-5° from normal, and the integrating sphere measurements are taken at normal incidence. There will therefore be slight - predominantly horizontal - spatial distortions in the matching of map and image which will skew the alignment of the ARS imaged points with corresponding integrating sphere map regions. As such, some leniency must be afforded for high scatter points detected nearby, but not exactly matching up with where they were recorded in one experiment or the other. That being said the integrating sphere maps and the ARS images clearly match up well, with only 3-4 high scatter points (red circles) on the matched region measured in each experiment seemingly not appearing in one measurement or the other of the same optic.




Figure 4.18: ARS First Contact cleaned images for the two titania-silica 69.5% high-n layer HR stack samples overlapped with the integrating sphere (IS) scatter maps found in Figure 4.12. The white spots are points of high scatter as imaged on the ARS, the coloured region represents the overlay of the IS map where hotter colours correspond to higher scatter. The bottom two images duplicate the top images but with highlighted regions indicated by red circles, showing the few high scatter points that were seen in one experiment but not the other. The image is cropped at the primary ARS ROI, ~ 10 mm wide, which the 4.8 mm wide integrating sphere maps fit well within.

There are multiple factors which could lead to high scatter points being detected in one of these scatter measurements and not the other. Firstly, there might be points that only seem to highly scatter very close to normal incidence (integrating sphere) and not at $3\text{--}5^\circ$ (minimum ARS angle). Also, since different wavelengths of light were used in the two experiments, and the stack was specified for high reflectivity at 1064 nm, one might expect to see some points in the ARS images which actually originate from some 1056 nm light that transmitted through the front face, and were back-scattered off the back side and its AR coating boundaries - these are less likely to have been picked up with the integrating sphere which measured with 1064 nm light. Both sample surfaces could of course also be slightly damaged between measurements, though later microscope imaging suggested this is unlikely.

Lastly, though First Contact polymer cleaning was used in both experiments, and seemingly leads to the cleanest results, it can occasionally leave a small number of micro particles of residue when the polymer is peeled off the optic. These can be removed with subsequent rounds of cleaning, but can generally not be seen by eye, thus it is likely different regions of each optic were affected by a few small remaining particles of First Contact during each experiment. As such the apparent roughness and reflectivity changes this particulate could induce, can appear as high scatter. The same is also true of any dust that falls onto the optic after the cleaning process in each case.

4.3.6.3 ARS image high scatter features at larger polar angles

The evolution of the image scatter with polar angle rotation in these samples is perhaps best visualised in videos made via compiling all ARS images of the First Contact cleaned samples from around 3 - 80°. Five second movies can be viewed at the following links for the as-deposited sample [251] and for the 850°C heat treated sample [252] (recommended playback speeds of $\times 0.25$). The compiled images serve as an excellent visual tool, for interpreting how the scatter of the coating evolves at the many captured angles, and also highlights limitations brought on by other scattering sources.

Figure 4.19 shows the image of as-deposited First Contact cleaned sample at 29° rotation. A near solid halo of high scatter can be seen beyond the coating edge (the edge boundary indicated by the green square). At this angle this halo lies entirely within the primary region of interest defined by the solid blue ring, and is approximately 9.6 mm in diameter. To set the diameter of the coated area during deposition, a mask will be placed in-front of the optic at some distance away from it. One might expect some sputtered particles to land beyond the target coating surface as they individually approach the optic from many different angles during deposition, or even deflect off the mask. In doing so they will have a small extended range in which they can settle on the sample beyond the desired coating perimeter, based on the thickness of the mask and its distance from the face (among other factors). As many of these particles bombard the optic to build up the coating, one could expect statistically, all these independent interactions to build up a defined thin halo of




Figure 4.19: Image of the as-deposited First Contact cleaned sample, captured at a polar angle of 29° by the ARS camera. Inside the primary region of interest, a large circular perimeter of high scatter can be seen. The green square indicates the location of the front side coating edge witnessed with visual inspection (too dim to see in this image). The red circle highlights a visible section of high scatter originating from the back side AR coating edge. The orange diamond highlights scatter from the sample mount, which is bleeding into the final ROI. Lastly, the yellow triangle is highlighting scatter from the barrel.

particulate beyond the main coating edge as the optic rotates in the chamber and the high angle particles deflect past the mask edge. Of note, from these ARS images it was clear that this bright ring lay off centre by approximately 0.5 mm with respect to the main coating surface perhaps indicating a tilt in the mask.

The red circle in the figure highlights a similar mask shadow halo of high scatter on the back surface, which was seen to move opposite to the direction of rotation. The orange diamond shows bright scattering from the mount, in this image, which in general becomes more and more prominent at higher polar angles. Eventually, at large angles, light from here will cross into the regions of interest where counts are collected, increasing the apparent scatter, and limiting the ability to resolve scatter intrinsic to only the optic -

the same is true for light scattered off the barrel (yellow triangle) though in this particular image there is no significant barrel scatter coupling into the ROIs. These features act to shorten the resolvable polar angle range and at certain angles could significantly influence the measured scatter.

From the images obtained with the ARS, the differences produced by different cleaning methods were catalogued, with the best cleaning method to reduce the quantity of high scatter features seeming to be through the application of First Contact polymer solution. Images of the optics at many polar scattering angles were acquired. The observed high scatter points in the smallest angle-of-incidence images seemed to on the whole match quite well with the previous scatter maps acquired with an integrating sphere, with features exclusive to only one measurement, or the other, on the whole being attributed to differences in surface contamination, observing angle, and wavelength. Elevated scattering from the coating edge, and in a halo beyond this edge were also identified. Some back-scatter originating from the back face was seen to clearly be visible as well in both samples, suggesting a significant amount of transmission at 1056 nm at the ppm level.

4.3.6.4 ARS integrated scatter results for the titania-silica HR coatings

Images of the samples at multiple angles were acquired and analysed, with the positions of their high scattering features identified. Information was also gained about which high point scatterers of each sample truly originated from their front faces (and which were from the back face or arising from mounting features) through observing their spatial and intensity evolution with polar angle. With all this the scatter from the sample coatings could now be accurately quantified, and features arising in the BRDF and images through angle understood. First, the BRDF at each polar scattering angle was calculated as described in Section 4.18 using Equation 4.3.5. The larger the BRDF at a particular angle the more scattered light signal is being directed onto the camera.




Figure 4.20: BRDF data for the two titania-silica 69.5% high-n layer HR stack samples before and after cleaning with First Contact ('FC cleaned'). The as-deposited sample was drag wipe cleaned with isopropyl alcohol and sprayed with ionised nitrogen immediately prior to the first measurement ('DW cleaned'), whereas the 850°C heat treated sample was only sprayed with ionised nitrogen for the first measurement ('not cleaned').

Figure 4.20 shows the impact different cleaning techniques had on the BRDF. For the heat treated sample, cleaning with First Contact resulted in a reduction in the BRDF of one-to-two orders of magnitude across the whole range of polar scattering angles, when compared after only being blasted with a top gun. For the as-deposited sample, which was freshly drag wipe cleaned before the first measurement, the improvement from First Contact cleaning was less, but it still had a significant impact on the scattering below angles of 40°. These results serve to show the strong effects of surface contamination on scatter measurements.

Out of the three cleaning methods tested, clearly First Contact cleaning yielded the best results for both samples, as was already gleaned from the acquired images. The error bars on the BRDF measurements in Figure 4.20 account for the fractional uncertainties in the input power meter measurement, CCD noise, and the corresponding calibration factor

errors of these. The fluctuations on measured scattered light power from the power meter swapped in place of the camera as part of the calibration process, and the standard error from the fitted CCD counts versus ROI area are also considered. The fractional errors on all of these are added in quadrature to produce the final BRDF error bars.

Between around 48° - 51° for the as-deposited First Contact cleaned sample a noticeable sudden drop in the BRDF can be seen. Comparing with the acquired CCD images [251] provides some explanation of this. At the start of this angle range, a high scatter spot present in one of the outer regions of interest was witnessed to suddenly vanish, and just after 51° another high scattering point appears in the primary (smallest) ROI. High point scatterers suddenly appearing/disappearing can lead to significant shifting in the BRDF. If a feature of high scatter seemingly moves through the different ROIs, defined by the multiple blue circular rings which can be seen in Figure 4.17, one would expect a change in the measured scatter signal. This is because the CCD counts, which are scaled to obtain the BRDF via Equation 4.18, are also scaled by the solid angle/area subtended by the elliptical region of interest in which they are detected as per Equation 4.23, and the ellipses all encompass different sized areas. As the optic was rotated in polar angle some high scattering features could also be seen to move opposite to the direction of rotation (and, thus, cross through different ROIs), indicating that these points indeed originated from the back side, shining back through the front face and also coupling into the measurements. These will have varying influence on the measurement, depending on the ROI in which they reside for a given angle.

A general increase in scatter towards lower angles can be observed for all samples – as one might expect, more of the light is scattered at angles closer to the angle-of-incidence of the light source illuminating the sample. However in comparing both First Contact cleaned samples, we can see distinctly different BRDF trends at lower scattering angles. The heat treated sample has over an order of magnitude higher scatter at the lowest angles below $\sim 10^\circ$. As discussed in Section 4.3.6.1, this comes from the feature seen at the coating edge for that sample. Below 10° for the heat treated sample a very high scatter vertical streak was present, which can be viewed in the 3° image of Figure 4.15 and in the compiled video [252]. A feature of similar brightness was not present in the as-deposited images [251].




Figure 4.21: Integrated scatter data for the as-deposited and 850°C heat treated samples taken with the ARS after cleaning with First Contact and an ionised nitrogen.

The BRDF values were used to calculate the integrated scatter via Equation 4.24. The integrated scatter is a cumulative measurement of the BRDF values acquired for all measured polar scattering angles, and represents the total level of scatter from each sample. Figure 4.21 shows the integrated scatter values obtained for both optics, with each point representing the sum of all the previous BRDF values at each polar scattering angle, plus the BRDF value obtained at that angle converging to the ‘total’ integrated scatter.

The heat treated sample shows (5.2 ± 1.8) ppm TIS between angles of 3-80°, and the as-deposited sample exhibits (2.9 ± 1.0) ppm TIS across the same angular range. Assuming nominally identical surface polishing in both samples, as no scatter data exists for the as-deposited state of the heat treated sample, it appears that the 850°C 100 hour heat treatment has significantly increased the scatter. As the TIS values also do not at all sharply increase or spike at the high polar angles, and they appear to follow a clear defined trend, it shows that sufficient decoupling from the effects of diffuse mount and barrel scatter (very prominent at these angles) was achieved. The uncertainty range in

each value here represents a 35% random error of the ARS system, which is derived from variance in previous repeated measurements of a sample in the system under identical conditions. This error is significantly large enough that the calculated errors in the BRDF shown in Figure 4.20, and in the polar angle measurement can be ignored.

4.3.7 Comparison of the measured scatter from the different techniques

The scatter measurements recorded with the integrating sphere and ARS setups agree to within error for both the as-deposited and heat treated samples. Table 4.2 collates these integrated scatter values obtained for the 69.5% Ti titania:silica/silica HR stack samples between the integrating sphere, ARS and CTN cavity setups. Regardless of beam size used, there is general consistency for each sample measured between the setups. The same seems true, regardless of slight differences in the TIS angle capture ranges, and also regardless of the fact that the ARS measurements were performed at a slightly different wavelength than the integrating sphere and the cavity CTN setup. All of this evidence suggests that these differences between the integrating sphere and ARS setups only lead to small overall change in measured value and do not dominate the measurement over the intrinsic surface/material properties of the samples. The fact that the scatter values agree between setups where different beam sizes were employed (the largest encompassing 65% of the total coated surface area), perhaps suggests that if the heat treated coating was indeed crystallised, any crystals present were much smaller than the smallest beam size used (0.1 mm diameter) - i.e. micro or nanocrystallites. This was indeed found to be the case as will be discussed in Chapter 5.

A clear increase in the scatter between the as-deposited and the 850°C, 100 hour heat treated samples was observed, though the larger error range on the ARS values makes this increase less definitive from just this measurement alone. It is, however, clear from the integrating sphere measurement. Figure 4.18 also showed that both dedicated scatter systems resolved regions of high point scatterers that by and large agree in their spatial distributions. Outlying high points not seen in one measurement or the other perhaps result from differences in surface contamination during measurement or from limitations in comparing the integrating sphere TIS map with only one ARS image at the lowest angles.

Table 4.2: Comparisons of the measured scatter of the 69.5% Ti titania-silica mix HR stack coatings. Upper bounds from the CTN cavity setup, and estimates of the TIS from the integrating sphere (IS) and ARS setups are displayed. Current aLIGO ETM coating values obtained from the IS and a CASI setup are also shown. As-deposited values are highlighted in blue, with optimal CTN heat treatment values in red.

| sample and setup | beam diameter [mm] | λ [nm] | TIS [ppm] |
|--------------------------------------|--------------------|----------------|-----------|
| as dep. (CTN cavity) | 0.1 | 1064 | < 10 |
| as dep. (IS) * | 0.3 | 1064 | 2.8±0.6 |
| as dep. (ARS) ** | 5.2 | 1056 | 2.9±1.0 |
| 850°C, 100 hr (CTN cavity) | 0.1 | 1064 | < 10 |
| 850°C, 100 hr (IS) * | 0.3 | 1064 | 4.6±1.0 |
| 850°C, 100 hr (ARS) ** | 5.2 | 1056 | 5.2±1.8 |
| aLIGO ETMs 500°C, 10 hr (IS) * [123] | 0.3 | 1064 | 9.5±2.0 |
| aLIGO ETMs 500°C, 10 hr (CASI) [124] | 2.0 | 1064 | 4.9±1.5 |

* For the IS, integrated scatter is given for scattering angles between $1^\circ < \theta_S < 75^\circ$ [123].

** For the ARS, integrated scatter is estimated from integrating the BRDF over the partial hemisphere of backscattering of approximately $3^\circ < \theta_S < 80^\circ$, assuming isotropic scattering in azimuthal angles. [246, 247].

The integrating sphere TIS map is constructed with the system seeing point scatterers across the entire sampled angle range at once, whereas the ARS images can only show the high scatterers present at the one discrete image angle, and many were witnessed to appear/disappear with rotation. Both setups also showed that for each sample there are a small number with points of much higher scatter present than the average scatter of the rest of the surface.

The current Advanced LIGO optics have all been characterised using the same integrating sphere setup at Caltech [123] and using a Complete Angle Scatter Instrument (CASI) setup at the coating manufacturer LMA [124]. Table 4.2. Though at first these aLIGO test mass numbers from the two setups may seem quite different, this is a result of the different analysis methods each employs. The integrating sphere setup's mean TIS is calculated by measuring the TIS (i.e. scatter at all angles) for each mapped point individually, and then the values for all points are averaged. Instead, the CASI constructs a scattering map at only one fixed scattering angle for all points. An average value of the BRDF at that one fixed scatter angle for all points in the CASI map is found. Then, to convert this measurement into TIS, a point representative of this average BRDF value is chosen, and a complete BRDF measurement as a function of the scattering angle for just that point

alone is then made. In integrating the BRDF curve for that point versus angle, the TIS is estimated [253]. The CASI analysis does not, therefore, factor in that (e.g.) the higher scatter points (say from contamination, damage etc.) may not produce the same BRDF trend versus angle.

When comparing the integrating sphere map histograms of many of the test masses (similar to the ones in Figure 4.13) one can observe that, the peak of the histograms (typically lower than the reported mean value) agree well with the stated values from the CASI system, usually to within 1ppm or better. As example, the scatter measured with the integrating sphere setup from two of the super polished ETM (“SPETM”) aLIGO optics, showed the sample denoted SPETM02 had integrating sphere measured mean TIS of 7.2 ppm, but the most commonly occurring scatter values from the surface map (peak of the histograms) are between 4-5 ppm [254] which is consistent with the published CASI measurements for that sample [124]. Likewise another sample, SPETM04, had a mean TIS value of 10.5 ppm as measured with the integrating sphere, but with the mode from the mapped points being between 7-8 ppm – this is again consistent with the CASI results of 8 ppm for this optic [255]. So in the literature there is in fact agreement between the integrating sphere maps and the CASI’s TIS values generated from a representative point, despite the initial appearance of disparity from the values in Table 4.2. As the ARS constructs a full angular dependant BRDF from scattered light of a ~ 5 mm illuminated surface at once, it is quite reasonable that it produced quite comparable values to the integrating sphere ~ 5 mm diameter map mean measurement in this study. However one may observe if sampling from one point on a much larger coated area optic, such as a GWD test mass with this ARS, that it may then align more with CASI value due to the relatively similar size of beam used (assuming that the point is representative of the average scatter).

In comparing the current GWD mirror coatings to the titania-silica based HR coatings (which have full ETM specified reflectivity), it can be seen that the candidate material performs just as well or better than the coatings currently installed in Advanced LIGO and Virgo in terms of scatter. Comparing the optimum heat treatment numbers for both coatings measured on the same integrating sphere setup yields the most direct comparison, and the scatter from the titania-silica coatings in this case is around half of what was

measured across multiple aLIGO test mass optics. In also comparing the peak value of the scatter histogram obtained for the heat treated titania-silica sample of ~ 2.4 ppm (see Figure 4.13) to the LMA CASI obtained current aLIGO ETM coating values, again this factor of two improvement can be seen. Note that great care is still needed when comparing these one inch test samples to the aLIGO/adV mirrors, due to the large difference in map area, as well as potential differences in the substrate surface quality, as both substrates will have undergone different polishing procedures. Nevertheless, these low scatter results for the optimal CTN heat-treated coating are very promising indeed.

4.4 Titania-silica optical loss study conclusions

In these dedicated optical loss studies, initial estimates of the optical absorption and scatter made by collaborators at MIT using their CTN cavity setup were able to be refined more precisely and found to a much greater degree of confidence. The author undertook and coordinated dedicated absorption characterisation studies of these titania-silica mix coatings using the PCI technique. Consistent results from two independent setups for the 69.5% Ti content HR stack were found, yielding (0.82 ± 0.11) ppm after being heat treated at an optimum CTN temperature of 850°C for a dwell duration of 100 hours. This equated to approximately an 86% reduction for the coating absorption versus its as-deposited state. It was also around half the initial estimate for the heat treated absorption from the CTN cavity setup.

Two optical scatter characterisation investigations were also launched, using an integrating sphere and ARS to quantify the integrated scatter over a wide angle range. The scatter results from these also agreed with each other rather well and yielded integrated scatter values at heat treatment temperature for optimum CTN of ≈ 5 ppm. This is comparable to the level of, or lower than, the current coatings used in room temperature detectors, and comfortably meets the target requirement for the next detector upgrades, being at half the required design specification [124].

Both of these results highlight the great potential of titania-silica mixes as a viable mirror coating material candidate for the next generation of gravitational wave detectors. They certainly appear to reach well below the average scatter requirements for future detectors [124] and an average absorption value of below 1 ppm, though not quite the 0.5 ppm target, still shows great promise - especially for a preliminary investigation. These optical loss studies along with the previous mechanical loss and CTN investigations made up the bulk of the work leading to a Physical Review Letter first author publication in October 2023 [171], and has also resulted in follow up studies being commissioned on behalf of the aLIGO and adV collaborations, where the GWD coating manufacturer LMA has begun production of this material for investigation [224].

Chapter 5

Titania-silica HR coating crystallisation investigations

5.1 Introduction

Raman spectroscopy is an analytical measurement technique using scattered light from a sample to measure its vibrational energy modes, and probe its chemical and structural composition. In the case of the titania-silica mix coatings that have been under investigation through Chapters 3 and 4, this technique was used to determine primarily the presence of crystallisation. The onset of crystallisation in amorphous coatings has previously been reported to worsen the optical properties such as absorption and scattering [82, 256]. It was hoped that for the high-n layers in the stacks, the mixing of titania with silica, which is known to not crystallise until much higher temperatures, would somewhat suppress the crystallisation of the mixed material, above the crystallisation temperature of pure titania. Pure IBS titania is known to crystallise above around 350°C, and in 1987 C.Y. She demonstrated the crystallisation suppression of single layer IBS titania-silica mixes versus their pure IBS titania counterparts made via the same process [257]. They found that a mixing with 10% silica content was sufficient to suppress crystal formation to a heat treatment temperature almost 200°C higher than for pure IBS titania. As the optimum heat-treatment temperature for CTN found for these HR stacks is several hundred degrees above the pure titania crystallisation temperature, and this optimum CTN heat treatment was accompanied with low absorption values and no significant increasing absorption trend versus lower heat treatment temperature measurements, determining the structure of the HR coatings became of great interest.

The culmination of the studies presented here, as well as those in Chapters 3 and 4, resulted in a first author Physical Review Letters publication in 2023 [171].

5.1.1 Raman spectroscopy brief principles

Raman spectroscopy extracts chemical and structural information of a sample, through the detection of Raman scattering. Any sample exposed to monochromatic light of some frequency, will reflect, transmit, absorb, and scatter the incident light. The vast majority of this scattered light will have the same frequency as the incident light - Rayleigh scattering [258, 259]. However some scattered light will also be shifted in frequency with respect to the inbound light - this is known as the Raman effect. This inelastic scattering of light was first postulated by A. Smekal in 1923 [260], and it would be five more years before it was experimentally observed by C. V. Raman and K. S. Kishnan [261], for which Raman was awarded the 1930 Nobel Prize in Physics [262]. This Raman scattering is much rarer than Rayleigh scattering; about 100 times less intense than Rayleigh scattering [263] and only occurring for approximately one in 10^6 - 10^8 incident photons [264, 265].

One can visualise scattering simply as a two photon process [263]. When a photon of incident monochromatic light interacts with the sample, some energy is absorbed by the sample, allowing one of its electrons in a particular energy state/vibrational level to ascend to a higher virtual energy state. This virtual energy state is not stable, and as such, the configuration is very short-lived, with the electron almost immediately falling back to a lower energy state, losing energy and subsequently re-emitting a scattered photon as it does [266, 267]. In the case of Rayleigh scattering, the electron will fall back to its original energy level, and therefore the energy (and hence wavelength) of the incident photon is conserved, with only the direction being changed on emission. In Raman scattering, the process is inelastic and the electron will fall back to a different energy state than it was in prior to interaction with the light field. In doing so, the energy lost by the electron is different from the energy it absorbed from the incident photon, and consequently the emitted and incident photon energies, and therefore wavelengths, are not equal - giving rise to the Raman effect [268]. The inequality condition for incident (*i*) and scattered (*s*) photons in Raman scattering can be stated as:

$$\frac{h}{\lambda_i} \neq \frac{h}{\lambda_s}, \quad (5.1)$$

where λ represents the photon wavelengths and h is Planck's constant. When the wavelength of the emitted photon is greater than that of the incident photon, the electron has net absorbed energy in the interaction - this is known as Stokes Raman scattering. Inversely, if energy is lost by the sample as the electron settles into a lower vibrational level than it was in initially, the scattered photon wavelength decreases and its energy increases, and this is known as anti-Stokes Raman scattering. Quantum mechanically Stokes and anti-Stokes are equally likely processes [269]. However, with an ensemble of molecules, the vast majority will be in the ground vibrational level (Boltzmann distribution) and anti-Stokes scatter is less statistically probable [264]. As such the Stokes Raman scatter is always more intense than the anti-Stokes, and therefore it is nearly always Stokes Raman scatter which is measured through Raman spectroscopy.

A material will have a characteristic 'Raman fingerprint' with spectral features based on its molecular structure and chemical composition. In practice, Raman scattering is represented as a spectrum as a function of intensity versus the Raman shift. Raman shift is simply the difference between the reciprocals of the scattered photon wavelength and the excitation wavelength. Plotting this way allows for visual comparison of a spectrum to other spectra even when different laser excitation wavelengths are used. Typically, Raman shifts are reported in wavenumbers, which have units of inverse length. The Raman shift in wavenumbers, $\Delta\tilde{\nu}$ is given by

$$\Delta\tilde{\nu} = \frac{1}{\lambda_i} - \frac{1}{\lambda_s}. \quad (5.2)$$

Most commonly, the wavenumber is given in units of inverse centimetres (cm^{-1}).

Following Placzek [270, 271], the intensity of Raman scattering at a particular Raman shift can be expressed in the form

$$\mathcal{I}_s = 4\pi^2 a^2 \Delta\tilde{\nu}^{-4} \mathbf{I}_e |\tilde{\mathbf{e}}_e \alpha \mathbf{e}_s|^2 d\Omega, \quad (5.3)$$

where a represents the fine-structure constant $\approx 1/137$. The scattered energy per unit time (intensity) into a solid angle $d\Omega$ is given by \mathcal{I}_s , while I_e represents the energy per unit area per unit time (irradiance) of the excitation incident on the sample. The unit vectors $\tilde{\mathbf{e}}_e$ and \mathbf{e}_s define the directions of the electric fields of the exciting and scattered radiation respectively, and α is the scattering tensor of a given material, which will depend on its composition, structure, and the size of the particles or features within it.

The main purpose of this investigation was to determine, after various heat treatment stages, if the titania-silica coatings had remained amorphous or had begun to form crystalline structures. The characteristic Raman spectra of amorphous and crystalline solids of the same chemical composition can appear significantly different, primarily due to the total absence and presence, respectively, of spatial order and long range translational symmetry. Amorphous materials lacking any spatial order, can be thought of as collections of formula units of the same chemical composition, but possessing varying bond angles and lengths depending upon chemical bond interactions with their nearest neighbours [272]. Their spatial arrangement is disordered, and as such amorphous material Raman spectra tend to manifest very broad bands with widths of up to several hundred wavenumbers. Conversely, crystalline materials possess long-range translational symmetry and have unique - or at least a much narrower range of - bond angles, and through this far more ordered structure manifest tall and narrow Raman peaks [273, 274]. Therefore as a sample transitions from amorphous to become more and more crystalline through heat treatment, one expects the broad diffuse spectral bands present to converge into much narrower sharp peaks with increasing intensity over a decreasing spread in wavenumber [275].

5.2 Titania-silica mix Raman spectroscopy experiment

Raman scattering was used to track any crystallisation in both of the titania-silica mix HR stacks. These experiments were carried out on coated silica substrates of various geometries, ranging from 1-inch eighth-disk wedges to full 3-inch disks taken to various stages of heat treatment. The measurements were performed during a research placement at Université Claude Bernard Lyon 1 and continued upon return to the University of Glas-

gow. In Lyon, Raman spectra were recorded using a LabRAM HR Raman spectrometer equipped with two EDGE filters - ultra steep longpass filters - meaning that only Stokes Raman measurements were possible. The incident light source was a 473 nm laser. In Glasgow, Raman spectra were recorded using a LabRAM HR Raman spectrometer equipped with two EDGE filters, with a 532 nm laser light source. Raman spectra were recorded from shifts of 100-1200 cm^{-1} , and the incident laser power was limited in each system to 0.6 mW (Lyon) and 1 mW (Glasgow). This was done to avoid any laser induced damage or crystallisation of the coatings, as it is well known that both sample heat treatment, and heat transfer from too high incident laser power can induce crystallisation in optical coatings [257]. Therefore it is important that care is taken to not impart too high laser powers and induce local crystallisation/burning when probing the structure via Raman spectroscopy (and indeed in other measurements).




Figure 5.1: Comparison of the 69.5% Ti high-n layer HR stack coating as dep and bulk substrate Corning 7980 SiO_2 Raman spectra measurements, with spectra of as-deposited amorphous IBS SiO_2 [276] and TiO_2 [277] films, and anatase (crystalline TiO_2) [278] [A][276],[B] [277],[C] [278].

Figure 5.1 shows the measured Raman spectra for the as-deposited 69.5% Ti high-n layer HR coating and the uncoated silica substrate, as well as spectra from the literature for amorphous IBS TiO_2 and SiO_2 , and the crystalline anatase form of TiO_2 for comparison. One can see that the Raman spectrum for the HR stack (black line) contains features arising from both amorphous IBS TiO_2 and SiO_2 , as expected. The coatings, which are designed for high reflectivity at 1064 nm, transmit about 90% of light at the Raman excitation wavelengths used in this study, meaning that the substrate will also produce some Raman signal. The significant influence of SiO_2 substrate signals on the Raman spectra of thin IBS films in the absence of masking materials was well documented by L. S. Hsu et al. [274]. Specifically they, in fact, investigated its effects on TiO_2 films.

In Figure 5.1, one can see that of the five broad flat peaks emerging in the measured HR spectrum, the first and third primarily result from the TiO_2 present in the coating. The second, fourth, and fifth prominences in the HR spectrum similarly seem to arise primarily from the influence of SiO_2 . Figure 5.1 also shows an example spectrum of anatase - a crystalline form of TiO_2 . Unlike with the amorphous spectra, the anatase has much narrower peaks, with maximum heights significantly higher than, and in some cases orders of magnitude, above the background. As expected the HR coating has no signs of these tall sharp peaks, confirming that no crystallisation is present in the as-deposited coating.

In Figure 5.2 a comparison of the spectra of the 69.5% and 63.2% Ti high-n layer HR stacks is shown. The absolute magnitude of the peaks between the spectra here is unimportant, rather, comparing the relative heights and slopes of the broad peaks and features within one spectrum, and contrasting these trends with those obtained in the other spectrum will allow for useful conclusions to be drawn. The spectra, on cursory inspection, look very similar. As well as having a proportionately higher SiO_2 content, the 63.2% coating is also $\sim 15\%$ thinner, so one might also expect a slightly larger impact on the spectra from the substrate. Both factors in this case would lead to an increase in the influence of SiO_2 on the 63.2% HR spectrum.




Figure 5.2: Comparison of the Raman spectra of the 69.5% and 63.2% Ti high-n layer HR stack coatings. In the inlay, the 63.2% data is re-scaled so that the second broad band between $400\text{-}500\text{ cm}^{-1}$ has equal intensity in both spectra to aid in visual comparison. The main plot spectra have underwent no specific scaling.

The relative height of the first and second peaks is very similar for both coatings (see Figure 5.2 inlay). However, the relative height of the third peak (compared to the second) is smaller in the 63.2% Ti spectrum than in the 69.5% Ti spectrum. This peak arises from TiO_2 (see Figure 5.1), and this decrease in relative height is consistent with the lower TiO_2 content. Along with this, the slope between peaks 3 and 5 - which is mostly influenced by the TiO_2 - is less steep in the coating with proportionally less TiO_2 . The fourth and fifth peaks primarily arising from SiO_2 are also more pronounced in the 63.2% spectra, which again suggests there is fractionally more SiO_2 in the thinner coating stack with less Ti. Following these baseline measurements, the heat treatment investigations could begin.

5.2.1 Titania-silica mix Raman spectroscopy heat treatment study

The Raman spectra of samples heat treated for the optical absorption, mechanical loss and coating thermal noise studies were measured at a variety of heat treatment temperatures. Figures 5.3 and 5.4 show the Raman spectra acquired for the 69.5% and 63.2% Ti high-n layer HR stack coatings. The emergence of narrower sharp spectral peaks of increasing




Figure 5.3: Raman spectra acquired for the 69.5% Ti high-n layer HR stack coating through multiple stages of heat treatment. Displayed on the inlay are the 850°C and 950°C spectra (from the samples used to measure CTN) presented on a log scaled y-axis.




Figure 5.4: Raman spectra acquired for the 63.2% Ti high-n layer HR stack coating through multiple stages of heat treatment. The crystallised spectra are displayed on the inlay shifted in y to better visualise peak evolution.

intensity at higher heat treatment temperatures indicates crystallisation and its growth in these samples. For the 69.5% Ti HR coating, the spectral indicators of crystallisation manifest at 575°C, whereas for the lower titania concentration, they are not observed until 600°C. In both coatings, at the manifestation points we see the clear emergence of 4 peaks at approximately 150, 400, 530, and 630 cm^{-1} . These are all consistent with anatase crystalline titania formation, modulated by the amorphous silica and titania signals. This is perhaps best visualised from comparison of the spectra in Figure 5.1. As the samples are progressively heat treated beyond the initial crystal formation point, the spectra show stronger evidence of titania crystallisation; the peak intensities increase, and most peak positions shift slightly in wavenumber, to then generally align better with values from the literature for crystalline titania [279, 280, 281]. At higher temperatures another small characteristic anatase titania peak also clearly emerges in the spectra at around 200 cm^{-1} .

Table 5.1 shows comparisons of the spectral peak locations found for both stacks after different stages of heat treatment, and values from the literature. The introduction of silica seems to have suppressed titania crystal formation to higher temperatures than expected for pure titania by at least 100°C, as expected from the literature [257]. The coating with a higher silica concentration crystallised at a slightly higher temperature (approximately 25°C higher), which is consistent with silica suppressing crystallisation. The continued growth of the peaks seen in all measurements up to 950°C also suggests that none of the coatings are fully crystallised and still retain some amorphous structure.

5.2.2 Titania-silica mix Raman spectroscopy heat treatment study - extended discussions

Interestingly, the 63.2% HR coating behaves differently to its counterpart. Another weak peak can be seen manifesting at approximately 250 cm^{-1} both at the onset of crystallisation and as it persists throughout all subsequent stages of heat treatment. It is not a characteristic anatase peak, but rather indicates the presence of rutile crystalline TiO_2 in the coating [279, 280, 281]. Even after heat treatment at 850°C peaks are seen to continue to rise in intensity, suggesting the structure is not yet fully crystallised, but if heat treated to even higher temperatures beyond what is necessary for our purposes, one might also

Table 5.1: Raman shift peak positions recorded for the two HR coating stacks after crystallisation emergence and as heat treatment progressed, compared with previous Raman measurements of crystal powders. The rutile peak row is highlighted in grey.

| Crystalline TiO ₂ Raman peak positions [cm ⁻¹] | | | | | | | Expected around: | Assignment |
|---|-------------|-------------|-------------|-------------|-------------|---|-----------------------------------|------------|
| 63.2% 600°C | 63.2% 650°C | 63.2% 800°C | 69.5% 575°C | 69.5% 850°C | 69.5% 950°C | | | |
| 146 | 148 | 143 | 148 | 140 | 140 | 146 ^[X] /147 ^[Y] | E _g | |
| - | 197 | 197 | - | 198 | 191-202* | 194 ^[X] /198 ^[Y] | E _g | |
| 240 | 248 | 245-249* | - | - | - | 235 ^[X] / ^[Y] /247 ^[Z] | - | |
| 399 | 399 | 397 | 404 | 398 | 393 | 395 ^[X] /398 ^[Y] | B _{1g} | |
| 512-521* | 520 | 520 | 521 | 517 | 516 | 514 ^[X] /515 ^[Y] | A _{1g} + B _{1g} | |
| 604-620* | 623 | 634 | 615 | 633 | 634 | 636 ^[X] /640 ^[Y] | E _g | |

* - low spatial resolution measurement

[X]^[279], [Y]^[280], [Z]^[281]

expect a fully crystallised titania anatase structure to partly convert to rutile [273]. The different structural behaviour in these two high-n coatings which differ by $\sim 6\%$ Ti content through heat treatment is fascinating and noteworthy, but further characterisation is beyond the scope of this thesis.

Table 5.1 shows the evolution of the crystalline spectral peaks through heat treatment. There are many factors which could lead to slight disagreements when comparing Raman spectra of identical materials between different setups, one being the room temperature at which a particular measurement was taken [282]. In the case of a multi-layer coating stack, one could also expect the light field penetration through the coating stack for different wavelengths of laser light to differ, leading to varying levels of influence from the substrate on the final spectra, though as discussed later Figure 5.5 shows this was not the case in this study. The enhancement of the Raman signal from thin films is highly dependent on the complex index of refraction of the individual coating layers and substrate materials [283]. Among other things, the index of refraction determines the period and amplitude of the variation in the Raman intensity with thickness. This all being said, good agreement between these two coatings Raman peaks with expected peak locations coming from bulk crystalline powders is observed. In all bar the primary peak, a general shift towards better agreement between the measured thin film data and the bulk powder values is observed with increasing heat treatment (as the crystal structure grows) - particularly for the three lower intensity anatase-TiO₂ peaks at the highest Raman shifts. These three peaks, coincide with of a region of high signal arising from SiO₂. The total signal will therefore

have a larger SiO₂ contribution, potentially distorting the positions of TiO₂ peaks. As the crystalline TiO₂ signal increases with heat treatment, the effect of the SiO₂ on the overall signal is diminished. Since the coatings do not appear to be fully crystallised, there will also be some residual amorphous TiO₂ influencing the measurements as well.

In general, as the level of crystallisation in these coatings continues to increase through heat treatment, the amorphous signal should become relatively lower as the peaks increase in intensity and narrow, and one might expect the thin film results to conform more to the bulk crystal powder values. Through heat treatment we may also hope to relieve stress in the coatings induced by the thermal mismatch of the coating and substrate properties, which can also affect the Raman spectra [284]. The vibrations of a crystal structure are described not in terms of the vibrations of individual atoms but in terms of collective motions in the form of waves, known as lattice vibrations [285]. It is these lattice vibrations that can give rise to Raman scattering and through measurement we gain information of its 'Raman active' lattice vibration frequencies [286]. In this way incident light can interact with the crystal to induce or destroy one or more lattice vibration quanta - phonons. With the presence of a strain/stress in the crystal, its lattice vibration changes [284], and hence the spectral peak positions may shift significantly. The thermal expansion coefficient of SiO₂ is generally positive and near zero in the range of $(10^{-6} - 10^{-7})/K$ [287]. Glass mixes of TiO₂:SiO₂ have previously been shown to achieve both negative and positive thermal expansion coefficients before and after heat treatment [287, 288, 289]. This expansion mismatch in the alternating coating layers, and indeed potential swapping of sign, through heat treatment could lead to significant stress changes in the coating. Indeed a coating stack with an overall negative thermal expansion paired with the SiO₂ substrate could lead to significant tensile stress on the sample and could also factor into the cracking witnessed on some of the samples [206].

Overall, the peak positions in the samples heat-treated at the highest temperatures do agree quite closely with crystal powders. However the primary anatase peak at 140 cm^{-1} at the highest heat treatment appears to be significantly red shifted (lower Raman shift than expected) with respect to the expected values. This could perhaps, as discussed, result from thermal (tensile) stress effects [290], but it could also perhaps hint at an evolution of the O/Ti stoichiometry in the high refractive index layers, which would

influence the crystal structure and Raman shift. Parker and Siegel noted previously in a series of studies [291, 292], that for the primary anatase-TiO₂ Raman peak, purely oxygen content, and not crystal growth, induced a large peak shift. Specifically, between samples with relatively similar O/Ti ratios of 1.96 and 2, a decreasing 9 cm⁻¹ Raman shift was noted with the increasing O content. Such a small change in stoichiometry clearly led to a very significant red shift. They achieved this increase in O content through in-air heat treatment - the same process used in this study. There is potential for a similar effect to be occurring here, where the mixed coatings are absorbing more oxygen through heat treatment, and could even be absorbing more than the desired amount and becoming oxygen rich at the higher heat treatments. Another means of potential oxygen content shift in the crystallising high-n layers could also come from inter-diffusion of elements between the SiO₂ and TiO₂:SiO₂ layer boundaries. Inter-diffusion at adjacent layer boundaries has also been seen previously to lead to stress changes in some multi-layer films [290, 293]. Ultimately either oxygen absorption into the coating or inter-diffusion through heat treatment could be plausible mechanisms of increasing oxygen content in the high-n layers, or inducing stress changes, and potentially explain the significantly red shifted main Raman peak with heat treatment.

One can also observe trends emerge between the point groups to which the optical phonon modes belong (see Table 5.1). The A_{1g} mode is one which undergoes out-of-plane motion, has singular degeneracy, even parity under inversion, and is symmetric with respect to the principle axis of symmetry undergoing out of plane vibrations [294, 295]. The B_{1g} mode conversely undergoes in-plane motion, has singular degeneracy, even parity under inversion, and is anti-symmetric with respect to the principle axis of symmetry [294, 295]. An E_g mode also undergoes in-plane motion has double degeneracy, with a 2 dimensional irreducible representation [294, 295, 296, 297]. The rutile Raman peak at around 247 cm⁻¹ does not coincide with any theoretical calculations for the fundamental modes allowed by symmetry in this phase, and is known to be the compound vibration peak due to the multiple-phonon scattering processes [279, 298].

From Table 5.1, the in-plane vibration E_g and B_{1g} anatase modes have their Raman bands all red shifted after the maximum heat treatment with respect to bulk crystalline powders. (In this observation we exclude the Raman band around 200 cm^{-1} , which is challenging to localise due to its close proximity to a much larger intensity peak.) The only band containing elements of the A_{1g} and out-of-plane motion (516 cm^{-1}) experiences a slight Raman blue shift. Though this band also has elements of B_{1g} motion, the final blue shift observed concurs with previous optical thin film studies by Lee et al., which indicated that the rate of Raman shift evolution in A_{1g} peaks is greater than that of the evolution of peaks manifesting from in-plane vibrations [299]. Red and blue shifting of the peaks are potential indicators of tensile and compressive strain respectively. Opposing red/blue shift trends between E_g and A_{1g} peaks witnessed here in our thin film coatings with respect to crystal powders were also witnessed by Lee. It was also noted that as more layers of the same thin film material were deposited on the same sample, all spectral peak locations eventually converged to the bulk case. However we can not so easily compare our titania-silica mix/silica bi-layer stack coatings to this as in that previous study the same crystalline material was sequentially layered. With each bi-layer of our stack we introduce more pure-silica which remains amorphous, so we would, again, not necessarily expect increasing the number of bi-layers in the titania-silica mix/silica stack to bring convergence to the bulk material properties Lee et al. witnessed.

5.2.2.1 Correction of spurious Raman spectra features

In the 69.5% HR coating 850°C and 950°C spectra a peak at 480 cm^{-1} can also be seen, which does not coincide with any form of crystalline titania: see Figure 5.3. Those two spectra were the first taken on the Glasgow Raman system, and the artifact was later witnessed on all samples measured in Glasgow, but in none of the same samples measured in Lyon, and seemingly originated from the main band of silica as can be seen in Figure 5.1. Given that a Raman spectrum of the same material should in principle be identical regardless of acquisition wavelength, the presence of this ‘spike’ highlighted something was significantly affecting the acquired data. One hypothesis was that a much greater portion of the light field could be transmitting through the thin-film structure (designed for reflectivity at 1064 nm) to the substrate at 532 nm (used for excitation in Glasgow) than at




Figure 5.5: Modelled transmission spectra for wavelength range 400-550 nm, for the 63.2% HR coating.

473 nm (used in Lyon), and hence the silica substrate could have a greater influence on the measurements taken in Glasgow. However after running transmission models using the TFCalc [300] optical modelling software, the transmission to the substrate at these wavelengths (for both coatings) was found to be very similar - $\approx 93\%$ transmission for green light, and $\approx 90\%$ for blue as can be seen in Figure 5.5. This small increase in transmission would not account for the emergence of such a prominent spectral feature in the Glasgow data.

A systematic study was then carried out to identify experimental conditions that could suppress the spike. The study involved testing the effect of focus depth on a sample spectrum, by measuring the crystallisation onset 600°C 63.2% HR sample and adjusting the focus point of the laser beam on/through the sample by $\pm 6\ \mu\text{m}$. In doing this for different focus depths, above, inside, and below the $\approx 5\ \mu\text{m}$ thick coating, the changing influence of the silica substrate could be clearly witnessed in the spectrum between $300\text{-}700\ \text{cm}^{-1}$. The coating signal could be much more clearly resolved when the focus was optimised to coincide better with the coating surface. Figure 5.6 shows the difference of focusing just $3\ \mu\text{m}$ into the coating versus the top surface with an objective lens of $\times 50$ magnification and numerical aperture of 0.5. The stronger influence of the silica substrate




Figure 5.6: Raman spectra acquired using different objectives for the 63.2% Ti HR stack at 600°C heat treatment (left), and for the 69.5% Ti HR stack at 850°C heat treatment (right). For the 63.2% coating, two spectra taken with the 0.5 NA objective are shown; A. focus point 3 μm below the top surface, B. focus point at the top surface.

is clearly seen in spectrum A between Raman shifts of 200-500 cm^{-1} with the laser beam waist focused deeper into the coating, than in spectrum B when it was focused on the top surface. However, even after this optimisation, the spike at 480 cm^{-1} was still observable in the Glasgow measurements.

As the spike was still present and seemed to coincide in wavenumber with the main Raman band of silica, a second investigation was launched in swapping the Glasgow objective lens to one with a higher numerical aperture. All Raman measurements in Glasgow were previously conducted with an objective lens of $\times 50$ magnification and numerical aperture of 0.5, whereas with the Lyon measurements a $\times 100$ magnification 0.9 numerical aperture objective was used. The numerical aperture is a measure of the solid angle of light collection of the microscope objective [301]. A large value numerical aperture corresponds to a greater solid angle. In our case, this quantity has implications for the laser beam profile incident on the sample. Ideally, one would like to fill the back aperture with the laser beam such that it comes to a diffraction limited spot in the field of view [302]. Not all the rays travelling down the barrel of the microscope objective lens will reach it at the

same angle of incidence. Objectives with the greater numerical apertures will cause the rays passing through the lens near the perimeter to refract at the larger angles. Therefore, one can envision a distribution of angles of incidence on the sample that is dependent on the numerical aperture [271]. Specifically the numerical aperture (NA) is defined as [301, 303]

$$\text{NA} = n \sin\theta, \quad (5.4)$$

where n is the refractive index of the immersion medium between the objective and the sample of interest (in our case for air $n \approx 1$), and the θ is maximum light cone angle.

The depth resolution (z) of a confocal microscope (focusing a probe laser of wavelength λ) is inversely proportional to the square of the numerical aperture and given by [301, 303]

$$z = \frac{n\lambda}{\text{NA}^2}. \quad (5.5)$$

From Equation 5.5 it can be calculated that swapping objectives from 0.5 to 0.9 NA should have an effect of reducing the depth resolution of the system by a factor of 3.24, and hence greater suppress substrate influence once focused on the sample top surface - care was also taken to ensure the lenses were cleaned and free of particulate left by other users. Upon acquiring new data for the 600°C 63.2% HR and the 850°C 69.5% HR samples with the $\times 100$, 0.9 NA objective, the spike at 480 cm^{-1} was no longer present, proving the second investigation successful. These results can be seen in Figure 5.6 and comparisons made with Figures 5.3 and 5.4. To show even more explicitly the impact of the objective on thin films, Figure 5.6 also includes Raman spectra recorded using a $\times 10$, 0.25 NA objective typically used only for initial surface focusing. Due to the $\times 13$ greater depth resolution of the 0.25 NA objective versus the 0.9 NA objective, the signal using the former appears to strongly resemble the blank amorphous silica substrate. Using the 0.25 NA, the only distinguishable features of titania being present is the crystalline peak around 140 cm^{-1} , and in the case of the 850°C sample undergoing advanced crystallisation, also the modulation of the spectrum from the 634 cm^{-1} titania peak.

5.3 Titania-silica mix supplementary GI-XRD measurements

From the Raman spectroscopy investigations, the presence of crystallisation in these coatings, and some key aspects of their nature, were characterised. However information on the crystal size was not possible to acquire. Grazing incidence x-ray diffraction (GI-XRD) measurements were later used to quantify this. These measurements performed on samples heat treated to varying final temperatures would also serve as a verification of the crystallisation temperatures found via Raman spectroscopy.

Colleagues at the University of Strathclyde performed GI-XRD measurements on heat treated samples of both HR stacks, and confirmed the crystallisation temperatures for both stacks found in Raman, serving as a further verification of the results already obtained. Through these measurements, it was also discovered that the nature of the crystallisation in both cases was from the formation of titania nano-crystallites in the coatings. The crystallites were initially around 2-3 nm in size at the onset of crystallisation in the respective stacks (575°C and 600°C), but were found to have grown to be between 6-10 nm after 850°C heat treatment.

5.4 Titania-silica mix crystallisation conclusions

These Raman measurements were performed to ascertain whether the HR coatings under investigation had crystallised, and if so, whether the mixing of silica with the titania suppressed crystallisation until higher temperatures. Evidence was witnessed for both, as well as indicators of significant strains and/or oxygen acquisition in the coatings at the higher heat treatments of interest for optimal coating thermal noise. The 69.5% Ti high-n HR coating crystallised at 575°C as anatase, and the 63.2% Ti high-n HR coating crystallised at 600°C as an anatase-rutile mix. This result was not necessarily expected for coatings differing in Ti content by 6%, and, understanding the exact factors that led to it may be of interest and have application for fields beyond gravitational wave observation

astronomy. For instance, studies of mixed crystal phase titania have previously shown that in certain configurations they can yield greater photoactivity than pure anatase and pure rutile [304, 305], which means a mixed crystal titania coating could potentially be of interest in an anti-reflective (AR) configuration for solar cell production.

Further to the results achieved, it would be of interest to study more IBS titania-silica mix ratios in order to better understand how the mixing ratios affect crystal formation and structure. Performing oblique incidence Raman spectroscopy on a modified setup could also yield improved results with the potential to significantly eliminate even more the substrate influence on the spectra. Doing so would allow for much thinner coatings to be well characterised through Raman spectroscopy, as the high incident angle would result in a much shorter depth profile into the samples. This in turn would allow for much more accurate information to be obtained for, say, single layer coatings. Previous measurements on pure IBS titania coatings on silica and silicon substrates with a modified oblique angle technique yielded a 50-fold reduction of substrate influence on the final spectra compared with the standard near-normal incidence measurements [257, 274], and illustrates how implementing this setup configuration change could be very beneficial for measurements of these and other thin film coatings in the future.

Ultimately these results, combined with the studies presented in Chapter 4, led to an interesting outcome. A coating stack previously composed of two fully amorphous materials, had its high-refractive index material undergo advanced stages of crystallisation after heat treatment, yet its optical losses remained relatively low and unchanged. Understanding the underlying physical mechanisms that lead to the continued low levels of absorption and scatter witnessed after the formation and growth of nano-crystallites, should provide key insights for the design of future proposed state-of-the-art optics. This, potentially, has large implications for a plethora of high precision experiments beyond gravitational wave detection, such as those which involve optical atomic clocks.

Chapter 6

Mechanical loss studies of titania-germania coatings

6.1 Introduction

This chapter describes research undertaken at the California Institute of Technology (Caltech) LIGO Laboratory under the supervision of Dr Gabriele Vajente. During this study the author characterised the mechanical losses of numerous single layer and HR stack coated disks comprising of 44% Ti, titania-doped-germania, and amorphous silica deposited onto SiO_2 disk substrates.

A mixed material of 44% titania (TiO_2), and 56% germania (GeO_2) by cation concentration, henceforth referred to as $\text{TiO}_2:\text{GeO}_2$, was identified by Vajente, et al. in 2021 [148] as a highly promising replacement high-refractive index material to pair with pure silica (SiO_2) with potential to meet improved thermal noise targets for future upgrades to the current generation of ambient temperature gravitational wave detectors (GWD). The initial motivation to investigate coatings based on GeO_2 arose from the discovery of a correlation between the fraction of edge-sharing versus corner-sharing polyhedra in the atomic structure of materials and their room-temperature mechanical loss, as reported for $\text{ZrO}_2:\text{Ta}_2\text{O}_5$ in [150]. GeO_2 shares some key similarities with the current room temperature GWD low-index material, SiO_2 . Both are known to exhibit a prevalence of corner-sharing, and SiO_2 is the amorphous oxide which exhibits the lowest known room-temperature loss across the acoustic frequency range [82, 155, 156]. Additionally at cryogenic temperatures (~ 100 K), a peak emerges in the mechanical loss of amorphous GeO_2 [306], similar to the one found in amorphous SiO_2 [307] and indeed other amorphous oxides [308, 309, 310].

Earlier studies by Yang, et al. [151] on pure GeO_2 confirmed that the atomic packing could further be altered to improve the atomic structure via elevated temperature deposition, and post deposition heat treatment of the coatings. As such, it was deemed an interesting material to investigate.

The primary ethos for doping GeO_2 with TiO_2 was so that a mixed material with larger refractive index could be produced, and hence a thinner HR coating stack than if made incorporating pure GeO_2 , could eventually be achieved. Vajente, in collaboration with colleagues at Colorado State University (CSU) and others in the LIGO scientific collaboration, concluded that after trialling numerous cation concentration mixes of single layers of $\text{TiO}_2:\text{GeO}_2$, that 44% Ti doping seemed the most promising composition for reducing coating thermal noise via its improved (lowered) mechanical losses. Specifically, it had optimally low mechanical loss after 600°C , 108 hr post deposition heat treatment, whilst still remaining amorphous [148]. The work undertaken by the author involved characterising the mechanical losses of both new optimised deposition single layer IBS coatings produced by CSU, and characterising various multi-layer stacks incorporating both $\text{TiO}_2:\text{GeO}_2$ and SiO_2 . Alongside these investigations, sister studies of the mechanical losses of the first runs of single layer $\text{TiO}_2:\text{GeO}_2$ and multilayer coatings produced by the GWD coating manufacturer LMA were also initiated.

All samples discussed in the following studies were made from Suprasil® 313 [311] fused- SiO_2 , nominally ~ 75 mm diameter, ~ 1 mm and polished to an optical grade [147]. The mechanical losses of each were characterised in the Caltech GeNS system described in Section 6.2, using the analysis methods discussed in Section 3.4.2–Section 3.4.4 for general GeNS measurement. The samples were measured in their uncoated state, as-deposited, and following varying rounds of post-deposition heat treatment, with fixed heating/cooling rates of $100^\circ\text{C}/\text{hr}$ to, and from, the desired temperature. Each sample underwent multiple GeNS suspensions at each temperature step, with multiple resonant modes from 1–30 kHz measured, multiple times during each suspension, and coating loss calculated. The mean coating loss values acquired for a particular resonant mode, for a particular suspension of the disk, were then compared with those values acquired from the other suspensions in the same post-deposition heat treatment state of the sample in order to acquire the ‘best’ approximation of the loss of the coatings, as in Section 3.4.2.3. For any identical

coatings from the same batch, deposited onto more than one substrate, checks were also carried out to verify their extracted coating losses agreed with each other (an example is shown in Figure 6.8), with the ‘best’ coating losses for the vibrational modes of each sample displayed in all the following results plots in this chapter.

6.2 Comparisons between mechanical loss measurement procedures at Caltech and Glasgow

The mechanical losses of coatings and substrates were measured in much the same way at Caltech as they were at the University of Glasgow. The GeNS system at Caltech had four GeNS platforms upon which disk samples are suspended inside a vacuum chamber, as shown in Figure 6.1. The motion of each disk is measured by a HeNe laser and a quadrant-photodiode. Unlike the process used in Glasgow, a high voltage amplifier is driven with broadband white noise, such that many vibrational modes of the disk are simultaneously excited.

6.2. Comparisons of loss measurement procedures at Caltech and U of G 190





Figure 6.1: (Top) Labelled image of the four GeNS platforms mounted inside a vacuum chamber at Caltech. (Bottom) Topside view of the layout of the four GeNS platforms, showing the paths of the impinging laser beams in red, before and after striking a suspended sample. Image and diagram shared by Vajente with the author and adapted [312].

6.2. Comparisons of loss measurement procedures at Caltech and U of G 191




Figure 6.2: Comparison of the quadrant-photodiode signals taken before and after the excitation is applied to a suspended uncoated silica disk. X and Y denote disk displacement measured with the optical lever readout across the orthogonal vertical and horizontal axes of the quadrant-photodiode respectively. Post-excitation, the X and Y spectra together clearly show all the resonant modes of the disk $\lesssim 30$ kHz. Below 2 kHz the signals are limited mainly by acoustic and seismic disturbances, and above this they are shot noise limited. However, it is clear from the amplitude of the peaks post-excitation that this noise background level is not a limitation for the measurement. The fundamental mode of this disk can be seen at around 1.1 kHz.

In the Caltech GeNS setup no initial mode-search is carried out by performing targeted sweeps through frequency space. The white noise excitation allows for simultaneous measurement of the disk motion induced at all mode frequencies below 32.5 kHz, limited by the sampling frequency of 65 kHz. The resonant modes of the disk are instead identified in the post-excitation frequency spectrum, by comparing it to the spectrum pre-excitation. Figure 6.2 show the pre and post-excitation signals detected by the X and Y channels of one of the quadrant-photodiodes for one such disk, and the excited peaks are clearly apparent.

6.2. Comparisons of loss measurement procedures at Caltech and U of G 192

The ratio of the spectrum post-excitation divided by the spectrum pre-excitation is used to automatically identify all excited resonances: all frequency bins where the ratio of post/pre spectrum is larger than a certain threshold (typically 10) are identified. The evolution with time of the amplitude of each of the identified peaks is then tracked by computing Fourier transforms over time periods which can range from minutes to hours depending on the level of loss of a given sample. The analysis method outlined previously in Section 3.4.2.1 is then applied to the beating amplitude decays witnessed for a given mode of a sample to extract its measured mechanical loss.

There are a few differences with the analysis methods employed across the Caltech and Glasgow setups. With reference to the flow diagram shown in Figure 3.7 of the process for fitting of a beating decay to the acquired amplitude ringdown data at Glasgow, the Caltech procedure adds an additional branch, this is shown in Figure 6.3. Specifically, it also fits a model which forces the two fitted loss values of a degenerate beating pair of modes to be equal, and compares this to fits of the data made assuming both initially have different losses. Between these two, the model with minimised residuals between the fit and the data is taken as the estimate of the measured losses.

This analysis used to be part of the Glasgow analysis code, however it was removed as it was found that the initially non-equal model, after several iterations, would generally converge on loss values of a mode pair being equal, if indeed that was the best fit to the data. This was removed from the Glasgow process to improve computation speed, with thresholds being placed that if the two fitted loss values agreed within 10%, the higher loss value would be overwritten with the lower of the two. Extensive work was previously performed by Glasgow colleagues which confirmed that removing this analysis branch indeed did not significantly impact the fitted results [313].




Figure 6.3: Flow diagram showing the major processes used to perform the optimised fit of Equations 3.13–3.14 to the beating ringdown data at Caltech. This is similar to the analysis procedure for Glasgow GeNS measurements shown in Figure 3.7, but with an additional branch highlighted in yellow.

6.2. Comparisons of loss measurement procedures at Caltech and U of G 194

There are differences in the analysis pipelines implemented between Caltech and Glasgow which pertain to material property selection. At Glasgow, coating material properties, such as the Young's modulus, Poisson ratio, and film density are initially estimated with values from the literature for a given material, or indeed from direct measurements of the material properties where they are available. For a mixed material, the composite averaging method shown in Section 3.4.3.2, is also used. These estimates are then input into a finite element analysis model of a coated disk of specified coating and substrate thicknesses, which was built upon from a previous blank disk model that had modelled geometry and mode frequencies which closely match the measured mode frequencies of that disk while it was blank.

In the Caltech analysis, for each sample, the resonant frequency shifts (δf_i) for every mode resulting after coating the disk are recorded, as are the measured losses of the coated disk $\phi_{M,i}$ and uncoated substrate $\phi_{S,i}$. A model has been developed of coated disks, which can be used to compute the excess loss measured (versus the uncoated substrate) as a function of the material parameters: Young's modulus (Y), Poisson ratio (ν), and density (ρ). This model is based on finite element simulations of the blank substrates, performed with COMSOL Multiphysics [191]. This is similar to the process at Glasgow, except as opposed to building a bespoke model for each disk, the Caltech analysis code uses the observed frequency shifts of the modes, upon coating, to fit the data to a pre-generated library of material properties which would induce this frequency shift, and hence compute the energy ratios/dilution factors for each mode. To assist with this analysis, the coating density and thickness of single layer coatings are independently measured through Rutherford backscattering spectrometry and ellipsometry prior to this analysis, as such only the Young's modulus and Poisson ratio are refined in this process.

Ultimately, though the processes may differ, after optimisation of the material parameters, one would expect to see equivalent mechanical loss results achieved from either analysis pipeline/ measurement. Verifying this consistency is an active area of research between both institutions. Initial checks of uncoated samples have yield promising agreement, however a full comparison of any significant discrepancies (or lack thereof) that arise from a plethora of uncoated and coated samples measured and analysed on both systems is interesting planned work in the near future.

6.3 Mechanical loss studies of LMA titania-germania

An investigation was undertaken into measuring the mechanical losses of single layers of 44% Ti, TiO₂:GeO₂ and pure SiO₂, as well as a coating stack comprising of both, deposited by the Laboratoire des Matériaux Avancés (LMA), in their ion-beam-sputtering (IBS) ‘Grand Coater’ system. This is the same chamber in which current aLIGO/adV coatings are manufactured, and these studies would serve as an important step for verifying the candidacy of TiO₂:GeO₂ as a future GWD coating.

Specifically, two single layer TiO₂:GeO₂ coated GeNS measurement samples, one single layer of SiO₂, and one 20-layer stack of quarter-wave-layer thickness at 1060 nm were made. All samples, save one, were single-side coated, with the first of the two TiO₂:GeO₂ samples being coated with around 500 nm of material on both sides. Details of the coating thicknesses, as well as other deposition parameters for these coatings can be found later in Table 6.1. These coatings were all deposited at 100°C, with the single layers being deposited after the chamber had reached a base pressure of $<8 \times 10^{-6}$ mbar, whereas deposition of the stack began only once the chamber had reached a base pressure of $<5 \times 10^{-7}$ mbar. The higher chamber base pressures for the single layers were selected primarily due to timing requirements for this path-finder study, and to observe if they influenced defect formation in coatings made from this chamber, as a sister study with coatings manufactured by Colorado State University had indicated [314] - see Section 6.4. All samples were measured as-deposited and after three heat treatment stages at 500°C, and 600°C, with dwell times of 10 hr, and a final heat treatment again at 600°C, but for a longer 108 hr dwell.

One unexpected occurrence that arose at the onset of this study, was that the double-sided coated TiO₂:GeO₂ single layer sample was observed to contain a cloudy ring of darker appearing material on one face, encompassing the outermost few mm of the sample. Figure 6.4 shows an image of the sample. Through optical microscopy this ‘halo’ was confirmed to be present on only one of the two sides, and after talks with the manufacturer, and inspecting of its distinct and off-centre orientation, its source was traced to a washer/spacer the sample rested on while in the deposition chamber. When the side coated first with TiO₂:GeO₂ was flipped to be in direct contact with the washer to en-




Figure 6.4: (Left) Image taken of the disk coated with $\text{TiO}_2:\text{GeO}_2$ on both sides under green light after heat treatment at 500°C - the visual difference of the outer most region ('halo' feature) is apparent. (Right) Zoomed in darkfield microscope image taken of the boundary between both regions.

able coating of the second side, their interaction led to the imprinted ring observed. It could not be removed with chemical cleaning with isopropanol and acetone. The halo also became more visually pronounced after the very first heat treatment at 500°C , with the material of the ring appearing to have crystallised (see the right hand side of Figure 6.4).

The coating loss results from the double-side coated sample with the halo were, however, very similar to those from the single side coated sample. Figure 6.5 shows the coating mechanical loss results obtained from the double-side (DS) and single-side (SS) coated $\text{TiO}_2:\text{GeO}_2$ samples, as well as the single-side coated SiO_2 and 20-layer stack samples




Figure 6.5: Graphs of the coating losses measured via GeNS of the four samples, as deposited (left), and after final heat treatment at 600°C for 108 hr (right) of all modes. Clover modes are represented by unfilled markers, and mixed modes with filled markers.

as-deposited and after the final heat treatment at 600°C for 108 hr (replicating the optimum heat treatment for other single layer $\text{TiO}_2:\text{GeO}_2$ coatings deposited by CSU found previously [148]). Mechanical loss results from both the single-side and double-side coated $\text{TiO}_2:\text{GeO}_2$ single layers are broadly equivalent both as-deposited, and especially after final heat treatment, perhaps suggesting that the halo structure does not impinge too deep into the coating, meaning, at most, only a very small relative volume of altered material versus unaltered would be present in the layer.

One can also see that the loss results for the 20-layer stack are consistent with it being composed of the two other materials, given that its losses consistently lie between the two. There are two distinct features that inform us that the $\text{TiO}_2:\text{GeO}_2$ is dominating the loss of the stack over the SiO_2 ; the overall loss level, and the overall loss trends. The loss of the stack is generally closer in value to that of the $\text{TiO}_2:\text{GeO}_2$ single layers, and as the




Figure 6.6: Individual graphs showing the coating losses measured via GeNS of each of the four samples, through all stages of heat treatment separately. Clover modes are represented by unfilled markers, and mixed modes with filled markers.

TiO₂:GeO₂ begins to exhibit a strong frequency dependence in the losses at the final heat treatment (while the SiO₂ losses remain broadly flat) the stack also exhibits a noticeable but weaker frequency dependence. Figure 6.6 shows the losses of all four samples at all heat treatments (i.e. excluding the as-deposited state) with bespoke graph scaling so as to better observe their values and trend, this is particularly useful for the significantly lower loss SiO₂ single layer coating (shown in the bottom left plot of Figure 6.6).

Of great note as well are the single layer SiO₂ losses characterised in this study. The losses were low and near-constant with frequency, and in line with previously deposited LMA SiO₂ from this system after 500°C, 10 hr heat treatment [82]. Large SiO₂ loss improvements following 600°C, 10 hr and 108 hr heat treatments were recorded. After the 600°C, 10 hr heat treatment the losses had fallen to 58% of the previous 500C, 10 hr values of $(3.3\pm 0.7)\times 10^{-5}$ and, further, after 108 hrs at this final temperature they were of the level $(8.0\pm 0.7)\times 10^{-6}$, which were only 24% of the average values found after 500°C, 10 hr - showing a reduction of 76% in the loss of the low index material. This is significant as 500°C, 10 hr is the current heat treatment temperature and duration of the current GWD coatings recipe. Therefore the benefits of heating SiO₂ at this slightly higher temperature for near $\times 11$ greater duration are apparent, and this study serves as a first such characterisation of SiO₂ from the same deposition system as the GWD test masses are coated. Loss results comparable to these final SiO₂ values have been achieved previously from LMA's SPECTOR IBS deposition system, but after the coatings underwent a much higher temperature 900°C heat treatment for a much shorter (standard) duration of 10 hr [165]. These results on the whole show great promise for meeting future GWD noise reduction goals.

The single layer TiO₂:GeO₂ results also show great promise. The current aLIGO/adV coating high-index material is 27% Ti, TiO₂:Ta₂O₅ made in the same chamber also exhibits a similar frequency dependence in the loss when measured via GeNS. The TiO₂:GeO₂ results at around the same lowest measurement frequencies of these TiO₂:Ta₂O₅ studies (~ 1.1 kHz) have almost a factor of $\times 4$ lower loss than the TiO₂:Ta₂O₅ [82] and have a decreasing trend with reducing frequency. Performing a simple power law fit to the single layer TiO₂:GeO₂ data where $\phi_c = af^b$, using linear regression methods similar to as was done for the current TiO₂:Ta₂O₅ [82] one extracts a and b values of around 3.2×10^{-5} Hz^{-b}

and 0.13 respectively. This fit would produce losses in the 100 Hz range where CTN is dominant, also of around a factor of $\times 4$ less than $\text{TiO}_2:\text{Ta}_2\text{O}_5$ [149] with only a $\sim 10\%$ reduction in refractive index. Further to this, these single layer results can be compared directly to values acquired by Vajente et al. in [148] for the average loss of all modes measured between 1–30 kHz. These fully heat treated $\text{TiO}_2:\text{GeO}_2$ single layer samples investigated in the study have an average loss of $(1.1 \pm 0.2) \times 10^{-4}$ compared with previous $\text{TiO}_2:\text{GeO}_2$ single layer coatings which had an average loss of $(1.0 \pm 0.2) \times 10^{-4}$, where these error ranges represent the standard deviation from all measured loss values. Note that the frequency dependence of the losses appears significant so simple averaging like this will not take this into account.

No sample exhibited coating cracking at any point in this study, but both $\text{TiO}_2:\text{GeO}_2$ single layer coatings were found to grow a few blisters after heat treatment at 600°C , for 10 hr. On the other hand the SiO_2 single layer, and the stack containing both materials, but deposited at a lower base pressure, did not. This suggests that the significantly lower base pressure used during the stack deposition assisted in suppression of the stresses that form these kinds of defects. However this stack sample did begin to exhibit small sites of what appeared to be delamination in various regions of the coating, suggesting further adjustment of the deposition parameters is needed to completely suppress defect formation [314]. Regardless of these defects, the losses still continued to fall after heat treatment at this and the subsequent heat treatment step.

6.4 Studies of titania-germania single layers made via different processes

Along with the study of the LMA coatings, loss studies of coatings deposited in CSU's SPECTOR IBS system were also undertaken. CSU had found after many iterations of their $\text{TiO}_2:\text{GeO}_2$ coatings, that they were able to prevent blistering and cracking in the material. Unlike many of their own previous runs [148], and those of the runs from LMA discussed in the previous section, these CSU coatings were deposited only after the chamber reached an elevated temperature of 200°C , and base pressure of $< 1.6 \times 10^{-7}$ mbar were achieved.




Figure 6.7: Comparison of the coating losses measured via GeNS of the different $\text{TiO}_2\text{:GeO}_2$ single layer coatings with different deposition parameters, as deposited (left), and after heat treatment at 600°C for 10 hr (right). Clover modes are represented by unfilled markers, and mixed modes with filled markers. Deposition details can be found in Table 6.1.

This single layer material (CSU) was compared with both the previously discussed run from LMA (LMA-1), and a new run from LMA (LMA-2). The new LMA run, informed by the CSU studies, also trialed an elevated temperature deposition, but at the maximum chamber temperature the Grand Coater could reach of 150°C . Details of parameters for these coating runs, along with the previous coatings are shown in Table 6.1, and the best loss results for all three coatings as-deposited, and after 600°C , 10 hour heat treatment are shown in Figure 6.7.

From the as-deposited and 600°C , 10 hr results it would appear that both LMA, and the CSU coating all have comparable losses. As-deposited LMA-2 has lower losses than LMA-1, which is consistent with the 50°C elevated temperature producing amorphous coatings of more ordered structure and consequently lower loss [151, 315, 316]. The higher temperature coating ions will have more energy to move about the substrate surface during deposition and as such over time a more ordered structure closer to an ideal-glass forms. Upon heat treatment to (the much higher temperature of) 600°C both of the LMA coating losses can be seen fall even further in line with each other. On the other hand, these CSU coatings appear to have slightly higher losses as-deposited, perhaps appearing to break the trend of elevated temperature deposition reducing the loss. However, care must be taken when comparing coatings from different systems, as features such as slight difference in chamber cleanliness, or ion source, can lead to significantly different losses [208]. Though, in counter to this, the CSU coatings were deposited with approximately fifty times lower




Figure 6.8: Comparison of the coating losses measured via GeNS of the different $\text{TiO}_2\text{:GeO}_2$ single layer coatings with different deposition parameters, that were able to be annealed at the final step of 600°C for 108 hr - i.e. LMA-2 run is excluded. Clover modes are represented by unfilled markers, and mixed modes with filled markers. Deposition details can be found in Table 6.1.

base pressure than the LMA single layer runs so in theory the environment in the CSU chamber may have less particulates. It would subsequently be discovered that the CSU coatings were deposited less stoichiometric than expected, and this could also contribute to the observed higher as-deposited losses.

The CSU coating, after 600°C 10 hr heat treatment, appears to align with the loss of the LMA coatings above 5 kHz, but seems to possess slightly higher losses at lower frequency. However, unlike in the as deposited case, these slightly higher losses at low frequencies potentially result from not finding the true loss level during the two 600°C , 10 hr measurements for the CSU coated samples. The primary evidence supporting this is that when the subsequent 600°C , 108 hr heat treatment was performed, the coating losses of LMA-1 and CSU both actually yield virtually equivalent results across the entire measured frequency band - see Figure 6.8.




Figure 6.9: Image showing the haziness the LMA-2 single layer sample developed some time after the 600°C, 10 hr measurement, as a consequence of absorbing water from its surroundings.

Unfortunately, the LMA-2 $\text{TiO}_2:\text{GeO}_2$ coating was observed to turn hazy in the time between the 600°C, 10 hr step and the planned 600°C, 108 hr step so was not further heat treated - see Figure 6.9. $\text{TiO}_2:\text{GeO}_2$ is hydrophilic, and as such is known to absorb water from the surrounding atmosphere at its surface, resulting in a change of the optical and mechanical properties [147]. However, this is not as large a concern for a multi-layer HR coating stack capped by a SiO_2 half-wave layer. To combat this in the $\text{TiO}_2:\text{GeO}_2$ single layers, the coatings, between, and after, measurement are stored in vacuum sealed containers. However the LMA-2 sample still became contaminated. Poor vacuum sealing of this particular sample when stored, or perhaps the presence of excess water in the chamber after alignment, potentially as a result of a small spillage from the alignment water dish, and then spread through the chamber when the vacuum pump was turned on, could be potential causes. However as its 600°C, 10 hr losses align with the first LMA single layer of $\text{TiO}_2:\text{GeO}_2$, it is likely that its 108 hr results would not have differed significantly from those of LMA-1 presented in Figures 6.5–6.6.

At the final heat treatment step, the CSU sample losses align with the LMA-1 results within measurement spread in most cases. The argument could also be made that at the lowest recorded frequency ~ 1.1 kHz the CSU coatings have slightly better losses than the LMA-1 coatings. The losses from both CSU $\text{TiO}_2:\text{GeO}_2$ coated samples are plotted separately in Figure 6.8, to show how well they agree with each other, and is representative of the general agreement seen between all GeNS samples coated with the same coating

6.4. Studies of titania-germania single layers made via different processes 203

across all 21 disks for which measurements are presented in this chapter. Notably, however, there is a slight but observable disagreement between the measured losses at ~ 1.1 kHz from both CSU disks, which affects the trend of the loss/frequency slope, which would ultimately influence the extracting of losses at lower frequency.

Aside from the hazing, the LMA-2 samples were observed to manifest no defects after the 600°C , 10 hr heat treatment. The CSU samples also showed no blistering or cracking, making these the most promising future GWD detector candidates in this study, as they achieved their extremely low loss results free of defect. Performing a simple power law fit to the data from both CSU coated disks separately after 108 hr heat treatment yields estimated loss levels at around 100 Hz of only $(18.3 \pm 1.7)\%$ of the current GWD high-index material $\text{TiO}_2:\text{Ta}_2\text{O}_5$. The fact that these results are achieved alongside significantly improved low-index SiO_2 levels shown in Figure 6.6 when compared with current 500°C , 10 hr heat treatment, highlights the great prospects for these materials after 600°C , 108 hr heat treatment for future detector coating designs. However, to produce comparable reflectivity, these boons would be mitigated by requiring a stack of overall $\sim 46\%$ greater thickness (see Equation 2.6) than the current ETM GWD coating, as many more bi-layers (26 total) would be required, with each having high-index layers of roughly $\sim 10\%$ thicker than $\text{TiO}_2:\text{Ta}_2\text{O}_5$. Still, the prospects of significantly improved CTN results these materials could yield is apparent from the greater relative reduction in loss versus required increase in thickness. From Equation 2.31 one could roughly estimate a level of improvement in the amplitude spectra density of the CTN of these coatings, as the loss of the stack is dominated by the high- n layer losses, of $\delta\text{CTN} \approx \sqrt{0.18 \times 1.46} \approx 50\%$. This estimate is very much in line with previous predictions for $\text{TiO}_2:\text{GeO}_2$ [148] and is also neglecting the fact that the SiO_2 loss was shown here to improve by around an order of magnitude with this same heat treatment. This work has thus shown great promise for advances in future detector coatings, and the ability of these coatings to achieve these results whilst free from defects.

Table 6.1: List of all single layer coatings and one stack design measured in the initial TiO₂:GeO₂ investigations of this thesis. The sample ID, materials, number of layers, thicknesses, deposition temperature T_{DEP} , and chamber base pressure during deposition p are shown. The final column also lists the defects which arose through heat treatment, if any, with a note of any contamination during or after deposition in brackets.

| sample ID | N ^o of layers | material | thickness [nm] | T_{DEP} [°C] | p [mbar] | defects |
|--|--------------------------|------------------------------------|----------------|-----------------------|--------------------|-----------------|
| LMA-1 st coating runs | | | | | | |
| SS-SiO ₂ | 1 | SiO ₂ | 701 | 100 | 8×10^{-6} | none |
| 20L-stack | 20 | both | [1410+1830] | 100 | 5×10^{-7} | delamination |
| DS-TiO ₂ :GeO ₂ | 2 | TiO ₂ :GeO ₂ | [501/498] | 100 | 8×10^{-6} | blisters (halo) |
| SS-TiO ₂ :GeO ₂ | 1 | TiO ₂ :GeO ₂ | 564 | 100 | 8×10^{-6} | blisters |
| LMA-2 nd coating run | | | | | | |
| LMA-2-TiO ₂ :GeO ₂ | 1 | TiO ₂ :GeO ₂ | 539 | 150 | 8×10^{-6} | none (hazing) |
| CSU single layer coating run | | | | | | |
| CSU-TiO ₂ :GeO ₂ | 1 | TiO ₂ :GeO ₂ | 483 | 200 | 2×10^{-7} | none |

6.5 Mechanical loss studies of titania-germania HR stacks

As well as the aforementioned studies in characterising single layer coatings from LMA and CSU, and one stack comprising of 10-quarter-wave-bi-layers of TiO₂:GeO₂ and SiO₂, the author investigated the mechanical losses of two HR stack coating designs of these two materials. Specifically, two stacks, meeting the GWD end test mass (ETM) coating reflectivity requirement, deposited in CSU's IBS system were studied. These stacks comprised of 52 alternating layers of TiO₂:GeO₂ and SiO₂. The stacks had deposition parameters identical to the CSU TiO₂:GeO₂ discussed in Section 6.4, optimised in an attempt to prevent blister/cracking formation upon heat treatment.

Each stack was designed for use at 1064 nm, with total thicknesses of $\sim 8.6 \mu\text{m}$. The final deposited layer of each stack was SiO₂, so as to protect the TiO₂:GeO₂ from exposure to the lab environment, as has been found to be detrimental previously (see Figure 6.9). This final layer is also a near-half-wave of silica so as to not reduce reflectivity. Besides these similarities, the two topologies are quite different. The first stack, dubbed 52L-A,

had layers of constant, near quarter-wave (QWL) thickness, though in the design phase had undergone thickness/material optimisation, similar to as is done for current GWD coatings [82], where the $\text{TiO}_2\text{:GeO}_2$ and SiO_2 layer thicknesses were approximately 11% thinner and 7% thicker than 1064 nm QWLs respectively. One benefit of this optimised thickness design is that it should reduce the total coating mechanical loss/CTN of the stack when compared with a stack comprising of QWLs, as these are dominated by the higher-loss $\text{TiO}_2\text{:GeO}_2$ - see Section 6.3.

The second stack, dubbed 52L-B, was designed to produce the same level of reflectivity of the first stack, but with reduced optical absorption. Research into IBS- $\text{TiO}_2\text{:GeO}_2$ at present has shown it to have higher absorption at 1064 nm than SiO_2 made under the same conditions [314], and as most of the impinging laser light is reflected back from the top few bi-layers, with less and less light present with depth into the stack, this design sought to decrease the relative thickness of the top-most $\text{TiO}_2\text{:GeO}_2$ layers. As such it has a gradient of increasing $\text{TiO}_2\text{:GeO}_2$ and decreasing SiO_2 layer thickness the deeper into the stack. The success of this design has recently been shown by colleagues at Stanford University, to achieve absorption at 1064 nm of around 3.3 ppm, and 3.8 ppm after 600°C 10 hour, and 100 hour heat treatment respectively [314, 317]. These absorption results for 52L-B were directly compared with those of 52L-A, and were seen to be approximately ~40%-50% lower than the stack of near constant layer thicknesses [314, 317]. Visualisations of each stack design are shown in Figures 6.10–6.11.




Figure 6.10: Layer structure of the 52L-A coating designed for high reflectivity at 1064 nm. One can observe the near constant layer thickness of ~ 125 nm $\text{TiO}_2\text{:GeO}_2$, and ~ 197 nm of SiO_2 , plus a half-wave SiO_2 cap layer at the top of the stack.




Figure 6.11: Layer structure of the 52L-B coating designed for high reflectivity at 1064 nm, with reduced optical absorption [317]. One can observe the decreasing $\text{TiO}_2\text{:GeO}_2$ thickness towards the top of the stack.

6.5.1 Titania-germania HR stack loss results

Two SiO₂ disk samples (75 mm diameter by 1 mm thick) were coated with coating 52L-A, and one disk was coated with 52L-B. The mechanical losses of each were characterised in the Caltech GeNS system described in Section 6.2, using the analysis methods found here and delved into at greater depth previously through Section 3.4.2–Section 3.4.4 for general GeNS measurement. The samples were measured in their uncoated state, as-deposited, and following 600°C, 108 hr dwell post-deposition heat treatment. For the two 52L-A samples, it was verified that they produced the same losses as each other within measurement spread, with the best loss values being presented in the following results.

Figure 6.12 shows the coating mechanical loss results obtained for both HR stack designs as-deposited and after post-deposition heat treatment (modelling each as a composite layer for extracting the energy ratios as in Section 3.4.3.1). Both stacks were found to have identical coating losses to each other before and after heat treatment, with losses at the final (600°C, 108 hr) heat treatment reducing by around a factor of $\times 6$ for most modes when compared with the as-deposited numbers. The high level of agreement between 52L-A and 52L-B was expected given that, by thickness/volume, they have identical amounts of TiO₂:GeO₂ and SiO₂ (to within $\sim 1\%$) made under the same conditions in the same system - see Table 6.2 later.




Figure 6.12: Coating mechanical losses of the two TiO₂:GeO₂/SiO₂ HR stacks measured as-deposited (left) and after heat treatment at 600°C, for 108 hrs (right). The losses of clover and mixed modes are represented with hollow and filled markers respectively.

Generally, as-deposited, the coating losses are broadly flat with frequency, with values around $\sim 3.2 \times 10^{-4}$. Here only modes up to 20 kHz were sampled instead of the standard 30 kHz due to an error with the measurement software. Though with the observed flat loss trend through frequency, not sampling the higher frequencies should not be detrimental.

Upon heat treatment, above 15 kHz the HR stack losses are also observed to be broadly flat with frequency at a level of $\sim 5.2 \times 10^{-5}$. However, as shown in Figure 6.12, at lower frequencies the mixed modes appear to exhibit a significant frequency dependence, with their losses decreasing with frequency up to 10–15 kHz. On the other-hand, the clover modes - which have a lower proportion of bulk strain energy stored during vibration than the mixed modes - still exhibit broadly flat losses across frequency. Previously (see Figure 6.6) the 20-layer stack deposited at LMA showed similar results at this heat treatment for the mixed modes, but to a much lesser extent. However, the generally flat trend of the loss of the clover modes below 10 kHz seen in these 52-layer coatings was not witnessed on this other stack coating and this was quite a surprising result given that:

- (A) The 20-layer stack deposited by LMA conversely showed a significant increase in the losses of these clover modes from 1 kHz–10 kHz at comparable heat treatment following the trend of the LMA single layer $\text{TiO}_2:\text{GeO}_2$ coatings - Figure 6.6.
- (B) The 52-layer stacks should be dominated by the loss of the $\text{TiO}_2:\text{GeO}_2$ layers. Single layers of $\text{TiO}_2:\text{GeO}_2$ made with the same recipe, in the same coater were shown to exhibit a large frequency dependence in the clover modes at this heat treatment. However, this is not observed in these 52 layer stacks.

As the previous 20-layer stack and single layer results also show that the heat treated losses from both mode families generally converge at higher frequencies, these HR stack results potentially suggest the presence of a mechanism leading to reduction of loss improvement upon heat treatment below 10 kHz. Indeed, for these HR coatings, both the slope of the clover modes previously witnessed has vanished, and the only mixed mode with frequency < 10 kHz, has proportionally higher loss than witnessed on any prior samples. This is a significant result, because if the mechanical losses of the coating are higher than expected below 10 kHz, it would suggest that these HR designs could produce higher CTN than

previous estimates of around 45% of current detector coatings [148], and no longer meet the target of 50% CTN improvement. Indeed, colleagues at MIT would later directly measure the CTN of these coatings to be higher [318] than estimated from previous CSU single layer predictions [148], which will be discussed further in Section 6.7.

These 52-layer coatings are more than double the thickness of the previously measured stack, with more than double the number of interfaces. Either could be a potential source for less loss reduction on heat treatment. Generally, thicker coatings can lead to greater deformation of the substrate on which they are coated, which can lead to excess losses on single-side coated samples, versus those coated on both sides to counteract this deformation [149]. Since, the substrates in this study are much thicker and with comparable coating thickness (ratio of HR coating/substrate thickness of $\sim 1/120^{\text{th}}$) making the deformation if any likely quite small.

Increasing the number of interfaces could lead to more sites of potential interdiffusion, which is a process of diffusional exchange of atoms across two materials that are in contact, and is driven by the chemical potential gradient existing across the boundaries [319]. The present Ti, Ge, O, and Si atoms (as well as any impurities such as Ar, or C) bonding in different configurations than intended at and near the interfaces, would result in small amounts of undesired material, with potentially much higher loss. Interdiffusion is routinely observed when bonded dissimilar materials are put to high temperature, such as with thermal barrier coatings (e.g. in turbine materials), and is also observed to emerge in as-deposited multilayer coatings due to deposition conditions [319]. The latter is less likely to be the case here given that the loss trends in the 20-layer and 52-layer stacks, made in different systems with different parameters, are similar as-deposited. In the literature no strong evidence of significant interface losses between layers has been found for GWD detector IBS coatings [149, 172], but that does not mean it is necessarily negligible in $\text{TiO}_2\text{:GeO}_2/\text{SiO}_2$ based stacks upon heat treatment.

Towards verifying these potential sources leading to a reduction in loss improvement on heat treatment, the author assisted LIGO colleagues in devising a study to observe potential changes in mechanical loss with the number of interfaces or with the total thickness of the coating of many different $\text{TiO}_2\text{:GeO}_2/\text{SiO}_2$ stack topologies made under identical conditions to each other. The results of this study acquired by the author and Vajente are described in the next section.

6.6 Various topology stacks excess loss study

In order to investigate any effects on the coating mechanical loss of changing the number of interfaces and total thickness, a study was devised in collaboration with Caltech and CSU, where multiple different stack topologies were deposited onto fused- SiO_2 disk substrates. Specifically, five different stack designs of nominally equivalent total thickness, thinner than the 52-layer HR coatings, were made with varying numbers of interfaces, and differing volume ratios of the two materials.

Table 6.2: List of all $\text{TiO}_2\text{:GeO}_2/\text{SiO}_2$ coating stack designs explored in this study, showing their number of layers, total thicknesses, average individual material layer thicknesses, and the total coating volume comprised of SiO_2 expressed as a percentage. Thicknesses were measured via spectrophotometry. For 52L-B's structure see Figure 6.11.

| ID | N ^o of layers | plot marker | total stack thickness (d_T) | layer structure | $d_{\text{Ti:GeO}_2}$ | d_{SiO_2} | volumetric SiO_2 % |
|-------|--------------------------|-------------|---------------------------------|-----------------|-----------------------|--------------------|-----------------------------|
| 2L | 2 | ○ | ~1.6 μm | constant | ~808 nm | ~800 nm | ~50% |
| 8L-A | 8 | △ | ~1.6 μm | constant | ~205 nm | ~190 nm | ~48% |
| 16L | 16 | □ | ~1.6 μm | constant | ~97 nm | ~103 nm | ~52% |
| 8L-B | 8 | ▽ | ~1.7 μm | constant | ~106 nm | ~309 nm | ~74% |
| 10L | 10 | ◇ | ~1.7 μm | constant | ~137 nm | ~201 nm | ~60% |
| 52L-A | 52 | ○ | ~8.6 μm | constant | ~125 nm | ~197 nm | ~62% |
| 52L-B | 52 | ○ | ~8.6 μm | varying | varying | varying | ~63% |

All of these new stacks were designed to be $\sim 1.6\ \mu\text{m}$ thick, i.e. $\times 5.4$ thinner than the 52-layer coatings, with constant layer thickness. In order to investigate the effect of potential excess interface losses, three stacks of near-equal total thicknesses/volume of both materials present, but with increasing number of interfaces; 2-layers, 8-layers, and 16-layers were made. A second 8-layer stack but with instead SiO_2 layers $\times 3$ thicker than the $\text{TiO}_2:\text{GeO}_2$ layers, with nominally the same total stack thickness as the first was also deposited, in order to better perceive how the volumetric split of the two materials influences loss trends in the stacks. Finally a 10-layer stack of nominally equivalent individual layer thicknesses/material volume split as 52L-A (which is also the same material volumetric split of 52L-B) was made, and its losses characterised, in order to be compared against them. The actual layer thicknesses were confirmed for each coating via spectrophotometry (with measurements also being made via ellipsometry for the stacks with < 10 layers) by collaborators at CSU. The details of all stack topologies are shown in Table 6.2.

For each of the five new coatings, two GeNS samples were coated per run, with each giving broadly consistent and equivalent results to the other, with the best losses obtained (see Section 3.4.2.3) presented in the following for each measured mode.

6.6.1 Comparing two 8-layer stacks of different material content

Firstly, to view the influence in the stack loss trends the two individual materials exhibit, the coating losses were quantified for the two 8-layer stacks before and after heat treatment - these results are shown in Figure 6.13.

For the single layer coatings characterised previously in Section 6.3, the SiO_2 generally exhibits a flat trend with frequency both as-deposited and heat treated, whereas the $\text{TiO}_2:\text{GeO}_2$ as-deposited had been seen to exhibit a slight positive frequency dependence in its loss, though is also mostly flat with frequency. These qualities are reflected in these 8-layer stack results, with 8L-B, which is comprised of 74% SiO_2 by volume, having flat as-deposited losses. Whereas 8L-A, which has almost double the amount of $\text{TiO}_2:\text{GeO}_2$ within the same total thickness, exhibits a slight increasing loss trend with frequency.




Figure 6.13: Coating mechanical losses of the two $\text{TiO}_2\text{:GeO}_2/\text{SiO}_2$ 8 layer stacks measured as-deposited (left) and after heat treatment at 600°C , for 108 hrs (right). The losses of clover and mixed modes are represented with hollow and filled markers respectively. The percentages denote the percentage of volumetric SiO_2 in the coating.

However the influence of the $\text{TiO}_2\text{:GeO}_2$ is most apparent after heat treatment, where much stronger frequency and mode family dependence's of the losses of 8L-A are witnessed than the 8L-B coating, suggesting this dependence comes primarily from the $\text{TiO}_2\text{:GeO}_2$. This was expected given the trends seen previously in Sections 6.3–6.4 for the single layer coatings, which are reminiscent of these 8L results. Indeed another general point is that the losses of 8L-A are all higher than 8L-B, which again is consistent with the $\text{TiO}_2\text{:GeO}_2$ being a higher loss material than the SiO_2 , as seen previously.

With the losses of these two stacks characterised and compared with each other and with previous single layer results, the ability to observe the general influences of the relative volumetric material split in a stack was confirmed. This comparison can be made with isolation from varying interface and thicknesses, given that both were (by-and-large) the same for each 8-layer coating. Also, unlike the 52-layer coatings, both of these 8-layer coatings exhibit increasing frequency dependence in their losses after heat treatment from the low to high frequencies, even though 8L-B has proportionally lower $\text{TiO}_2\text{:GeO}_2$ volumetric content than both the 52L stacks made with the same process which exhibit flat loss trends through frequency. This, in-tandem with the previous analysis, ultimately suggests that the losses measured at low frequencies for the 52-layer coatings had not improved as well as one might have expected after heat treatment.

6.6.2 Investigating the effects of increasing interfaces on loss

In order to determine if the number of interfaces led to excess losses for stacks comprised of the 44% Ti, $\text{TiO}_2\text{:GeO}_2$ and SiO_2 , results from three different stacks were compared. Alongside coating 8L-A characterised in the last section, a 2-layer and a 16-layer stack (dubbed 2L and 16L respectively) were studied. All three of these $\sim 1.6\ \mu\text{m}$ stacks were measured to have the same total thickness within a 28 nm (i.e. $<2\%$) spread, with approximately the same $\text{TiO}_2\text{:GeO}_2$ to SiO_2 ratio of $(50\pm 2)\%$. Therefore a comparison of the loss results from these coatings before and after heat treatment would yield insights into potential interface losses, isolated from both thickness and composition.

Figure 6.14 shows the mechanical loss results obtained from these coatings. Interestingly, both as-deposited and after heat treatment a slight, but measurable, increase in loss with number of interfaces is observed. As-deposited, this trend is seen across the whole frequency space, with the 16L losses being generally larger than the 2L ones, with the 8L-A losses being in between these (or in some cases equivalent to the 16L or 2L losses within error). As the deposition parameters were nominally identical, between these runs, and the other topology parameters were kept the same, this suggested some small residual loss related to the number of interfaces.




Figure 6.14: Coating mechanical losses of the three $\text{TiO}_2\text{:GeO}_2\text{/SiO}_2$ of identical thickness and relative material contents, but with varying number of interfaces. The as-deposited and post heat treatment at 600°C , for 108 hrs results are shown on the left and right respectively. The losses of clover and mixed modes are represented with hollow and filled markers respectively.

Upon heat treatment, this effect is observed to still be present to an extent, though above 10 kHz the losses of the 16L and 8L-A samples become indistinguishable within error. For all but one of the 2L clover modes (represented by the hollow black circles), the losses were still measured to be slightly lower than their 16L and 8L-A counterparts. The only one that did not was the 12.8 kHz (0,7) mode which was observed to have slightly uncharacteristically higher loss than the trend of the other modes would suggest in its two heat treated suspensions. It is believed this slightly high measured loss of this mode arose as a result of imperfect suspension/sample preparation and its true loss is probably slightly lower, and could likely be resolved with more suspensions/ sample cleaning.

Overall, the effect on the loss of increasing the number of interfaces seems quite small, and proportionally does not appear to change greatly with heat treatment, with most losses of the 16L modes measured to be, at worse, $\sim 10\%$ higher than their 2L counterparts both before and after heat treatment. The final heat treatment temperatures are significantly ($\times 3$) higher than the deposition temperatures, providing much more energy for atomic interdiffusion, and yet no great evolution in the loss differences between the three stacks is witnessed on heat treatment, so this does not appear to be the likely culprit for the disparity. The overall substrate deformation due to stress from the coatings should also be comparable for each of these three stacks. Perhaps then, the slight excess losses witnessed, unchanging on heat treatment, is brought about by small amounts of contamination in the process from switching between ion targets in chamber for coating successive layers. Indeed, in any case the effect seems minimal, even at the lower frequencies where it appears to be the most prominent. The 16L stack has $\times 4$ the number of interfaces as the 2L, as such the disparity in loss trends at low frequencies seen previously with the 52-layer CSU HR coatings compared with these and the previous 20-layer LMA stack does not seem likely to be caused purely by an interface effect. In all cases for the 2L, 8L-A and 16L stacks the loss trends through frequency of the form also witnessed in the single layer $\text{TiO}_2:\text{GeO}_2$ can still be resolved, with the much thicker full reflectivity HR stacks being the only ones in these studies to exhibit far more uniform loss trends through frequency upon heat treatment - see Figure 6.12.

However, as a noticeable but small effect is still perceived from differing numbers of interfaces, a further study into characterising the loss of, say, a 52-layer stack of ~ 30 nm thick layers of each material, such that its total thickness is the same as the three presented here could yield further insights. Equally a 4-layer stack of comparable total thickness and material split could aid in quantifying the effect of interfaces. Concurrently, also characterising a stack comprising of 8, $\sim 1.1 \mu\text{m}$ thick layers, such that its thickness is comparable to the HR coatings could yield further information for disentangling the effect of the number of interfaces versus the total thicknesses actually required for full GWD reflectivity on the loss, and hence CTN. As these three coatings, post heat treatment, can be seen to have lower losses at lower frequencies, still following the trends of single layers, returning a few of these samples back to chamber for coating of a few sets of much thicker layers, making the re-coated samples have total thickness of $\sim 8.6 \mu\text{m}$, with minimal change in interfaces, could be trialled to see if they lose their lower loss with lower frequency trend when more layers are coated on top (after a subsequent round of post-deposition heat treatment). This could also help determine whether the less improved 52L HR stack losses at lower temperatures result primarily from an increase in overall bulk material.

6.6.3 Comparisons of the HR coating losses to a nominally identical material-split 10-layer stack

To complement the previous results, a 10-layer stack (10L) of similar thickness to the 2L, 8L-A, 8L-B, and 16L coatings, but with relative $\text{TiO}_2:\text{GeO}_2/\text{SiO}_2$ volumetric split similar to the 52L-A, and 52L-B HR coatings was produced and characterised. The loss results for this in comparison to these HR stacks and 8L-B of the highest SiO_2 content are displayed in Figure 6.15 both as-deposited and after heat treatment, with the heat treated clover and mixed mode trends also plotted separately in Figure 6.16 to aid in visual comparisons.

As expected from the analysis of the other stacks, the 10L, 52L-A, and 52L-B stacks were observed to have identical losses as-deposited across the whole frequency range, and upon heat treatment the thicker stacks with more layers, showed less loss improvement below 15 kHz than the 10L coating. After heat treatment the thinner 10L stack showed signifi-




Figure 6.15: Coating mechanical losses of the two 52 layer $\text{TiO}_2\text{:GeO}_2/\text{SiO}_2$ stacks, and a thinner 10 layer stack containing nominally identical proportions of the two materials, and the 8 layer stack (8L-B) with fractionally more SiO_2 . These were measured as-deposited (left) and after heat treatment at 600°C , for 108 hrs (right). The losses of clover and mixed modes are represented with hollow and filled markers respectively.

antly improved losses versus 52L-A, and 52L-B below 15 kHz, with its losses flattening to be in line with them beyond these frequencies. Though these results were predicted from the studies in the last two sections, confirming it served as an important step in assessing the apparent detrimental-loss effect observed in the full HR stack. The heat treated loss of the 10L coating at the lowest measured frequency of ~ 1.1 kHz is around $\sim 30\%$ – 32% lower than the losses of the 52 layer coatings, with the gap seeming to further widen towards lower frequencies. Along with these, the 8L-B coating (with more SiO_2) results are plotted to show that the material volumetric composition still dominates overall.

For the HR stacks, it is apparent that the clover modes exhibit a relatively flat trend over the entire measured frequency band, and the mixed mode losses decrease with frequency. The inverse is seen on the 10L stack with the mixed modes seemingly possessing flat loss with frequency and the clover losses increasing with frequency - all these results can be seen in Figure 6.16. On subsequent analysis of the fractional strain energies stored in bulk versus shear in the coating for each mode, it was found that the clovers in all three stacks had similar proportions of bulk energy, but the bulk energy in the mixed modes had all proportionally increased in the HR stacks with respect to the 10L. This increase in bulk strain-energy storage in the 52L coatings was around 5% versus the 10L coating for all mixed modes, except the first mixed mode at ~ 6 kHz where a much larger increase of 17% in bulk energy storage was found.




Figure 6.16: Coating mechanical losses of the two 52-layer stacks and 10-layer stack of similar material splits, at the final heat treatment, with the clover mode losses plotted on the left and the mixed mode losses plotted on the right.

Further analysis was carried out in finding the relative loss improvement of the 10L stack through frequency. Upon grouping the modes of different families appearing at similar frequencies, one finds the measured loss improvement is consistent between the families. For example, both the clover and mixed modes of the 10L stack at ~ 6 kHz exhibit loss improvement of 17%, and similarly, both clover and mixed modes at around ~ 20 kHz, and ~ 29 kHz each respectively exhibit 7% and 8% improvement when compared with 52L-B. This suggested that perhaps the level of relative loss improvement with frequency was independent of proportion of bulk/shear strain energy storage in a given mode.

6.7 Titania-germania HR CTN estimates

Ultimately it is desired to extract CTN values from these loss results for the full HR stacks, to gauge the level of improvement they would yield in a GWD (at around 100 Hz). As discussed, the heat treated loss trends of these stacks are not straightforward. Ideally one would want to quantify the frequency dependence of the bulk and shear losses for both the constituent materials of the stack, and use a CTN model such as the effective medium (Equation 2.45) to estimate the CTN. In the absence of heat treated single layer SiO_2 data deposited by the same process used for the SiO_2 layers in the HR stacks, this becomes more

challenging. As shown in Chapter 3 for $\text{TiO}_2\text{:SiO}_2$ based HR coatings, good agreement can be reached with direct CTN measurements when the losses of the SiO_2 layers are not known. To do this previously the CTN model proposed by Yam (Equation 2.41) was adopted, making the assumption that the coating has equal bulk and shear losses.

This study on $\text{TiO}_2\text{:GeO}_2\text{/SiO}_2$ showed that the HR stacks exhibit quite different loss trends with frequency than thinner stacks (Figure 6.16), and that the loss of the full HR stacks improve less on heat treatment at lower frequencies than the thin stacks do. Inspecting the post heat treatment losses of modes belonging to different families (i.e. with different bulk/shear energy splits) at similar frequencies between the thinner and thicker stacks, it was seen that the difference in losses were comparable. Thus, the higher losses of 52 layer stacks versus the 10 layer stack would not appear to be strongly related to bulk/shear effects, and seem more likely to arise from another effect - perhaps related to the total overall thickness of the coating as interface effects seem small. As a result, a coating thermal noise analysis which assumed equal bulk and shear losses was carried out, and it seems likely that this approach may still provide a reasonable approximation for these coatings. A fuller analysis, including measurements of the loss of the silica layers, is of interest as future work to allow refinement of the thermal noise predictions.

Yam's equation for CTN (Equation 2.41) shall be used, modified slightly since the SiO_2 layer losses are not known. Every individual layer is approximated to have equivalent loss and elastic properties which correspond to the values for the stack that induced the observed frequency shift in the samples versus their uncoated state. In doing this, and still designing the stack structure with correct thickness and refractive index corresponding to the single layer properties of the two materials, one can make an approximation of the CTN. However, unlike with the $\text{TiO}_2\text{:SiO}_2$ coatings studied in previous chapters, these $\text{TiO}_2\text{:GeO}_2$ HR stacks have significant frequency dependence in the losses, which need to also be accounted for when inferring what the loss levels would be at around 100 Hz - the frequency at which we wish to evaluate the CTN. Simple linear extrapolations from the loss data points were carried out to yield estimates of the losses of each stack at 100 Hz.

Table 6.3: Coating thermal noise estimates from mechanical loss results of the 52-layer $\text{TiO}_2\text{:GeO}_2$ and SiO_2 based coating stacks, the key material properties needed to calculate these values are also given. The directly measured CTN of each stack in the folded-cavity setup at MIT is also shown in the bottom row. All values are represented as a percentage of the nominal aLIGO CTN at 100 Hz of $\sim 6.8 \times 10^{-21} \text{Hz}^{-\frac{1}{2}}$.

| CTN estimates of $\text{TiO}_2\text{:GeO}_2$ HR Stacks and key parameters | | | | |
|---|------------------|------------------|---------------------|---------------------|
| Properties | as-deposited | | 600°C, 100 hr | |
| | 52L-A | 52L-B | 52L-A | 52L-B |
| ν_C | 0.22 | 0.22 | 0.16 | 0.16 |
| Y_C [GPa] | 74^{+16}_{-16} | 74^{+16}_{-16} | 75^{+17}_{-17} | 74^{+17}_{-17} |
| $\phi_C^{\text{just-clovers}}$, @100 Hz [$\times 10^{-5}$] | - | - | $5.0^{+0.5}_{-0.5}$ | $4.9^{+0.2}_{-0.2}$ |
| $CTN_{@100\text{Hz}}^{\text{just-clovers}}$ [% of CTN_{aLIGO}] | - | - | 62^{+5}_{-3} | 62^{+3}_{-2} |
| $\phi_C^{\text{all-modes}}$, @100 Hz [$\times 10^{-5}$] | 32.1 | 33.0 | $5.3^{+1.0}_{-1.0}$ | $5.2^{+0.4}_{-0.4}$ |
| $CTN_{@100\text{Hz}}^{\text{all-modes}}$ [% of CTN_{aLIGO}] | 156^{+7}_{-3} | 164^{+8}_{-8} | 64^{+8}_{-5} | 63^{+5}_{-2} |
| Directly measured [318] | | | | |
| $CTN_{@100\text{Hz}}^{\text{MIT}}$ [% of CTN_{aLIGO}] | N/A | N/A | 72 | 66 |

However, due to the quite different loss–frequency trends witnessed for the clover and mixed modes, as an additional check an extraction of the loss down to 100 Hz was also done for just the clover modes, to evaluate the effect on the overall estimates of the CTN. Given the generally lower losses of the clover modes versus the mixed modes toward lower frequencies, with both families falling in line at higher frequencies, this second estimate would serve as a lower limit on the overall CTN.

As mentioned in Section 6.2, the Young’s modulus and Poisson ratio of coatings could be inferred through the frequency shift observed between the samples in their coated and uncoated state. Using these values for the stack coated samples along with the extrapolated losses at 100 Hz the CTN could be extracted for the as-deposited coatings and for the coatings after 600°C 100 hr heat treatment. The 52L-A, and 52L-B coatings in their heat treated states were also measured in the MIT direct CTN measurement setup described in Section 3.5.1, thus allowing for verification of these predictions. Table 6.3 shows the results of these calculations, along with the losses, Young’s moduli and Poisson ratios that were used in Equation 2.41.

The estimates of the CTN levels of these coatings from the loss results at 600°C 100 hr heat treatment are: for 52L-A $64^{+8}_{-5}\%$ of the current aLIGO coatings, and for 52L-B $63^{+5}_{-2}\%$. When taking just the clover modes to estimate the CTN, these values both drop by around only 2% - see Table 6.3, giving some confidence in the robustness of this analysis approach.

For both coatings, the heat treated CTN values represent a reduction by around 60% from their as-deposited CTN levels. Both coatings having very similar CTN to each other makes sense given that the total difference in thickness/volume between the two materials in the stacks was only around 1%. Indeed given that the stack losses dropped to $\sim 16\%$ of their as-deposited values on heat treatment for both coatings this CTN reduction seems reasonable, as through Equation 2.31 for example, we can see the CTN amplitude spectral density is proportional to the square root of the loss and so $(1 - \sqrt{0.16}) \approx 60\%$. This result for the CTN of these stacks is exceptionally low, and almost provides the desired 50% reduction versus current detectors, but is still significantly higher than the 45% of aLIGO levels initially predicted from single layer loss measurements of $\text{TiO}_2:\text{GeO}_2$ [148].

As a verification of these CTN estimates, the influence of the only parameter without well defined uncertainties, ν_C , was further investigated. The code used to estimate the elastic properties of coatings from the frequency shift they induce in the resonances of disk samples often cannot well approximate the Poisson ratio. This is due to the fact that slight changes in Poisson ratio has very little impact on the frequency shifts, and hence the energy ratios/dilution factors extracted for inferring coating mechanical loss. This weak dependence on the Poisson ratio, is the likely source of the apparent change in the estimates from 0.22 to 0.16 post heat treatment. Both values, however, are still consistent with the stack being a composite mixture of IBS $\text{TiO}_2:\text{GeO}_2$ and SiO_2 with values found of 0.25 and 0.11 respectively [148]. Nevertheless, to test the impact of changing this parameter, estimates were made for 52L-B in keeping all other parameters fixed and allowing ν_C to vary from 0–0.3, which produced a variation in the final CTN estimates obtained of no greater than 6% between the two extremes of the range.

These CTN estimates were further verified through comparison with direct measurements made by collaborators at MIT, which yielded CTN of 72%, and 66% of current aLIGO levels for the 52L-A and 52L-B coatings respectively [318]. Both measurements are in agreement with the range of values predicted from the loss results. This agreement with the direct cavity CTN measurements gives further confidence in the validity of the values obtained, despite the quite different loss trends observed between the mode families. The fact that there was minimal change in the CTN estimates when running the analysis with just the clover mode losses also suggest that any potential large increase in the bulk losses

at lower frequency (as inferred from the steep increasing trends of the mixed modes) does not contribute to the CTN significantly. Another possible explanation for this is that some other factor is leading to the loss increase of the mixed modes, rather than an increase in bulk loss, such as unobserved defects or contamination, though this lies outside the scope of this analysis.

Ultimately, despite the significantly different loss trends which arose between single layers of $\text{TiO}_2\text{:GeO}_2$ and the full reflectivity HR stacks, the CTN of these samples has been characterised, and witnessed to be in agreement with direct thermal noise measurements. They both appear to provide a significant reduction in CTN versus the current aLIGO coatings, with 52L-B also having demonstrated a promising path forward for reducing the impact of the optical absorption of $\text{TiO}_2\text{:GeO}_2$ as well, with no detriment to the CTN performance.

6.8 Titania-germania loss study conclusions

A range of studies to investigate factors affecting the mechanical loss of $\text{TiO}_2\text{:GeO}_2$ -based coatings, and SiO_2 layers for use in these coatings, were carried out.

The mechanical losses of initial coatings produced by the GWD coatings manufacturer LMA were measured. These investigations identified an optimal heat treatment regime for single layers of $\text{TiO}_2\text{:GeO}_2$ of heating to 600°C for 108 hours. Following this heat treatment, the coating losses were around $\times 4$ lower than the currently-used GWD coating high refractive index material, $\text{TiO}_2\text{:Ta}_2\text{O}_5$. This level in loss reduction held both at the lowest measured frequency of 1.1 kHz, and after extrapolation down to 100 Hz where the CTN dominates the noise in current room temperature detectors. The loss of the $\text{TiO}_2\text{:GeO}_2$ (ϕ_C) was found to vary with frequency as $\phi_C = 3.2 \times 10^{-5} f^{-0.13}$. This reduction in loss for the high- n material was also accompanied by great improvements in the losses of SiO_2 coatings made via the same process. The SiO_2 improved in loss by around 76% versus the 500°C , 10 hour results, which are representative of current GWD SiO_2 coatings, seemingly reducing to $(8.0 \pm 0.7) \times 10^{-6}$ after 600°C , heat treatment for 108 hours, which in itself is an important and highly impactful result.

Subsequent investigations were launched into LMA, and CSU single layer coatings produced with elevated deposition temperatures at 150°C and 200°C respectively, and using an order of magnitude lower chamber base pressures of $\sim \times 10^{-7}$ mbar. These changes were implemented in order to mitigate defects seen to arise in the TiO₂:GeO₂ single layer coatings and stacks with heat treatment. Ultimately these studies successfully verified that those alterations led to single layers and stacks free from blisters or cracks after the final heat treatment, and had no detrimental impact on the mechanical losses of TiO₂:GeO₂. The losses of the defect free TiO₂:GeO₂ CSU single layers was seen to follow a near identical trend with frequency to the LMA single layers, though with slightly lower losses towards lower frequencies, and as such could be estimated to yield loss levels of $(18.3 \pm 1.7)\%$ of the currently employed TiO₂:Ta₂O₅ at 100 Hz.

Two full-reflectivity HR coatings were made of these two materials and measured via GeNS. One which had a uniform layer structure (52L-A), and one where the thicknesses of the layers was altered throughout the coating in order to have less of the higher absorption TiO₂:GeO₂ near the top of the stack (52L-B), thus less light would be interacting with it. Both these stacks, while still reducing overall in their losses with heat treatment, did not reduce as much as expected from the single layer results, towards lower frequencies. Both exhibited broadly flat loss across frequency for their clover modes, and conversely exhibited increasing loss towards lower frequency for their mixed modes after heat treatment. Because of this, subsequent studies were launched into characterising the effects of increasing the number of layer interfaces, total thickness, and fractional material composition. The results indicated that interface effects were relatively small, and that total thickness and material split seem to be the largest contributors to the observed differences in loss.

Ultimately CTN values for these full stacks were able to be extracted yielding values of $64^{+8}_{-5}\%$ and $63^{+5}_{-2}\%$ of current aLIGO levels for 52L-A and 52L-B respectively, which was also verified by collaborators at MIT. These results show both the great promise of this TiO₂:GeO₂ as a next generation GWD coating material, with it's exceptional mechanical losses. However, the result is also around $\times 1.5$ higher CTN than HR stacks of this coating were predicted to achieve, based on previous single layer estimates [148]. The smaller

improvement in the loss of these stacks versus single layers after heat treatment would appear to be the cause for this. Whilst the mechanism for this reduced loss improvement is not fully understood, attempts were made to characterise it, and hone in on the potential causes, laying the groundwork for future research.

Regardless, these coatings are still highly promising candidates for the next upgrades of GWDs, with exceptionally low CTN. Although their absorption is still at the few-ppm level [317] and requires further research investment into the deposition parameters to resolve sub-ppm absorption. That being said, this also means it would be a very interesting material to pair with the previously investigated $\text{TiO}_2\text{:SiO}_2$, or indeed current $\text{TiO}_2\text{:Ta}_2\text{O}_5$, both of which have sub-ppm optical absorption and low scatter. A stack comprising primarily of alternating $\text{TiO}_2\text{:GeO}_2$ and SiO_2 , but with the top few bi-layers of $\text{TiO}_2\text{:GeO}_2$ swapped with either of these other two materials, could indeed be one way forward for the very next detector improvements.

Chapter 7

aLIGO test mass characterisation studies

7.1 Introduction

This chapter describes investigations conducted as part of a visiting researcher placement at the California Institute of Technology (Caltech) from August 2022 - December 2022, alongside Dr GariLynn Billingsley, Dr Camille Makarem, and Dr Liyuan Zhang in the LIGO Caltech Core Optics Components division. During this placement the author characterised numerous parameters of LIGO test mass optics.

Absorption maps were produced of the two LIGO-Livingston (LLO) end test mass (ETM) optics, previously installed in the first through third observing runs of Advanced LIGO, in order to locate and characterise their so-called ‘point absorbers’, some of which will have affected detector duty-cycle while both ETMs were in use in the detector. An optical microscope survey of these point absorbers was also conducted, and a new cleaning procedure was developed for potential use in future optics in order to mitigate their influence, and in some cases remove them entirely.

Alongside this work, the author also assisted in the high precision geometric characterisation and analysis of three uncoated input test mass (ITM) optics, later to be coated and installed in future LIGO upgrades. The thicknesses, diameters, and barrel marker locations/orientations required for suspending the optics were determined. Their front-face surface profiles were also characterised to high precision using a custom ZygoTM VerifireTM interferometer [320], from which their radii of curvatures were also inferred. The latter of these geometric measurements are detailed in this thesis.

7.2 Characterisation of LLO ETM point absorbers

Optical absorption can be broadly split into two categories; uniform, and non-uniform [321]. Uniform absorption is characterized by an approximately spatially invariant absorption value across the sample. The aLIGO ETMs, which employ $\text{TiO}_2:\text{Ta}_2\text{O}_5$ and SiO_2 coatings to form a highly reflective (HR) mirror, possess a by-and-large uniform absorption of (0.27 ± 0.07) ppm [124] across their front surface. The influence of this absorption, and the anomalies in the ETMs which result via induced heating of the optic are well understood and certain effects can be well mitigated. Two such anomalies (as discussed in Section 2.2) are thermal lensing, and curvature changes. GWDs have thermal compensation systems [322, 323], as will be discussed later in this chapter, which provide low spatial frequency correction of the resulting net test mass deformation from these effects [324]. Non-uniform absorption, on the other hand, is any form of absorption with high spatial frequency arising from features much smaller than the laser beam, leading to much more localised effects that are randomly distributed [321, 325]. A salient example of this is a “point absorber”.

Point absorbers are sub-millimetre, highly absorbing points which are seemingly scattered stochastically across the surfaces of the primary optics of the Advanced LIGO and Advanced Virgo detectors [321]. The interaction of these point absorbers with high power laser light in the arm cavities, results in anomalies such as nanometre scale substrate lenses, and thermo-elastic deformations which significantly reduce the sensitivity of the detectors. The resulting reduction in performance arises directly through a reduction in the power-recycling gain, and indirectly through the feedback control system [321].

The existence of point absorbers in the aLIGO optics has been known since the first observing run [326]. However, their origin, and the exact mechanisms that give rise to them in the GWD optics are not well understood, and as such this is a highly active area of research [321, 326]. Some may arise from contamination during measurement, transportation, and installation at the detector sites, whereas others may come from features embedded inside the coating, deposited during the coating process, or indeed arising during the post-deposition heat treatment process [326]. Regarding those potentially forming during coating production, spectroscopic inspections of ‘witness’ samples coated in the

same runs as the aLIGO and adV test masses, have reported high concentrations of aluminium in regions of higher absorption in the coatings [321] – perhaps giving insight into the nature of some point absorbers. Regarding post-coating deposition contamination, some investigations have reported additional point absorbers seemingly appearing during the inspections of the optics [321]. Investigations conducted onsite in LLO and LHO, as well as those by the author at Caltech using an optical microscope, reveal a vast menagerie of point absorbers with myriad features on the test masses, many showing clear signs of having formed through different mechanisms.

7.2.1 The impact of point absorbers on gravitational wave detection

There are multiple laser power dependent loss mechanisms arising from the presence of point absorbers, that reduce the performance of a GWD. For example, local distortions of the mirror surface due to point absorbers can give rise to wide-angle scattering of the laser beam in the interferometer arms. This leads to a loss of power in the arm cavity as they deflect a portion of the light to the walls of the beam tube, stopping it cycling through the arm cavities [326]. The primary source of cavity losses actually comes from resonant enhancement of losses of intermediate higher order optical modes [321, 325]. These losses include scattering and resonant loss of light power from the fundamental optical mode into higher order modes (HOM). Point absorbers cause this due to the distortions of the mirror surface they induce, leading to subsequent distortions and phase shifts in the reflected wavefronts. The cavity will resonantly enhance/suppress those modes as a function of the round trip phase they accumulate in the cavity. HOM experience a greater degree of clipping at the edges of the optic and therefore high loss [321].

In a dual-recycled Fabry-Perot Michelson interferometer, such as aLIGO and AdV, any decay in power-recycling gain can be observed as a function of the input laser power P_{IN} and the stored laser power in the cavity P_{A} . The latter is itself a product of the former, the beam splitter transmission (nominally 50%), the power recycling gain G_{P} , and the optical cavity gain G_{A} defined as [321]

$$P_{\text{A}} = \frac{1}{2} G_{\text{P}} G_{\text{A}} P_{\text{IN}}. \quad (7.1)$$

The optical power absorbed by a point absorber, is dependent on its physical size, and the power stored in the arm cavities. The latter of which, from Equation 7.1, can be seen to be dependent on the input laser power. The amount of power absorbed will then influence the amount of localised deformation on the optic. As discussed, these deformations then lead to power loss in the arm cavity due to heat transfer, and to deflection of portions of the beam, leading to some light not traversing the whole cavity. It is essential to determine the resulting surface deformation induced from a point absorber arising both inside and outside the high absorbing region. With these, the full deformation induced by the point absorber on its host optic can be inferred.

Brooks et al. produced an equation framework for approximating the point absorber induced surface deformations [321] expanding on previous work by Winkler et al. [327]. They also went on to estimate the power loss into HOM from the fundamental mode as a result of these surface deformations, ultimately determining that the losses from point absorbers would increase with increasing cavity power [321]. The full mathematical treatment is not essential for the project undertaken by the author at Caltech. However, interesting future work may lie in modelling the deformations one might expect from the point absorbers that were discovered, based on their positions, sizes, and absorption values recorded, and seeing how the overall effect compares with the losses observed through O1–O3.

To illustrate that the loss of cavity power due to point absorbers will impact detector performance, Equation 7.1 will be rearranged whilst considering the following approximation. Since the ITM power transmission is very low, and the ETM power transmission is much lower ($T_{\text{ITM}} = 1.4\%$ and $T_{\text{ETM}} = 5 \text{ ppm}$ respectively) the following approximation holds for the optical gain of the Fabry-Perot arm cavity [321]

$$G_A \approx \frac{4}{T_{\text{ITM}}}. \quad (7.2)$$

Substituting Equation 7.2 into Equation 7.1 and rearranging, one can glean that any loss in the light power stored in the cavity will directly reduce the overall power recycling gain

$$G_P \approx \frac{1}{2} T_{\text{ITM}} \frac{P_A}{P_{\text{IN}}}. \quad (7.3)$$

One can use the power-recycling gain as a proxy for the overall interferometer performance at high frequencies. Brooks et al. [321] predict that the power-dependant reduction in GWD performance from point absorbers could significantly degrade maximum stored cavity power by up to 50%, and hence limit detector sensitivity. Though, with system wide corrections implemented, such as re-alignment of the input laser, they expect to mitigate this to a minimum 33% deficit in arm power build-up at nominal operating power of 125 W. Assuming no other adverse effects to interferometer operation, they predict a corresponding increase in the shot noise floor of 15%–20% above the nominal aLIGO design specification noise floor due to this effect [321]. Given that future GWDs demand a stored cavity power of approximately an order of magnitude or more higher than in Observing run 3 (O3) [98, 328], with similar laser intensity to full-power aLIGO [125, 321], reducing both the quantity and absorption of these high point absorbers, hence mitigating their influence on the optical fields of GWD, is of crucial and increasing importance [321, 326].

7.2.2 PCI absorption maps of LLO ETMs

The photothermal-commonpath interferometry (PCI) technique as described in Section 4.2.1 was used to map the absorption of the LIGO Livingston ETM optics. These were the first such measurements on “advanced” detector test masses, at the end of their lifespan, after having been removed from a gravitational wave detector. Images of a test mass installed in the setup can be viewed in Figure 7.1. It was hoped that these measurements, along with the subsequent imaging microscopy investigation conducted, would help identify and characterise the point absorbers that are present and arise throughout the lifespan of an optic. With the point absorbers located, alternative cleaning methods could then be trialled on those point absorbers that arose from contamination, to see if their impact could be reduced, or indeed if they could be completely removed.




Figure 7.1: Images of the Caltech PCI system with LLO ETM-Y installed. The left image is of ETM-Y seated in its translation stage, the right image highlights the paths of the pump and probe laser beams of the PCI, for a more detailed PCI schematic, see Figure 4.1.

For each ETM, a surface absorption map, of diameter 100 mm about the optic's centre was constructed by stitching together point absorption measurements at 0.2 mm intervals. The Caltech PCI setup used for these measurements was described previously in Chapter 4, Section 4.2.2. It should be noted that, in-situ, the main aLIGO laser encompasses a diameter of around 124 mm on the ETM face [328]. Therefore, there could be additional, unstudied point absorbers that may have affected detector performance, residing in this small unmeasured region beyond the area of the maps. However, due to the large time the map scans take to complete - around 17 days for 100 mm diameter - larger scan areas were impractical.

Histograms of the measured map points of each test mass are shown in Figure 7.2. Absorption values for around 200,000 points on each of the ETMs – dubbed ETM-X and ETM-Y – were acquired, but care should be taken when interpreting absolute absorption values from these initial measurements, as the coatings are thin and alignment can drift over the 2-3 week measurement. Also the pump laser was operated at relatively low intensity, during this scan, to protect the optic from potential laser damage, but consequently the signal-to-noise on the measurements is lower. Irrespective of the slight drifts, and larger relative noise, a measurement of a point absorber should still yield significantly higher than the average absorption of the vast majority of the surface and they can therefore be identified as candidates to investigate further once their spatial location was found. One can see from the histograms that the vast majority of points measured still appear




Figure 7.2: Histograms of the measured absorption, obtained via PCI, of each point of the 100 mm diameter maps about the centres of ETM-Y and ETM-X.

to produce absorption values between 0.2–3.0 ppm for both ETM-X and ETM-Y, and the threshold for which points to investigate further on this scale was those initially exhibiting greater than 50 ppm absorption. In this study, this resulted in around 70-90 initial point absorber candidates for each of the two ETMs.

Even though these measurements took place in a class 100 cleanroom, over the 17 day scan time, some dust would still fall onto the sample during measurement, and in fact most of these candidate points would originate from this dust or from particulate and surface contamination accumulated during the transportation, and handling of the optics, or perhaps leftover residue from previous rounds of cleaning. After the initial scan was conducted, each high-absorption candidate was then verified by returning to the relevant sample coordinates, adjusting surface position until a maximum absorption signal was found, and then attempting to clean this location by wiping with a polyester knit q-tip soaked with isopropanol (IPA).

If the absorption dropped to the normal levels of the surface after cleaning, then the candidate was considered spurious, originating from easily removable dust or contamination. Otherwise this cleaning would be repeated for that point, until no change/no further change in the absorption was observed after at least two consecutive attempts. If a higher

than the normal coating absorption signal persisted, the candidate was considered to be a genuine point absorber and the absorption estimates from the initial scan were updated with new estimates post-cleaning. From this, around twenty point absorbers for each ETM were identified. The pump beam used was $60\mu\text{m}$ diameter, so for any point absorbers physically larger in diameter than this, the measured absorption would represent only a lower-bound limit.

A final additional search for point absorbers was conducted, after all these candidates were verified. The author illuminated each ETM with a bright torch held close to the sample front face. In doing so any small points that shone brightly and which did not correspond to the coordinates of previously identified point absorbers, were then verified in the same way as above with the PCI and repeated cleaning with a polyester knit q-tip and IPA. This resulted in one additional real point absorber being confirmed, dubbed ETM-X_y, which resided approximately 1 cm outside the initially mapped area of ETM-X and was the second largest absorption point absorber found on ETM-X. The naming nomenclature “ETM-X a, b, c,...” follows the order in which each was confirmed as a ‘real’ point absorber.

Maps showing the relative locations of all the identified point absorbers are shown in Figures 7.3 and 7.4 for ETM-Y and ETM-X respectively. The relative sizes of the markers on the right-hand side plots of each figure represent the magnitude of that point absorber’s absorption with respect to the others. One can see there is no clear pattern nor trends in where they have manifested, although for ETM-X the highest point absorbers found do appear to be clustered together in the upper-central region of the map. Table 7.1 lists the absorption values in ppm measured from each point absorber.

After these measurements an optical microscope was then installed and a visual survey of the point absorbers conducted, which allowed for their physical sizes and features to be observed.

For ETM-Y 23 point absorbers were found, and for ETM-X 20 were found - all of which resided within the main interferometer beam area. Of these, four on each test mass were observed to have greater than 1% absorption at 1064 nm, and 5-6 points on each optic were found to be physically larger than the pump beam diameter - these are indicated with tick markers in Table 7.1. The three point absorbers exhibiting both the largest apparent absorption values, and largest physical sizes all appeared quite similar, and visually distinct from all other point absorbers found. These were ETM-X h , ETM-X y and

Table 7.1: All LLO ETM-Y and ETM-X point absorbers found via PCI mapping. The corresponding ID, absorption recorded at 1064 nm in ppm are shown. Also shown is whether the point absorber was physically larger than the PCI pump beam size of 60 μ m.

| № | ETM-Y ID | \varnothing 60 μ m | \varnothing > 60 μ m | ETM-X ID | \varnothing 60 μ m | \varnothing > 60 μ m |
|----|----------|--------------------------|----------------------------|----------|--------------------------|----------------------------|
| | | α_{1064} [ppm] | | | α_{1064} [ppm] | |
| 1 | ETM-Y o | 41700 | ✓ | ETM-X h | 64200 | ✓ |
| 2 | ETM-Y r | 22800 | | ETM-X y | 44300 | ✓ |
| 3 | ETM-Y n | 16800 | | ETM-X c | 20600 | |
| 4 | ETM-Y b | 10300 | ✓ | ETM-X l | 18000 | ✓ |
| 5 | ETM-Y w | 9000 | | ETM-X u | 5600 | |
| 6 | ETM-Y a | 8800 | | ETM-X e | 5500 | |
| 7 | ETM-Y f | 6700 | ✓ | ETM-X f | 4200 | |
| 8 | ETM-Y u | 4900 | | ETM-X r | 3500 | |
| 9 | ETM-Y t | 4700 | ✓ | ETM-X p | 3000 | |
| 10 | ETM-Y v | 4700 | | ETM-X s | 3000 | |
| 11 | ETM-Y h | 2800 | | ETM-X w | 2500 | |
| 12 | ETM-Y i | 2200 | | ETM-X d | 1900 | ✓ |
| 13 | ETM-Y k | 1900 | | ETM-X g | 1300 | |
| 14 | ETM-Y c | 1700 | | ETM-X o | 980 | |
| 15 | ETM-Y d | 1300 | | ETM-X x | 720 | |
| 16 | ETM-Y s | 900 | ✓ | ETM-X n | 650 | ✓ |
| 17 | ETM-Y m | 770 | | ETM-X k | 550 | |
| 18 | ETM-Y g | 450 | | ETM-X q | 470 | ✓ |
| 19 | ETM-Y p | 350 | | ETM-X t | 310 | |
| 20 | ETM-Y j | 280 | | ETM-X a | 90 | |
| 21 | ETM-Y e | 150 | | | | |
| 22 | ETM-Y l | 150 | | | | |
| 23 | ETM-Y q | 98 | | | | |




Figure 7.3: Spatial coordinates of the LLO ETM-Y point absorbers (left) plotted with respect to the centre of the optic. A corresponding bubble plot (right) shows the relative 1064 nm absorption of each point absorber. The large burn mark point absorber found is imaged on the inlay of the left-hand graph, and highlighted in red on the right-hand graph.




Figure 7.4: Spatial coordinates of the LLO ETM-X point absorbers (left) plotted with respect to the centre of the optic. A corresponding bubble plot (right) shows the relative 1064 nm absorption of each point absorber. The microscope images of the two large burn mark point absorbers found are imaged on the inlay of the left-hand graph, and highlighted in red on the right-hand graph. Note that point ETM-X-y was not found during the initial scan, as it was ~ 10 mm outside the initially mapped area.

ETM-Y o. Each appeared with the microscope to be ovular craters of ≈ 0.1 mm diameter. Two such craters were found on ETM-X: ETM-X h, and ETM-X y, with measured lower bound absorption of 6.4% and 4.4% respectively at 1064 nm. ETM-Y a had a lower bound estimate on its absorption of 4.2%.




Figure 7.5: Bright-field (top row) and dark-field (other rows) microscope images of the three large crater point absorbers of highest recorded absorption found on the LLO ETMs. On the bottom row for ETM-Xy the image is focused at a different depth from the row above, and for ETM-Yo, a longer exposure time was taken than for the image above it.

For the purposes of the following discussions on the microscopy studies we shall focus on a subset of the 43 point absorbers catalogued, primarily those which exhibited the highest absorption – the craters and the following highest absorption points with around 1% and greater measured absorption.

Each of the craters was observed to have significant depth, due to the necessity to significantly alter microscope focal height in order to clearly resolve the centre and edge of each to high precision. The centre of each crater was the point deepest into the ETM surface. None of the other point absorbers found exhibited such large downward depth profile, and none of the others (of sufficiently large enough physical size to image clearly) were as perfectly ovular/circular. Images of all three of these large point absorbers are shown in Figure 7.5.

The distinct lack of optically bright features near the centre of each of these crater point absorbers, as seen in their dark-field images, and increased scatter at the perimeter, and indeed the particulate that surround them (seen here with ETM-Y o), suggest some destructive process has perhaps occurred. One plausible explanation is that they are the remnant of an explosion or ‘burn’ which arose from interaction with whatever was previously in the centre of these craters – something with catastrophically high-absorption – and the high power LLO GWD primary laser beam/ as power was built up in the arm cavity. If this is indeed the cause there is not much that can be done to mitigate them once they arise. To better understand why they might form, scrapings of ETM-Y o and its surroundings, were sent to collaborators at NASA’s Jet Propulsion Laboratory for compositional analysis, but unfortunately their studies were inconclusive.

7.2.3 Cleaning off some of the highest absorption point absorbers

There were a few other point absorbers with similar, but slightly lower, absorption to the crater points. Most of these appeared as firmly affixed contaminants or as features slightly embedded in the coating. A new cleaning technique was trialled in attempt to remove many of these or at least significantly reduce their impact.

The current cleaning procedure involves wiping areas with a polyester knit q-tip soaked in IPA, or indeed pouring on First Contact Polymer solution (of similar cleaning capability to IPA) over wide areas of the optic, and peeling off the film that forms. It was decided to trial two more aggressive cleaning processes.

- Firstly, swapping the polyester q-tip for a firmer cotton-bud, still soaked in IPA.
- Secondly, swapping the IPA on the cotton-bud for a stronger solvent; acetone.

The soft polyester knit q-tips currently used, are great for cleaning without leaving fibres and residue on the test masses, like one would see from standard lens tissue or a cotton-bud. These fibres themselves could appear as their own point absorbers, burn, or otherwise contaminate the vacuum chamber of the detector when installed. However, the cotton-buds have both a rougher and firmer surface, so it was thought that they could be useful in attempting to remove firmer-affixed surface contamination point absorbers. Unlike the polyester knit q-tips, the cotton-buds can also be soaked in stronger solvents like acetone




Figure 7.6: Bright-field microscope images of, from left to right, the second-to-sixth highest absorbing point absorbers of ETM-Y before (top) and after (bottom) the new cleaning method trials with soaked cotton q-tips. Point-r and point-n were only cleaned using isopropanol, point-b only with acetone, and point-w and point-a, were cleaned with both isopropanol then acetone (with their images shown being after the latter cleaning).

without dissolving, and as such could facilitate deeper cleaning. The point absorbers on which these new cleaning methods were trialed in order of decreasing absorption were ETM-Y r, n, b, w, and ETM-Y a, which were the second-to-sixth highest absorption point absorbers found on ETM-Y. After the cleaning process their absorption levels were remeasured with the PCI. Microscope images of each point absorber, before and after these cleaning trials, can be seen in Figure 7.6.

Of these, all bar point ETM-Y a (henceforth referred to as point-a) were seen to improve by simply switching from a polyester knit q-tip to a cotton-bud. Point-a was the lowest absorption of the investigated points, and appeared to be a small blob of pink particulate ~ 10 micron in size resting atop the optic - perhaps being a melted piece of affixed surface plastic or of First Contact polymer from previous cleaning before being installed in the detector. Upon switching to acetone, point-a was almost completely removed, with only a thin remnant outline of it remaining (see Figure 7.6). After just IPA cotton-bud cleaning, point-r seemed to completely vanish and likewise point-n left only a small remnant. These two point absorbers were not cleaned with acetone in order to later see the impact on the absorption values of just switching to the cotton-bud.

Ideally after every round of cleaning a corresponding absorption measurement would have been acquired, but this was not practical as the microscope had to be mounted on the PCI optical stage, and constantly swapping it in/out and resetting up/aligning every measurement so as to not block the PCI optics would be far too time consuming and potentially introduce alignment errors and/or loss of spatial information. Thus, for point-n which did not appear to be fully removed, upon subsequent cleaning with acetone one might find further reduction in the absorption. The fifth largest absorption point absorber, point-w, was the only one not witnessed to change in appearance greatly, with either new cleaning method.

Table 7.2 shows the absorption results obtained with the PCI for the point absorbers where the new cleaning methods were trialled. One can see the large reduction in absorption for all of the point absorbers. The second and third highest point absorbers –point-r and point-n – which previously exhibited $\sim 2\%$ (i.e. 20000 ppm) absorption, were either completely removed (point-r) or at worst had improved by 2 orders of magnitude (point-n), just from switching to a cotton-bud while keeping the same cleaning solvent. Point-r and point-a visually appeared to be mostly removed (though only after acetone cleaning for point-a) and similarly no longer showed measurable absorption greater than the normal levels of the coating. Point-b and point-w which experienced the smallest visual and absorption changes, still improved in their level of absorption by 75% to 80% respectively. Both were expected to improve the least given that from the visual microscope inspections they had changed the least of the five, but their absorption reductions are still significant. Of note, point-b appears to have similar type of residue to point-a, but with a much brighter grain at the centre, perhaps indicating significant surface damage. Point-c, of initially much

Table 7.2: The absorption measured via PCI of the LLO ETM-Y point absorbers after cleaning with a cotton-bud and IPA/acetone.

| ETM-Y ID | $\alpha_{1064}(\varnothing 60 \mu\text{m})$ [ppm] before | $\alpha_{1064}(\varnothing 60 \mu\text{m})$ [ppm] after (new cleaning) | level of improvement |
|----------|---|---|-------------------------|
| ETM-Y r | 22800 | comparable to background | >99.95% |
| ETM-Y n | 16800 | 150 | 99% |
| ETM-Y b | 10300 | 2600 | 75% |
| ETM-Y w | 9000 | 1800 | 80% |
| ETM-Y a | 8800 | comparable to background | >99.95% |
| ETM-Y c | 1700 | comparable to background | >99.95% |

lower absorption than the others presented here, and not specifically targeted in these new cleaning trials, resided quite close to point-b (see Figure 7.3) and while cleaning point-b it was realised that the cotton-bud was also dragged over point-c as such it was also remeasured in the PCI, and found to have also vanished from the cotton bud cleaning.

7.2.4 LLO ETM point absorber conclusions

Through these investigations the ETMs previously installed in LLO were mapped and the point absorbers hindering the detectors performance located, imaged, and had their absorption levels quantified.

The results of these new cleaning trials on these point absorbers (though in ethos quite simple procedural changes) are of great significance for future GWDs. They show the great promise for reducing the influence of many of the highest-absorption point absorbers present on GWD test masses which hinder detector performance, or indeed completely removing them. Through this study, it was found that many of the highest absorbing point absorbers seem to originate from surface contamination and can be removed and mitigated by more aggressive cleaning methods. Being able to undertake these tests on optics at the end of their operational lifespan provided a unique opportunity to perform such characterisation on the actual optics of interest, with histories directly comparable to what can be expected for the future test masses.

However, surveying the optic with the microscope around the regions in which they were cleaned with the cotton-bud revealed, as expected, that many cotton fibres had shed onto the optics. These were able to be removed with subsequent First Contact cleaning without needing to disperse them to the lab environment with a nitrogen gun (something one would not want to do in the chamber of a GWD). However, First Contact cleaning cannot currently be done while the test mass is installed in a GWD as it needs to be poured onto a horizontal surface and then left to dry, and the ITMs and ETMs are suspended with their faces vertical. Nevertheless, these results seemed promising enough that methods for First Contact cleaning while the test masses are in-situ in the aLIGO detectors are currently being developed so that point absorber cleaning/removal techniques such as these could be implemented in the future, and the fibre left behind from the process removed [329].

7.3 ITM figure and radius of curvature studies

The surface profiles and radius of curvature of the three ITM optics were characterised with a custom Zygo™ Verifire™ interferometer [330, 320] (which is a high power Fizeau interferometer) optical measurement system at Caltech, henceforth referred to simply as the Zygo. This system is designed to measure, with high precision, the surface profile of myriad aLIGO optical components. A Zygo measurement involves comparing the surface of an optic of interest, with the known surface of a reference optic to which the system is calibrated. The reference optic's surface profile and radius of curvature are characterised to high precision prior to the measurement, with it being chosen as a reference optic due to its high surface quality free from large aberrations. In the Caltech system the reference optic, also known in common metrology terms as a transmission flat, or transmission sphere, is mounted on French cleats at the collimator output of the Zygo - allowing for repeatable reference subtraction.

The optic under investigation, referred to as the test optic, is placed in the path of the Zygo beam (a 1064 nm, 5 mW, 360 mm diameter beam), after it has passed through the transmission sphere, such that both the surface of interest and the reference surface are illuminated. Then, when light waves reflect back into the Zygo from the test optic's surface, they interfere with the reference wavefront reflected back into the Zygo from the transmission sphere. This interference pattern contains information about the shape of the test optic, allowing for precise measurement of its surface deviations from the ideal shape. A diagram showing a simplified setup can be viewed in Figure 7.7.




Figure 7.7: Diagram showing the Zygo surface profile measurement configuration, with the test optic being measured in the converging portion of the measurement beam. Adapted from [331].

The transmission sphere is essentially a lens through which the Zygo laser beam passes through before striking an optic of interest. It is used to transform the Zygo's collimated output laser beam into a spherical wavefront and also acts as a beamsplitter to divide the laser light into separate 'measurement' wavefronts (which are forward transmitted) and 'reference' wavefronts (reflected back into the Zygo) [331]. The last surface of the transmission sphere (closest to the optic of interest) acts as the reference 'master' surface. The quality of the reflected reference wavefront depends upon the figure quality of the master surface as the master surface reflects a portion of the laser beam back into the Zygo, forming the reference wavefront. The remaining laser light, now being a high quality spherical wavefront, acts as the measurement beam which converges to a focus in front of the transmission sphere and then diverges beyond this on the other side of focus. One can think of the diverging and converging parts of the beam as a library of precise spherical wavefronts with an infinite range of diverging radii and a limited range of converging radii [331] - all aLIGO test masses, being convex surfaces need to be measured in the converging portion of the measurement beam. By using a transmission sphere with a known shape, a Zygo can provide accurate and traceable measurements of optical surfaces, ensuring the quality and precision of optical components in aLIGO.

For surface profile measurements, it is beneficial for the transmission sphere's radius of curvature to be similar to that of the test optic. Matching their radius of curvatures helps optimize the interferometric measurements for accurate and reliable results. One reason for this is that it leads to more even spacing of the interference fringes generated during the measurement. Similarity in curvature, indeed, reduces potential distortions in the fringe patterns, making it easier to interpret and extract precise information about the test optic's surface. Aligning and calibrating the Zygo also becomes more straightforward when the transmission sphere's and test optic's curvatures closely align, which assists in the setup procedures of the system, enhancing the overall measurement reliability.

The Zygo captures and analyses the observed interference pattern in order to determine variations in the optical path length between the test optic and the reference surface. From this a detailed 3-D surface profile of the test optic face can be generated. The surface profile results can then be used to scrutinise the polishing procedure, observe defects on the surface, and otherwise evaluate the quality of a given optic, and if desired, also calculate its radius of curvature.

7.3.1 Zygo interferometer measurement setup procedure

For measuring the surface profiles of the three ITMs, a transmission sphere of diameter 362 mm and 2100 m radius of curvature was installed. When the transmission sphere is first installed in the system, it must be aligned to be as perpendicular as possible to the incoming measurement laser beam coming from the Zygo. In order to facilitate this, a retro reflector was placed after the transmission sphere in the beam path, and the interference pattern produced by this and the transmission sphere is then observed in the Zygo measurement software. The transmission sphere orientation was controllable via fine-adjustment knobs on the side of the Zygo which were used to align the transmission sphere horizontal and vertical tilt until the interference pattern was minimised, and the fringes were nulled - with this the transmission sphere could be considered aligned parallel to the Zygo measurement beam. This process is illustrated in Figure 7.8 (A).

Next, one wants to ensure any optic under investigation is installed parallel to the transmission sphere (i.e. also being perpendicular to the measurement beam). To ensure this, an auxiliary, visible-by-eye, green laser was reflected off the transmission sphere surface, and its reflected position recorded at a distance on the far wall of the lab ~ 4 m from the transmission sphere. This provides an easily visible reference point of the orientation of the transmission sphere, and is illustrated in Figure 7.8 (B). With the auxiliary laser locked in place, the test optic is then aligned to this reference point via the following procedure:

- The mounted test optic is brought to position approximately 0.1-0.3 m from the transmission sphere surface, care being taken to ensure that it, nor its mount clips the auxiliary laser path.

- The retro reflector, used previously, is then positioned in the auxiliary laser path, directing it onto the surface of the test optic, which it is then reflected back off onto the retro reflector, and will appear on the wall at the far side of the lab where the reference point from the transmission sphere has been marked.
- With this, the test optic's orientation can be adjusted through a combination of coarsely moving the mount, as well as fine tilt controls, until the return spot from the test optic overlaps with the reference point set previously by the transmission sphere - illustrated in Figure 7.8 (C).
- The test optic mount is then locked in place, the retro reflector removed, and the auxiliary laser switched off, with both the test optic and transmission sphere shrouded in a large fabric netting before the operator then leaves the laboratory.
- Next, the produced interference pattern between the transmission sphere and the test optic is observed with the Zygo software from a remote location, with final fine adjustments of the test optic alignment being made, via remote control of test optic mount actuators until the interference fringes are nulled. With this, the optic can be considered parallel to the transmission sphere, the setup is in the configuration previously shown in Figure 7.7, and measurements of its surface profile can be acquired.

Note that this final adjustment is carried out remotely in order to remove influence of vibration coupling, and temperature gradients from the presence of an operator in the laboratory. In general, before the final alignment the system and components must be allowed to thermalise, as temperature changes can lead to wavefront, and hence alignment distortions - this generally means a few hours wait time between leaving the laboratory and this final remote alignment. To mitigate vibration coupling leading to misalignment, the system and components are all placed on a vibration isolation table. Air current fluctuations are also mitigated by shrouding the test optic and transmission sphere in the netting before leaving. The final remote alignment step may have to be repeated multiple times until the fringes are observed to be stable and nulled. Generally the author waited a further 20 minutes after the remote alignment, observing if the fringe pattern on the

optic had changed significantly indicating alignment drift, and when it appeared stable (i.e. the fringes were still nulled), a measurement was taken. Examples of the interference pattern acquired for an optic at different stages of alignment before measurements were taken are shown in Figure 7.9.




Figure 7.8: Diagrams showing the alignment stages to facilitate the surface profile measurement. (A) shows the aligning of the transmission sphere to the Zygo measurement beam, (B) then shows the procedure for defining a reference point for the aligned transmission sphere, and (C) the orientation of a test optic's surface with that reference point, such that it is then parallel to the transmission sphere.




Figure 7.9: Images of the interference pattern produced by a test optic's surface reflections interacting with the transmission sphere reference wavefront, as recorded with the Zygo. The left image shows initial alignment drift (greater number of fringes) after the laboratory was left to thermalise for a few hours. The centre image fringes results from waiting 20 minutes post initial remote null fringe alignment when the test system was not yet stable, and the leftmost image shows the same, but once the system was stable. The various small dark blotches, mostly visible near the centre, result from factors Zygo/transmission sphere imperfections, but mostly due to unfortunate alignment laser damage on the Zygo camera.

7.3.2 Interpreting the test optic surface profile from Zygo wavefront data

With the optic of interest aligned with the transmission sphere, measurements of its surface profile were acquired. In order to interpret the results, the interfering wavefronts re-entering the Zygo are modelled by the software as the sum of a series of polynomials. We choose the Zernike polynomials [332] for this task, which is common practice in optical metrology measurements [333], largely due to their similarities with commonly seen optical aberrations [334, 335]. The Zernike polynomials have a highly useful property, being that they form an infinite series of orthogonal expressions over a continuous unit circle, which enables simple deconstruction of a wavefront into a series of different Zernike components [336]. The Zernike polynomials $z_n^m(\rho, \theta)$, used to interpret the returning wavefronts, are expressed in terms of the normalised radius ρ (defined such that $0 \leq \rho \leq 1$) and azimuthal angle θ with [332, 333]

$$z_n^m(\rho, \theta) = \begin{cases} R_n^m(\rho) \cos(m\theta) & \text{if } m \geq 0 \\ R_n^m(\rho) \sin(m\theta) & \text{if } m < 0, \end{cases} \quad (7.4)$$

where n and m are the radial and azimuthal orders of the polynomial respectively. The n values are non-negative integers, m are integers ranging from $-n$ to n , and $R_n^m(\rho)$ are radial polynomials defined as



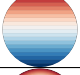
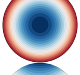
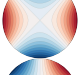
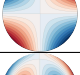
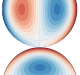
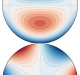
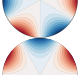
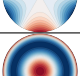
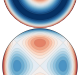
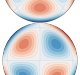
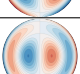
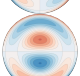
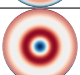

$$R_n^m(\rho) = \begin{cases} \sum_{k=0}^{\frac{n-m}{2}} \frac{(-1)^k (n-k)!}{k! (\frac{n+m}{2} - k)! (\frac{n-m}{2} - k)!} \rho^{n-2k} & \text{if } (n-m) \text{ is even} \\ 0 & \text{if } (n-m) \text{ is odd.} \end{cases} \quad (7.5)$$

These terms all describe distinct aberrations seen in the measured wavefronts arising due to differences in shape of the test optic surface with respect to the transmission sphere. A list of the aberrations corresponding to the first fifteen Zernike polynomial/coefficients, stating the aberration type, is shown in Table 7.3. The key one for measuring the radius of curvature of an optic is the ‘power’ coefficient also known as the ‘defocus’. Power is a

measure of the curvature of the surface/ wavefront without distinguishing between the horizontal (x) and vertical (y) dimensions. Its units are nanometres and it is equivalent to the height (z-displacement) difference between the centre point and the point farthest from the centre, derived from the best fit spherical surface [337].

In this study, the Zygo software was set to immediately account for the first three Zernike coefficients in all displayed surface profiles: i.e. piston (z-direction offset), and tilt (x/y-plane offset) corrections were applied, allowing for easy observations of the surface features. For objects with high curvature, power should also be accounted for, so as to not be dominated by the overall spherical shape, though this correction is not needed for the aLIGO optics with their large radii of curvature, typically around ~ 2000 m. As well as surface images, surface profiles (z-direction, height profiles) are taken by sampling linear cross sections of the measured surface height in the xy-plane from edge-to-edge passing through the centre.

Table 7.3: List of the first fifteen Zernike polynomials (by Wyant-fringe indexing [336]) used to describe the Zygo wavefront measurement/test optic surface aberrations, accompanied by a visual representation of each. Each representation is viewed in the x-y plane, with darker colours representing greater z-direction deviation, and red/blue representing negative (into page)/ positive (out from page) z displacement respectively.

| N ^o | n | m | $z_n^m(\rho, \theta)$ | Aberration | Visual |
|-----------------------|-----|-----|--|-----------------------------------|---|
| 0 | 0 | 0 | 1 | Piston |  |
| 1 | 1 | 1 | $\rho \cos(\theta)$ | Horizontal Tilt (x-axis) |  |
| 2 | 1 | -1 | $\rho \sin(\theta)$ | Vertical Tilt (y-axis) |  |
| 3 | 2 | 0 | $\rho^2 - 1$ | Power/Defocus |  |
| 4 | 2 | 2 | $\rho^2 \cos(2\theta)$ | Vertical Astigmatism |  |
| 5 | 2 | -2 | $\rho^2 \sin(2\theta)$ | Oblique Astigmatism (rotated 45°) |  |
| 6 | 3 | 1 | $(3\rho^3 - 2\rho) \cos(\theta)$ | Horizontal Coma |  |
| 7 | 3 | -1 | $(3\rho^3 - 2\rho) \sin(\theta)$ | Vertical Coma |  |
| 9 | 3 | 3 | $\rho^3 \cos(3\theta)$ | Oblique Trefoil |  |
| 10 | 3 | -3 | $\rho^3 \sin(3\theta)$ | Vertical Trefoil |  |
| 8 | 4 | 0 | $6\rho^4 - 6\rho^2 + 1$ | Primary Spherical Aberration |  |
| 11 | 4 | 2 | $(4\rho^4 - 3\rho^2) \cos(2\theta)$ | Secondary Vertical Astigmatism |  |
| 12 | 4 | -2 | $(4\rho^4 - 3\rho^2) \sin(2\theta)$ | Secondary Oblique Astigmatism |  |
| 13 | 5 | 1 | $(10\rho^5 - 12\rho^3 + 3\rho) \cos(\theta)$ | Secondary Horizontal Coma |  |
| 14 | 5 | -1 | $(10\rho^5 - 12\rho^3 + 3\rho) \sin(\theta)$ | Secondary Vertical Coma |  |
| 15 | 6 | 0 | $20\rho^6 - 30\rho^4 + 12\rho^2 - 1$ | Secondary Spherical Aberration |  |
| Higher order terms... | | | | | |

A final alignment drift check was performed during the measurement using the Zernike polynomial fits to the surface. This involved taking multiple surface profile measurements, each individually taking ~ 30 s to acquire, and assessing if any (undesired) evolution in the fourth–fifteenth Zernike coefficients (all with units of nanometres) occurred with time. Forty such surface profile images are acquired during a measurement run, with two average surface profiles; one of the first 20 measurements, and then one of the second 20 being constructed. The average of the first 20 measurements is then subtracted from the average of the second 20, which for a nominally stable stationary system, should result in no perceived changes in the wavefronts. However if the relative orientation of any of the optics changed, this would alter the fitted Zernike polynomials to the wavefronts, and upon subtraction, if any of the fourth–fifteenth Zernike coefficients were found to differ by more than 0.2 nm between the two average profiles, then there was determined to be unacceptably high alignment drift in the system. At this point the sample would then be realigned, allowed time to settle, and a new measurement run acquired until these twelve Zernike coefficients remained to within 0.2 nm on subtraction. Once achieved, the environment/optic was then determined to be sufficiently stable, and the average of all 40 of these ‘stable’ measurements were then taken as the acquired surface profile of the optic. Figure 7.10 shows a screenshot of the Zygo measurement software, showing the change of all of the desired Zernike coefficients within this 0.2 nm tolerance, post subtraction of the first dataset with the second.




Figure 7.10: Image showing the subtraction profile of the average of the first 20 surface profiles and average of second 20 surface profiles of ITM 11 orientated at 270°. Note the fourth–fifteenth Zernike terms (from Table 7.3) within the red boxes, all have values below 0.2 nm on subtraction indicating a stable measurement. The surface is corrected for both tilt and piston deviation. Note as well the black fringes and blotches near the centre, which are artefacts of the interferometer - the genesis of which is not perfectly anti-reflective internal optics distorting measurements of the phase and hence distorting surface profile measurements in those regions.

7.3.3 Removing artefacts in the Zygo data

As can be seen in Figure 7.10 some artefacts, appearing as black rings and blotches on the wavefront profile, can be observed in the surface measurements of the optic. These points arise due to not perfectly anti-reflective coated internal optics of the Zygo leading to additional light coupling into the signal and distorting measurements of the phase and hence surface profile measurements in those regions. In order to remove their influence and obtain good estimates of the entire surface profile one must measure optics at multiple roll angle orientations (where the artefacts are observed not to rotate with the test optic). Specifically, the three aLIGO ITMs being studied had measurements taken, as outlined in the previous two sections, at eight equally spaced roll angles with respect to the incoming beam. Figure 7.11 shows a diagram of the mount used to hold the ITMs during Zygo




Figure 7.11: Diagram of the mount used to hold test mass optics for Zygo measurement - taken from [338]. Some key components it possesses are raisable rollers which can lift the test mass in order to rotate it smoothly to the desired roll angle, and padded paddles for the optic to rest on during measurement, which reside on a tiltable platform inside the mount with fine control knobs to align it with the transmission sphere.

measurement. The mount possesses multiple rollers, which can be manually raised and lowered to hold the optic and rotate it to the desired roll angle, before lowering back down into the mount frame. The roll angle orientation for each test mass was defined with respect to the barrel etched, top fiducial marker, distinguished by an arrow head, such that 0° corresponded to when it resided at 12-o'clock, with 45° being when it resided near the 1-o'clock position and so on. Note even in the absence of the artefacts, multiple roll angle measurements would still be taken to eliminate the influence of any distortions/stress the mount causes on the optic. Even though the mount is designed to minimise this, the test mass optics are so heavy (~ 40 kg), they can distort under their own weight.

To facilitate matching of data acquired at different roll angles, before any measurement, three ~ 2 mm diameter circular dots of First ContactTM [241, 240] polymer were applied to the faces of the ITMs with a fine brush, approximately 5 mm in from the barrel. One was aligned with the top fiducial marker, and the others at the 90° and 225° points. These markers could be seen in the measurement wavefronts, allowing for the surface profiles at each roll angle to be rotated until the same markers overlapped in every profile. The average surface profiles at each roll angle (now realigned to match with each other) could then be compared and further averaged. This meant that for a given apparent surface region where an interferometer artefact appeared, it only appeared in that location in one-out-of-eight surface profiles (as they remain fixed upon rotation), and its impact therefore was significantly reduced once all eight roll angle measurements were averaged.

This average of all the different roll angle surface profiles is taken as the final surface profile of the optic, with the effect of the interferometer artefacts mitigated. Along with this, average profiles are also made from just the even angle (0° , 90° , 180° , 270°) roll angle measurements, and of the odd angle (45° , 135° , 225° , 315°) measurements. These can be subtracted from each other to provide an estimate of when drift has occurred in the measurement.




Figure 7.12: Surface images taken of the three ITM optics obtained with the Zygo interferometer, with corresponding surface (height) profiles shown below each. Each image represents the average of 320 images (40 taken at 8 different roll angles for each optic), with the surface profiles extracted from linear cross sections of the optic's face, approximately 300 mm in diameter, represented by the diagonal face-spanning lines in each image.




Figure 7.13: Zoomed in surface profiles from Figure 7.12 of the inner most 160 mm of each optic

7.3.4 ITM surface profile results

Following the measurement and analysis procedure previously outlined, images of the ITM surfaces were constructed, along with accompanying surface profiles, which are shown in Figure 7.12. Each image has tilt and piston deviations accounted for, in order for a level surface to be visualised, and is constructed from the average of 320 individual images (40 taken at 8 different roll angles). One can observe the high level of uniformity of the polished surfaces, with across the ~ 300 mm diameter sampled, each optic deviated in height by no greater than 4 nm. In fact, the largest deviations primarily emerge in the outermost 20 mm for each optic, where the bevel at the outer edge of the optic start influencing results in this edge region. The surface heights further in from this are typically uniform to within ~ 2 nm as can be seen further in Figure 7.13, showing zoomed in surface profiles of the innermost 160 mm of each ITM face.




Figure 7.14: Zoomed in image of the central 8 mm of ITM 11 shown previously in Figure 7.12. The faint residual interferometer artefact, its clones, and the interference they cause can be seen.

It is standard LIGO laboratory procedure to record the surface profiles of ITM and ETM optics at 300 mm diameter (full-face) and in the inner most 160 mm encompassing a region just beyond where the main aLIGO laser beam would impinge the optics. aLIGO is design specified such that the main interferometer beam will be of diameter 124 mm when making contact with the ETM front surface [328], and observing the inner most 160 mm diameter of the optic should encompass all portions of the surface which could interact with the laser within alignment tolerance. However, the main aLIGO laser will be even smaller than this, at 106 mm diameter when reaching the ITM front surface [328], this 160 mm observing region is kept the same for both families of test masses for consistency.

Each power-corrected ITM surface was by and large observed to be flat to within ~ 2 nm, but with a small bulge in the central 50 mm. In the central 8 mm region across all three optics, as well as slight increases in surface height towards the centre, there is also generally higher variability in the measured height signal. This central region can be seen to coincide with the location of an interferometer artefact as observed with Figure 7.14. Though its influence has been reduced through roll angle averaging, it is still observable faintly in the surface images, now with seven additional clones, broadly equally spaced, and with higher order fringes expanding out and interfering. This measurement feature, not of the optic, will impact the measurements of its surface, and would be most influential in this central region where the signal from all eight are overlapping - likely making them the cause for the observed excess variability within the central region on all three ITMs, though in any case it is a small effect.

7.3.5 ITM radius of curvature results

In order for a GWD such as aLIGO to function correctly the size and alignment of the beam must be well controlled. To facilitate the desired beam profile, the aLIGO ITM and ETM optics must have radii of curvature of (1934 ± 10) m [80], and (2240 ± 10) m [339] respectively. The radius of curvature of each optic must also be known precisely to within ± 3 m [80, 339]. The curvature of an optic's surface can introduce astigmatism, causing a laser beam impinging on it to focus differently in different meridians, which must be mitigated in a GWD. Any imperfections in the test mass curvature can also result in beam distortions, affecting the shape and uniformity of the reflected laser beam. As such, the full surface profiles (shown for the blank optics in the previous part) and curvature of the optics must be well characterised to be within specifications prior to installation. There also exist many effects which can hinder high power detector operation which must be accounted for. One of the major ones is thermal distortion, from which the dominant source is thermal lensing of the ITM substrates due to the laser power absorbed by the ITM coatings [340].

Coating absorption in both the ITM and ETM optics results in thermo-elastic distortions of the optics and a subsequent non-negligible change in their radii of curvature [340]. There exists thermal compensation systems (TCS) in aLIGO to help correct these distortion effects and maintain the radii of curvature, and hence maintain the beam profile within tolerance. These comprise of ring heater actuators affixed to the faces of all test masses to adjust their radii of curvature, spatially tuneable CO₂ lasers to heat the optic surfaces, as well as Hartmann wavefront sensors to monitor optical surface changes [341]. The TCS ring heaters are specified to provide radius of curvature compensation of > -116 m, such that a ± 10 m of absolute tolerance adjustment can be provided with at most 17% of this TCS dynamic range [323]. Requiring the test masses to be fabricated to within ± 10 m radii of curvature tolerance will allow the interferometer operating point sensitivity, at all meaningful power levels, with essentially full TCS dynamic reserve at maximum laser power [323]. Prior to installation each optic's radius of curvature must be known and verified to be within tolerance, with ideally, test masses of identical radius of curvature also paired together for joint operation in both arm cavities of a detector.

Using the surface profiles acquired of each ITM optic, their radii of curvature could be inferred. The surface profiles are fitted using the Zernike polynomials, and power being a measure of the curvature of the surface (equivalent to the height difference between the centre point and the point farthest from the centre of the best fit spherical surface) can be used to extract the radius of curvature. The equation for which, for these Zygo measurements, is as follows:

$$R = \frac{r_{\text{app}}^2 (R_{\text{TS}} + L_{\text{cavity}})}{r_{\text{app}}^2 + 2(R_{\text{TS}} + L_{\text{cavity}})P}, \quad (7.6)$$

where R is the radius of curvature of the test optic, R_{TS} the radius of curvature of the transmission sphere, L_{cavity} is the distance between the transmission sphere and the test optic being measured, r_{app} the radius of the circular aperture over which the data is collected, and P being the power (Zernike coefficient - see Table 7.3). In this study a 2100 m radius of curvature transmission sphere was used for all ITM measurements, and the cavity length was measured with a meter stick, with reading uncertainty of 0.5 mm, to be between 0.1–0.3 m for all measurements. The cavity length has a small but measurable influence on the overall radius of curvature, but its uncertainty can be considered negligible for the radius of curvature calculation given that the transmission sphere radius of curvature is five orders of magnitude larger. The fits of power were made over circular apertures of diameter ~ 160 mm just beyond where the main LIGO beam would impinge the ITMs when suspended in the detector.

The power fits across all 320 acquired surface images for each of the three ITMs is shown in Figure 7.15, each only differing within ~ 4 nm between all samples. The differences observed between different measurements of the same sample can arise from multiple factors. The ‘jumping’ witnessed between some successive sets of 40 measurements, which correspond to the data-sets from the different roll-angles, result from initially setting alignment to be on a bright or dark fringe: this can be seen for ITM 03 when comparing results 40–120 and the 40 measurements either side of these. The variation arises from different start phases and the slight offset this can induce on measurements. When the measurements are witnessed to gradually and continuously change, such as for most of the measurements of ITM 13, these instead are likely caused by temperature fluctuations of the lab environment, which results in slight changes in the curvature of the optics. Lastly




Figure 7.15: Fitted power values from the central 160 mm diameter of the three ITM optics, produced from 320 total surface images at 8 different roll angles.

when one set of 40 measurements is seen to have a constant gradient in the fitted power data, it could suggest the test mass has not settled in that mount position, an example of this can be seen for ITM 03 between measurements 200–240. However a result like this could result from a constantly increasing or decreasing temperature and decoupling the two factors in this case is not so trivial from power data alone. Ultimately all these effects are quite small and do not alter the power estimates by more than a few nanometres for a given measurement, which as can be seen later also does not lead to large curvature errors.

Taking the minimum, maximum, and mean of the power acquired for each ITM, Equation 7.6 can be used to infer their corresponding maximum, minimum, and mean radii of curvature. Also taking the standard deviation of the power data and adding it as an offset to the mean power value, and carrying this through to calculating the radius of curvature, and then subtracting the mean radius of curvature value, allows for an estimate of the environmental stability of the measurement in terms of radius of curvature change. These values are all reported in Table 7.4. One can see all three ITMs have identical radii of

curvature to each other, within error, and all meet the stringent aLIGO ITM radius of curvature specification, with all values falling between 1934 ± 10 m, and each with radius of curvature precision of better than ± 3 m, with all values obtained for a given optic differing by less than ± 2 m.

ITM 03 and ITM 11 each exhibit identical radii of curvature to their previous values prior to being coated, and then stripped of their coating in the re-polishing process. They are identical within the measurement spread, though the mean radius of curvature of ITM 11 now better matches with ITM 13 than it previously did with ITM 03, so if wanting to pair ITM optics for simultaneous use in a detector, to best approximation, ITM 11 and ITM 13 are the most suiting pair of the three, though all three are viable partners.

Further tests were initiated in confirming the radii of curvature of these ITMs with extended measurement apertures up to the full face of the optic (~ 300 mm diameter) with results not changing from those quoted in Table 7.4 by greater than 2 m for any test mass, highlighting the quality of the optics and polishing process. Similarly, multiple, cross sections were sampled with radii of curvature inferred, all yielding the same results as these within 2 m, and each again with low spreads of < 2 m. In general the re-polishing process has proven quite successful in maintaining ITM surface profiles/radii of curvature within specification without a reduction in surface quality.

Table 7.4: Radii of curvature R of the three ITM optics acquired with the Zygo interferometer. ΔR represents the environmental uncertainty in the measurement calculated from the standard deviation in the power data carried forward as a change in mean radius of curvature it would induce. The values pre re-polishing of ITM 03 [342] and ITM 11 [343] are shown as well for comparison.

| Measured ITM radii of curvature with key values | | | |
|---|------------------------|------------------------|------------------------|
| Property | ITM 03 | ITM 11 | ITM 13 |
| R_{TS} [m] | 2100 | 2100 | 2100 |
| L_{cavity} [m] | 0.169 | 0.170 | 0.190 |
| r_{app} [mm] | 80.2 | 80.2 | 80.2 |
| P [nm] | $127.0^{+1.5}_{-1.7}$ | $127.5^{+1.3}_{-1.5}$ | $127.6^{+1.7}_{-1.3}$ |
| R [m] | $1939.4^{+1.9}_{-1.8}$ | $1938.8^{+1.7}_{-1.6}$ | $1938.6^{+1.8}_{-1.9}$ |
| ΔR [m] | 1.0 | 0.8 | 0.8 |
| Old ITM radii of curvature pre re-polishing | | | |
| R_{OLD} [m] | $1939.3^{+0.7}_{-0.5}$ | $1939.2^{+0.5}_{-0.5}$ | |

In fact, in the case of ITM 11, a large surface feature, a few nanometres in height spanning the outer edge of the optic, previously present, was able to be removed upon re-polishing as can be seen in Figure 7.16. This particular feature would not affect detector performance, as it was well outside the main laser beam area, though should similar features emerge in future optics in problematic locations, these results give confidence in the re-polishing process to produce a high quality surface not impacted by defects previously present.

Ultimately all three blank ITM radii of curvature were found to lie within specifications, and close to the centre of the specification range. The deposition of coatings onto the aLIGO ITM and ETM substrates is known to alter the radius of curvature by typically <5 m so they will likely remain within specifications once they are coated as others have historically [344]. Once coated, each optic will of course need its surface re-characterised. A further requirement for the ITM coatings that will eventually be deposited on these test masses, is that the surface aberrations are not altered greatly from these high quality blank results, and based on current aLIGO constraints, they specifically must not add surface figure Zernike coefficients higher than second order with amplitude > 0.5 nm [345]. Having these substrate radius of curvature and surface profile results to compare to, will be invaluable for LIGO colleagues in assessing the figure properties of the future coatings, and inferring how they interact with the substrate.




Figure 7.16: Images of the surfaces of ITM 11 captured in the Zygo setup before installation in LIGO for operation in observing run 1 (old) and the ones taken as part of this project after being stripped of its coating. A large curved feature can be seen in the old data spanning the face of the optic in the upper north region, which has been removed upon re-polishing. The plot on the far right shows the difference in the surface profiles between the two data-sets, for the cross section shown in the surface images (d1) and of a surface profile taken tangentially to this and still cutting through the centre of the optic(d2).

7.4 Conclusions from aLIGO test mass studies

Through the work undertaken on this research placement at the California Institute of Technology LIGO Core Optics Components subgroup, many different features of past and future aLIGO test masses were characterised.

The surface profiles, and radii of curvature of three uncoated ITM test masses were characterised, with all being found to lie within design tolerance, all with $\sim 1939^{+2}_{-2}$ m. Two of these (ITM 11 and ITM 13) were previously installed at the LIGO Hanford detector and had since been stripped of their coating, and needed re-characterised for future upgrades. These results also served as a verification of the re-polishing process, showing that the radii of curvature of both these ITMs was unaltered within $\sim 0.1\%$ measurement uncertainty, from their uncoated results prior to being installed at Hanford for the start of observing run 1. The success of the re-polishing at yielding results un-influenced by any previous surface malformation was also demonstrated.

Alongside this work, a novel study involving mapping the absorption of GWD optics at the end of their lifespans was also launched. Specifically the two ETMs installed at LIGO Livingston through observing runs 1–3 were mapped in a PCI setup, with the location and relative magnitude of their point absorbers determined. After this, an optical microscope survey was launched to better understand the nature of the point absorbers, and a new cleaning procedure showing great potential to greatly reduce the influence of the point absorbers was developed. Many of the seemingly highest absorption point absorbers were found to arise from contamination of the optic surface. Of these, many were able to be removed with the new cleaning procedure, with those that still remained / partially remained, having their level of absorption drop by at least 75%. This study provided a potential path forward for mitigating the influence of point absorbers in future detectors, which are currently a major limiting factor for GWDs, and whose impact will continue to increase as detectors strive towards higher operating powers.

Chapter 8

Thesis conclusions

The first direct detection of gravitational waves in 2015 was one of the great feats of physics. Almost one decade on, we are now firmly in the advent of gravitational wave astronomy, where over 90 gravitational wave signals have been detected, and with knowledge gained from these observations having wide-spanning ramifications for the field of physics. These detections have ushered in the era of multi-messenger astronomy, freeing us from the confines of only electromagnetic surveys of astronomical events, and previously poorly characterised astrophysical systems, such as binary black hole and binary neutron star mergers have been able to be much better understood. These detections have also been able to provide validations of Einstein’s general theory of relativity that would not have otherwise been possible. Through all this our understanding of our Universe has been broadened, and we have thus been able to unlock many of its mysteries.

These detections have been made possible by the significant advances in technology and experimental techniques achieved over the last several decades, resulting in large scale laser interferometric detectors with the sensitivity required to measure relative displacements of the order of an atomic nucleus. They achieve this by continuously monitoring the relative position between highly isolated “test mass” mirrors suspended at great distances from each other. These gravitational wave detectors can currently achieve remarkable strain sensitivities of up to $\sim \times 10^{-24} \text{ Hz}^{-\frac{1}{2}}$, but are limited by a plethora of noise sources which reduce their sensitivity at different frequencies. Chief among these, in the most sensitive detection band of current detectors, is the thermal noise associated with their highly reflective mirror coatings.

The research presented in this thesis has focussed on measuring many key mechanical and optical properties of new and current coating materials. The primary goal was to find suitable new coating materials and coating designs with potential to improve upon the thermal noise levels in current detectors, for use in the very next round of upgrades, whilst still meeting the stringent optical requirements of the detectors. Studies have involved measuring, for different coatings, the key property influencing coating thermal noise: the mechanical loss, and from this, calculating the level of improvement they could yield. Substantial investigations were also carried out to characterise the optical properties and structure of one promising new material: $\text{TiO}_2:\text{SiO}_2$.

Studies of the current detector coatings were also conducted. Chapter 7 details work undertaken into mapping topographical features of various aLIGO test mass optics. Specifically the author mapped the optical absorption of the coatings of both the end test mass mirrors installed in the LIGO Livingston detector for the first-through-third observing runs. This was carried out via photothermal-commonpath interferometry, and from these maps the precise locations of various point-absorbers were identified. Point absorbers were a significant problem during the previous observing run, limiting detector performance. Subsequent microscope surveys provided insights into these point absorbers, with those with the largest magnitude absorption seeming to originate from points of damage in the coating, of which only 1–2 per mirror were found. Apart from these, many of the other high absorption point absorbers were found to originate from contamination, with a new cleaning method trialled to remove them. This yielded promising results for all the point absorbers it was tested on, seemingly reducing their total absorption levels by at least 75%, and some (notably the second highest absorption point absorber) even being removed from the mirror surface entirely. As such, this study provided a potential avenue for mitigating the influence of point absorbers going forward, which is of particular importance with the circulating laser beam powers in future gravitational wave detectors set to increase. Alongside these studies, the surface profiles of uncoated input test mass optics to be coated in the future were characterised and their curvatures found to be within tolerance of the LIGO detector specifications for future upgrades.

The two candidate coating materials for future upgrades primarily focused on in this thesis are $\text{TiO}_2:\text{SiO}_2$ and $\text{TiO}_2:\text{GeO}_2$; both are low-loss amorphous oxide candidates to replace the high-refractive index material of current detector mirrors $\text{TiO}_2:\text{Ta}_2\text{O}_5$. These investigations concluded that both are highly promising future materials.

Investigations of mixes of $\text{TiO}_2:\text{SiO}_2$ deposited by FiveNine Optics, showed that a highly-reflective stack consisting of a particularly promising mixture of $(69.5\pm 1.3)\%$ TiO_2 , $\text{TiO}_2:\text{SiO}_2$ as the high-refractive index material and SiO_2 as its low-index partner yielded optimal CTN levels of around 75% of current detector levels, whilst having ~ 0.82 ppm optical absorption and ~ 5 ppm scatter at 1064 nm, and whilst remaining free from defects after post deposition heat treatment at 850°C for a duration of 100 hours.

The exceptionally low optical losses were measured via several techniques. The absorption was verified via PCI on different setups with different beam intensities showing consistent sub-ppm absorption, and these were also compared to results acquired from a high finesse optical cavity setup, also showing a general agreement. The scatter was measured with the use of an integrating sphere and additionally with an angle-resolved-scatterometer, with both setups independently showing the scatter after heat treatment to be ~ 5 ppm, which is around a factor of two lower than current gravitational wave detector coatings measured in the same integrating sphere setup. These very encouraging optical loss results were, interestingly, achieved after the coatings had crystallised. Specifically, for the 69.5% TiO_2 mixture, the onset of crystallisation was found to begin at 575°C and by 850°C heat treatment for 100 hours, they had formed nanocrystallites 6-10 nm in size, which could be measured at all areas of the coating.

The CTN level of this coating was determined via two different methods: direct measurement in a folded-cavity setup and through measurement of the coating mechanical loss. The two techniques gave excellent agreement up until 600°C where defects in the form of blisters and cracks were seen to form on the loss samples. The direct CTN samples, on the other hand, did not manifest defects until a much higher temperature of 950°C ,

100°C beyond the optimum CTN heat treatment temperature. Whilst the origin of these defects on some samples requires further investigation, the fact that neither of the samples used for direct CTN measurement manifested any defects shows great promise for this material.

The other coating candidate investigated extensively was $\text{TiO}_2:\text{GeO}_2$. Specifically, investigations of the mechanical loss of various single layers, and of coating stack designs, using SiO_2 as a low index partner material, were carried out. Initial prototypes of these coatings from the current gravitational wave detector coating manufacturer LMA were measured, from which an improvement in loss of the single layers by around a factor of $\times 4$ with respect to current coatings was found. However, these single layer coatings were found to form blisters after the final heat treatment at a temperature of 600°C for 108 hours.

Complementary studies in collaboration with Caltech, Colorado State University, and LMA were undertaken to investigate the effect of deposition parameters on suppressing the formation of these blisters. Through this, it was verified that elevated deposition temperatures of at least 150°C and chamber base pressures of $\sim \times 10^{-7}$ mbar prevented blistering upon heat treatment in both stack and single layer coatings. The defect-free CSU $\text{TiO}_2:\text{GeO}_2$ was found to yield equivalent losses to the LMA coatings, but with slightly lower observed losses at low frequency. For all cases $\text{TiO}_2:\text{GeO}_2$ single layers were found to exhibit a strong frequency-dependent loss. The CSU coatings were estimated to yield loss levels as low as $(18.3 \pm 1.7)\%$ of the current detector coatings at 100 Hz where CTN dominates detector performance.

Various stack designs incorporating the optimised $\text{TiO}_2:\text{GeO}_2$ deposition recipe were produced. Two stacks meeting the full reflectivity requirement for gravitational wave detectors end test mass mirrors were deposited and their losses measured. One had a uniform layer structure (52L-A). The other had varying layer thicknesses to have progressively less of the higher absorption $\text{TiO}_2:\text{GeO}_2$ near the top of the stack (52L-B), thus reducing the amount of light interacting with it and decreasing the total absorption of the stack. Both of these designs were found to exhibit significantly less frequency dependent loss than the single layers and the thinner stacks comprised of $\text{TiO}_2:\text{GeO}_2$ and SiO_2 . Studies of the

effect of the number of layer interfaces, relative proportions of the two materials and total thickness on the frequency dependence of the losses were conducted, indicating that total stack thickness rather than number of interfaces seemed to be the primary driving factor for the excess loss witnessed at lower frequencies.

Ultimately, the CTN level of both these full reflectivity stacks was able to be estimated at $64_{-5}^{+8}\%$ and $63_{-2}^{+5}\%$ of current aLIGO levels, which was corroborated to be in line with direct CTN measurements of these coatings. These are exceptionally low CTN results, for a full reflectivity coating, and one free of defects. However, owing to their higher loss trends at low frequencies compared to single layers, the results are significantly higher than was predicted from previous single-layer based estimates of achieving 45% the CTN of current aLIGO.

SiO₂ single layers made by the current gravitational wave detector coating manufacturer LMA were investigated, yielding significantly improved loss results as a result of higher heat treatment temperature and duration versus current coatings. Specifically an average loss of $(8.0 \pm 0.7) \times 10^{-6}$ was measured for single layers of SiO₂ produced in LMA's Grand Coater system after being heat treated at 600°C, for 108 hours. When compared with the loss after a 500°C heat treatment for 10 hours – the current standard heat treatment recipe – this represented an overall reduction in loss of 76%. This shows the potential for significant benefits from extended duration heat treatment in these, and likely other, coatings.

Whilst producing very low CTN, the particular coating runs of TiO₂:GeO₂ investigated in this thesis were shown by collaborators to exhibit optical absorption levels of a few ppm, when optimised by having thinner TiO₂:GeO₂ in the top-most layers of the stack. This absorption level is unfortunately too high for adoption in current detectors. However, a multimaterial design comprised of both the primary materials investigated in this work may be of interest to exploit the most beneficial properties of both materials.

Both the TiO₂:GeO₂ and TiO₂:SiO₂ based coatings are highly promising candidates for the next upgrades of GWDs, with exceptionally low CTN. The absorption of the TiO₂:GeO₂ could perhaps be mitigated with the top few layers of TiO₂:GeO₂ replaced with much lower absorbing and much lower scattering TiO₂:SiO₂. An alternative design

could be made by swapping the top few layers of $\text{TiO}_2\text{:GeO}_2$ with the currently-used $\text{TiO}_2\text{:Ta}_2\text{O}_5$, but this is potentially less beneficial from a scattering perspective. Such multimaterial designs could also be used to enable the high absorption of SiN_x based coatings to be mitigated. $\text{TiO}_2\text{:SiO}_2$ and SiN_x have closely matched optimal heat treatment temperatures, suggesting this could be a promising multimaterial combination using $\text{TiO}_2\text{:SiO}_2$ in the top of the stack and SiN_x in the bottom of the stack [165, 168].

However, the results presented here for both $\text{TiO}_2\text{:SiO}_2$ and $\text{TiO}_2\text{:GeO}_2$ on their own merit further investigations into optimisation of both the deposition and post-deposition heat treatment procedures, for instance significantly longer duration anneals could negate the necessity for multimaterial topologies. Given the recent progress with both materials, it seems highly likely that one or both of these materials will be adopted in the imminent upgrades planned for Advanced LIGO and Advanced Virgo, and possibly for the planned room temperature Einstein Telescope and Cosmic Explorer detectors.

Bibliography

- [1] A. Einstein. “The Foundation of the General Theory of Relativity”. In: *Annalen der Physik* 49 (1916), pp. 769–822. URL: <https://einsteinpapers.press.princeton.edu/vol6-trans/158>.
- [2] A. Einstein. “On the Relativity Principle and the Conclusions Drawn from It”. In: *Jahrbuch der Radioaktivität* (1907), pp. 411–462. DOI: 10.1093/ajae/aaq063. URL: <https://einsteinpapers.press.princeton.edu/vol2-trans/266>.
- [3] A. Einstein. “On the Influence of Gravitation on the Propagation of Light”. In: *Annalen der Physik* 35 (1911), pp. 898–908. DOI: 10.1093/ajae/aaq063. URL: <https://einsteinpapers.press.princeton.edu/vol3-trans/393>.
- [4] A. Einstein. “The Formal Foundation of the General Theory of Relativity”. In: *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)* 1914 (1914), pp. 1030–1085. URL: <https://einsteinpapers.press.princeton.edu/vol3-trans/393>.
- [5] A. Einstein. “On the General Theory of Relativity”. In: *Royal Prussian Academy of Sciences (Berlin). Proceedings* (1915), pp. 778–786. URL: <https://einsteinpapers.press.princeton.edu/vol6-trans/110>.
- [6] A. Einstein. “On the General Theory of Relativity (Addendum)”. In: *Royal Prussian Academy of Sciences (Berlin). Meeting reports* (1915), pp. 799–801. URL: <https://einsteinpapers.press.princeton.edu/vol6-trans/120>.
- [7] A. Einstein. “Explanation of the Perihelion Motion of Mercury from the General Theory of Relativity”. In: *Royal Prussian Academy of Sciences (Berlin). Meeting reports* (1915), pp. 831–839. URL: <https://einsteinpapers.press.princeton.edu/vol6-trans/124>.
- [8] A. Einstein. “The Field Equations of Gravitation”. In: *Royal Prussian Academy of Sciences (Berlin). Meeting reports* (1915), pp. 844–847. URL: <https://einsteinpapers.press.princeton.edu/vol6-trans/129>.
- [9] J.A. Wheeler. *A Journey Into Gravity and Spacetime*. Library series. Scientific American Library, 1990. ISBN: 9780716750161. URL: https://books.google.co.uk/books?id=Wn_FQgAACAAJ.

- [10] J.A. Wheeler and K. Ford. *Geons, Black Holes, and Quantum Foam: A Life in Physics*. W. W. Norton, 2010. ISBN: 9780393079487. URL: <https://books.google.co.uk/books?id=zGFkK2tTXPsC>.
- [11] Patrick R. Brady and Jolien D.E. Creighton. “Gravitational Wave Astronomy”. In: *Encyclopedia of Physical Science and Technology (Third Edition)*. Ed. by Robert A. Meyers. Third Edition. New York: Academic Press, 2003, pp. 33–48. ISBN: 978-0-12-227410-7. DOI: <https://doi.org/10.1016/B0-12-227410-5/00299-4>. URL: <https://www.sciencedirect.com/science/article/pii/B0122274105002994>.
- [12] Matthew Pitkin et al. “Gravitational Wave Detection by Interferometry (Ground and Space)”. In: *Living Reviews in Relativity* 14.1 (2011), p. 5. ISSN: 1433-8351. DOI: 10.12942/lrr-2011-5. URL: <https://doi.org/10.12942/lrr-2011-5>.
- [13] Éanna É Flanagan and Scott A Hughes. “The basics of gravitational wave theory”. In: *New Journal of Physics* 7.1 (2005), p. 204. DOI: 10.1088/1367-2630/7/1/204. URL: <https://dx.doi.org/10.1088/1367-2630/7/1/204>.
- [14] R. A. Hulse and J. H. Taylor. “Discovery of a pulsar in a binary system.” In: 195 (Jan. 1975), pp. L51–L53. DOI: 10.1086/181708.
- [15] Joel M. Weisberg, Joseph H. Taylor and Lee A. Fowler. “Gravitational Waves from an Orbiting Pulsar”. In: *Scientific American* 245.4 (1981), pp. 74–83. ISSN: 00368733, 19467087. URL: <http://www.jstor.org/stable/24964580> (visited on 18/01/2024).
- [16] J. H. Taylor and J. M. Weisberg. “A new test of general relativity - Gravitational radiation and the binary pulsar PSR 1913+16”. In: *Astrophysical Journal* 253 (Feb. 1982), pp. 908–920. DOI: 10.1086/159690.
- [17] Nobel Media AB. *Nobel Prize Announcement*. <https://www.nobelprize.org/prizes/physics/1993/summary/>. Accessed on: 2024-02-21. 1993.
- [18] B. P. Abbott et al. “Observation of Gravitational Waves from a Binary Black Hole Merger”. In: *Physical Review Letters* 116 (6 2016), p. 061102. DOI: 10.1103/PhysRevLett.116.061102. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.061102>.

- [19] B. P. Abbott et al. “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral”. In: *Physical Review Letters* 119 (16 2017), p. 161101. DOI: 10.1103/PhysRevLett.119.161101. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.119.161101>.
- [20] B. P. Abbott et al. “Properties of the Binary Neutron Star Merger GW170817”. In: *Physical Review X* 9 (1 2019), p. 011001. DOI: 10.1103/PhysRevX.9.011001. URL: <https://link.aps.org/doi/10.1103/PhysRevX.9.011001>.
- [21] B. P. Abbott et al. “A gravitational-wave standard siren measurement of the Hubble constant”. In: *Nature* 551.7678 (Oct. 2017), pp. 85–88. ISSN: 1476-4687. DOI: 10.1038/nature24471. URL: <http://dx.doi.org/10.1038/nature24471>.
- [22] B. P. Abbott et al. “Tests of General Relativity with GW150914”. In: *Physical Review Letters* 116 (22 2016), p. 221101. DOI: 10.1103/PhysRevLett.116.221101. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.221101>.
- [23] B. P. Abbott et al. “Tests of General Relativity with GW170817”. In: *Physical Review Letters* 123 (1 2019), p. 011102. DOI: 10.1103/PhysRevLett.123.011102. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.011102>.
- [24] B. P. Abbott et al. “Tests of general relativity with the binary black hole signals from the LIGO-Virgo catalog GWTC-1”. In: *Physical Review D* 100 (10 2019), p. 104036. DOI: 10.1103/PhysRevD.100.104036. URL: <https://link.aps.org/doi/10.1103/PhysRevD.100.104036>.
- [25] R. Abbott et al. “Tests of general relativity with binary black holes from the second LIGO-Virgo gravitational-wave transient catalog”. In: *Physical Review D* 103 (12 2021), p. 122002. DOI: 10.1103/PhysRevD.103.122002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.103.122002>.
- [26] The LIGO Scientific Collaboration et al. *Tests of General Relativity with GWTC-3*. 2021. arXiv: 2112.06861 [gr-qc].
- [27] B. P. Abbott et al. “GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence”. In: *Physical Review Letters* 116 (24 2016), p. 241103. DOI: 10.1103/PhysRevLett.116.241103. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.241103>.

- [28] R. Abbott et al. “Observation of Gravitational Waves from Two Neutron Star–Black Hole Coalescences”. In: *The Astrophysical Journal Letters* 915.1 (2021), p. L5. DOI: 10.3847/2041-8213/ac082e. URL: <https://dx.doi.org/10.3847/2041-8213/ac082e>.
- [29] R. Abbott et al. “GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo during the Second Part of the Third Observing Run”. In: *Physical Review X* 13 (4 2023), p. 041039. DOI: 10.1103/PhysRevX.13.041039. URL: <https://link.aps.org/doi/10.1103/PhysRevX.13.041039>.
- [30] Todd A. Thompson, Matthew D. Kistler and K. Z. Stanek. *A High Rate of White Dwarf-Neutron Star Mergers Their Transients*. 2009. arXiv: 0912.0009 [astro-ph.HE].
- [31] P. Lorén-Aguilar et al. “Gravitational wave radiation from the coalescence of white dwarfs”. In: *Monthly Notices of the Royal Astronomical Society* 356.2 (Jan. 2005), pp. 627–636. ISSN: 0035-8711. DOI: 10.1111/j.1365-2966.2004.08472.x. eprint: <https://academic.oup.com/mnras/article-pdf/356/2/627/3976470/356-2-627.pdf>. URL: <https://doi.org/10.1111/j.1365-2966.2004.08472.x>.
- [32] Laurence E. H. Datrier. “Exploring the calibration of cosmological probes used in gravitational-wave and multi-messenger astronomy”. Available at <https://theses.gla.ac.uk/83608/>. PhD thesis. University of Glasgow, 2023.
- [33] Paul D. Lasky. “Gravitational Waves from Neutron Stars: A Review”. In: *Publications of the Astronomical Society of Australia* 32 (2015), e034. DOI: 10.1017/pasa.2015.35.
- [34] K. A. Thorne. *Searching for Gravitational-Wave Bursts with LIGO*. 2007. arXiv: 0706.4301 [gr-qc].
- [35] J. Aasi et al. “First low frequency all-sky search for continuous gravitational wave signals”. In: *Physical Review D* 93 (4 2016), p. 042007. DOI: 10.1103/PhysRevD.93.042007. URL: <https://link.aps.org/doi/10.1103/PhysRevD.93.042007>.
- [36] Brian Abbott. “*The Digital Universe Guide*”. Tech. rep. New York, NY, USA: American Museum of Natural History, 2012. URL: https://web.archive.org/web/20130317111255/http://www.haydenplanetarium.org/downloads/universe/Digital_Universe_Guide.pdf.

- [37] The LIGO scientific collaboration. “Introduction to LIGO and gravitational waves - stochastic gravitational waves”. Ed. by ligo.org. 2024. URL: <https://www.ligo.org/science/GW-Stochastic.php>.
- [38] Nick van Remortel, Kamiel Janssens and Kevin Turbang. “Stochastic gravitational wave background: Methods and implications”. In: *Progress in Particle and Nuclear Physics* 128 (Jan. 2023), p. 104003. ISSN: 0146-6410. DOI: 10.1016/j.pnnp.2022.104003. URL: <http://dx.doi.org/10.1016/j.pnnp.2022.104003>.
- [39] G. González and J. Pullin. “Gravitational Waves”. In: *Encyclopedia of Mathematical Physics*. Ed. by Jean-Pierre Francoise, Gregory L. Naber and Tsou Sheung Tsun. Oxford: Academic Press, 2006, pp. 582–586. ISBN: 978-0-12-512666-3. DOI: <https://doi.org/10.1016/B0-12-512666-2/00033-X>. URL: <https://www.sciencedirect.com/science/article/pii/B012512666200033X>.
- [40] B. P. Abbott et al. “GW170817: Implications for the Stochastic Gravitational-Wave Background from Compact Binary Coalescences”. In: *Physical Review Letters* 120.9 (Feb. 2018). ISSN: 1079-7114. DOI: 10.1103/physrevlett.120.091101. URL: <http://dx.doi.org/10.1103/PhysRevLett.120.091101>.
- [41] Gabriella Agazie et al. “The NANOGrav 15 yr Data Set: Evidence for a Gravitational-wave Background”. In: *The Astrophysical Journal Letters* 951.1 (2023), p. L8. DOI: 10.3847/2041-8213/acdac6. URL: <https://dx.doi.org/10.3847/2041-8213/acdac6>.
- [42] Fredrick A. Jenet and Joseph D. Romano. “Understanding the gravitational-wave Hellings and Downs curve for pulsar timing arrays in terms of sound and electromagnetic waves”. In: *American Journal of Physics* 83.7 (July 2015), pp. 635–645. ISSN: 0002-9505. DOI: 10.1119/1.4916358. eprint: https://pubs.aip.org/aapt/ajp/article-pdf/83/7/635/13081100/635_1_online.pdf. URL: <https://doi.org/10.1119/1.4916358>.
- [43] R. W. Hellings and G. S. Downs. “Upper limits on the isotropic gravitational radiation background from pulsar timing analysis.” In: 265 (Feb. 1983), pp. L39–L42. DOI: 10.1086/183954.
- [44] *LIGO’s Interferometer (IFO)*. <https://www.ligo.caltech.edu/page/ligos-ifo>. Accessed: [2024-02-19. LIGO-Caltech, 2024.

- [45] Kentaro Somiya. “Detector configuration of KAGRA—the Japanese cryogenic gravitational wave detector”. In: *Classical and Quantum Gravity* 29.12 (2012), p. 124007. DOI: 10.1088/0264-9381/29/12/124007. URL: <https://dx.doi.org/10.1088/0264-9381/29/12/124007>.
- [46] B. P. Abbott et al. “GW150914: The Advanced LIGO Detectors in the Era of First Discoveries”. In: *Physical Review Letters* 116 (13 2016), p. 131103. DOI: 10.1103/PhysRevLett.116.131103. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.131103>.
- [47] Andrew P. Spencer. “Advanced techniques in laser interferometry for current and future gravitational wave detectors”. Available at <https://theses.gla.ac.uk/81266/>. PhD thesis. University of Glasgow, 2020.
- [48] K. A. Strain and B. J. Meers. “Experimental demonstration of dual recycling for interferometric gravitational-wave detectors”. In: *Physical Review Letters* 66 (11 1991), pp. 1391–1394. DOI: 10.1103/PhysRevLett.66.1391. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.66.1391>.
- [49] G. Heinzl et al. “Experimental Demonstration of a Suspended Dual Recycling Interferometer for Gravitational Wave Detection”. In: *Physical Review Letters* 81 (25 1998), pp. 5493–5496. DOI: 10.1103/PhysRevLett.81.5493. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.81.5493>.
- [50] Michael Zucker. “*Vacuum Requirements for Gravitational Wave Interferometers*”. California Institute of Technology. 9th Jan. 2008. URL: https://indico.cern.ch/event/660772/contributions/2696462/attachments/1517339/2368521/Zucker_CERN_1Sep17_G1701685-v2.pdf.
- [51] L. Barsotti et al. “*The updated Advanced LIGO design curve*”. Tech. rep. The LIGO Scientific Collaboration, 2018. URL: <https://dcc.ligo.org/LIGO-T1800044/public>.
- [52] D. V. Martynov et al. “Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy”. In: *Physical Review D* 93.11 (June 2016). ISSN: 2470-0029. DOI: 10.1103/physrevd.93.112004. URL: <http://dx.doi.org/10.1103/PhysRevD.93.112004>.

- [53] Diego Bersanetti et al. “Advanced Virgo: Status of the Detector, Latest Results and Future Prospects”. In: *Universe* 7.9 (2021). ISSN: 2218-1997. DOI: 10.3390/universe7090322. URL: <https://www.mdpi.com/2218-1997/7/9/322>.
- [54] M. G. Beker et al. “Improving the sensitivity of future GW observatories in the 1–10 Hz band: Newtonian and seismic noise”. In: *General Relativity and Gravitation* 43.2 (2011), pp. 623–656. ISSN: 1572-9532. DOI: 10.1007/s10714-010-1011-7. URL: <https://doi.org/10.1007/s10714-010-1011-7>.
- [55] Jennifer C. Driggers, Jan Harms and Rana X. Adhikari. “Subtraction of Newtonian noise using optimized sensor arrays”. In: *Physical Review D* 86 (10 2012), p. 102001. DOI: 10.1103/PhysRevD.86.102001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.86.102001>.
- [56] M Coughlin et al. “Towards a first design of a Newtonian-noise cancellation system for Advanced LIGO”. In: *Classical and Quantum Gravity* 33.24 (2016), p. 244001. DOI: 10.1088/0264-9381/33/24/244001. URL: <https://dx.doi.org/10.1088/0264-9381/33/24/244001>.
- [57] Iain W. Martin. “Studies of materials for use in future interferometric gravitational wave detectors”. Available at <https://theses.gla.ac.uk/1517/>. PhD thesis. University of Glasgow, 2009.
- [58] M.K.M. Bader. “Seismic and Newtonian noise modeling for Advanced Virgo and Einstein Telescope”. English. PhD-Thesis - Research and graduation internal. Vrije Universiteit Amsterdam, 2021.
- [59] Irene Fiori et al. “Environmental Noise in Gravitational-Wave Interferometers”. In: *Handbook of Gravitational Wave Astronomy*. Ed. by Cosimo Bambi, Stavros Katsanevas and Konstantinos D. Kokkotas. Singapore: Springer Singapore, 2020, pp. 1–72. ISBN: 978-981-15-4702-7. DOI: 10.1007/978-981-15-4702-7_10-1. URL: https://doi.org/10.1007/978-981-15-4702-7_10-1.
- [60] S M Aston et al. “Update on quadruple suspension design for Advanced LIGO”. In: *Classical and Quantum Gravity* 29.23 (2012), p. 235004. DOI: 10.1088/0264-9381/29/23/235004. URL: <https://dx.doi.org/10.1088/0264-9381/29/23/235004>.

- [61] F Matichard et al. “Seismic isolation of Advanced LIGO: Review of strategy, instrumentation and performance”. In: *Classical and Quantum Gravity* 32.18 (2015), p. 185003. DOI: 10.1088/0264-9381/32/18/185003. URL: <https://dx.doi.org/10.1088/0264-9381/32/18/185003>.
- [62] Jim Hough, Sheila Rowan and B S Sathyaprakash. “The search for gravitational waves”. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 38.9 (2005), S497. DOI: 10.1088/0953-4075/38/9/004. URL: <https://dx.doi.org/10.1088/0953-4075/38/9/004>.
- [63] W. Heisenberg. “Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik”. In: *Zeitschrift für Physik* 43.3 (1927), pp. 172–198. ISSN: 0044-3328. DOI: 10.1007/BF01397280. URL: <https://doi.org/10.1007/BF01397280>.
- [64] Stefan Hild. “A Basic Introduction to Quantum Noise and Quantum-Non-Demolition Techniques”. In: *Advanced Interferometers and the Search for Gravitational Waves: Lectures from the First VESF School on Advanced Detectors for Gravitational Waves*. Ed. by Massimo Bassan. Cham: Springer International Publishing, 2014, pp. 291–314. ISBN: 978-3-319-03792-9. DOI: 10.1007/978-3-319-03792-9_11. URL: https://doi.org/10.1007/978-3-319-03792-9_11.
- [65] Carlo Beenakker and Christian Schönenberger. “Quantum Shot Noise”. In: *Physics Today* 56.5 (May 2003), pp. 37–42. ISSN: 1945-0699. DOI: 10.1063/1.1583532. URL: <http://dx.doi.org/10.1063/1.1583532>.
- [66] J. Abadie et al. “A gravitational wave observatory operating beyond the quantum shot-noise limit”. In: *Nature Physics* 7.12 (2011), pp. 962–965. ISSN: 1745-2481. DOI: 10.1038/nphys2083. URL: <https://doi.org/10.1038/nphys2083>.
- [67] T. P. Purdy, R. W. Peterson and C. A. Regal. “Observation of Radiation Pressure Shot Noise on a Macroscopic Object”. In: *Science* 339.6121 (Feb. 2013), pp. 801–804. ISSN: 1095-9203. DOI: 10.1126/science.1231282. URL: <http://dx.doi.org/10.1126/science.1231282>.
- [68] Carlton M. Caves. “Quantum-Mechanical Radiation-Pressure Fluctuations in an Interferometer”. In: *Physical Review Letters* 45 (2 1980), pp. 75–79. DOI: 10.1103/PhysRevLett.45.75. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.45.75>.

- [69] Rana Adhikari. “Sensitivity and Noise Analysis of 4 km Laser Interferometric Gravitational Wave Antennae”. Available at https://gwic.ligo.org/assets/docs/theses/Adhikari_Thesis.pdf. PhD thesis. MASSACHUSETTS INSTITUTE OF TECHNOLOGY, 2004.
- [70] V. B. Braginsky and F. Ya. Khalili. “Quantum nondemolition measurements: the route from toys to tools”. In: *Reviews of Modern Physics* 68 (1 1996), pp. 1–11. DOI: 10.1103/RevModPhys.68.1. URL: <https://link.aps.org/doi/10.1103/RevModPhys.68.1>.
- [71] Thomas Corbitt and Nergis Mavalvala. “Quantum noise in gravitational wave interferometers: an overview and recent developments”. In: *Fluctuations and Noise in Photonics and Quantum Optics*. Ed. by Derek Abbott, Jeffrey H. Shapiro and Yoshihisa Yamamoto. SPIE, May 2003. DOI: 10.1117/12.507487. URL: <http://dx.doi.org/10.1117/12.507487>.
- [72] Carlton M. Caves. “Quantum-mechanical noise in an interferometer”. In: *Physical Review D* 23 (8 1981), pp. 1693–1708. DOI: 10.1103/PhysRevD.23.1693. URL: <https://link.aps.org/doi/10.1103/PhysRevD.23.1693>.
- [73] V.B. Braginsky and F.Ja. Khalili. “Gravitational wave antenna with QND speed meter”. In: *Physics Letters A* 147.5 (1990), pp. 251–256. ISSN: 0375-9601. DOI: [https://doi.org/10.1016/0375-9601\(90\)90442-Q](https://doi.org/10.1016/0375-9601(90)90442-Q). URL: <https://www.sciencedirect.com/science/article/pii/037596019090442Q>.
- [74] H. Miao. *Exploring Macroscopic Quantum Mechanics in Optomechanical Devices*. Springer Theses. Springer Berlin Heidelberg, 2012. ISBN: 9783642256400. URL: <https://books.google.co.uk/books?id=Myo8jXpvx2sC>.
- [75] Robert Brown and Richard Taylor. *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 the general existence of active molecules in organic and inorganic bodies*. eng. [London], 1828. URL: <http://haab-digital.klassik-stiftung.de/viewer/epnresolver?id=3523373050>.
- [76] Peter G. Murray. “Measurement of the mechanical loss of test mass materials for advanced gravitational wave detectors”. Available at <https://theses.gla.ac.uk/565/>. PhD thesis. University of Glasgow, 2008.

- [77] B. Jancovici. *Statistical Physics and Thermodynamics [by] B. Jancovici: Problems by Yves Archambault. Translated from the French by L.J. Carroll.* Wiley, 1973. URL: <https://books.google.co.uk/books?id=ewcwngeECAAJ>.
- [78] Gregory M. Harry et al. “Thermal noise from optical coatings in gravitational wave detectors”. In: *Applied Optics* 45.7 (2006), pp. 1569–1574. DOI: 10.1364/AO.45.001569. URL: <https://opg.optica.org/ao/abstract.cfm?URI=ao-45-7-1569>.
- [79] Lukas Terkowski. “Silicon Based Mirror Coatings For Gravitational Wave Detectors”. Available at <https://ediss.sub.uni-hamburg.de/handle/ediss/9852>. PhD thesis. University of Hamburg, 2022.
- [80] G. Billingsley. “*Advanced LIGO Input Test Mass (ITM) A+ version*”. Tech. rep. Pasadena, CA, USA: The LIGO Scientific Collaboration, 2021. URL: <https://dcc.ligo.org/LIGO-E080511/public>.
- [81] Eiichi Hirose et al. “Characterization of Core Optics in Gravitational-Wave Detectors: Case Study of KAGRA Sapphire Mirrors”. In: *Physical Review Applied* 14 (1 2020), p. 014021. DOI: 10.1103/PhysRevApplied.14.014021. URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.14.014021>.
- [82] M Granata et al. “Amorphous optical coatings of present gravitational-wave interferometers*”. In: *Classical and Quantum Gravity* 37.9 (2020), p. 095004. DOI: 10.1088/1361-6382/ab77e9. URL: <https://dx.doi.org/10.1088/1361-6382/ab77e9>.
- [83] R. Abbott et al. “GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo during the Second Part of the Third Observing Run”. In: *Physical Review X* 13.4 (Dec. 2023). ISSN: 2160-3308. DOI: 10.1103/physrevx.13.041039. URL: <http://dx.doi.org/10.1103/PhysRevX.13.041039>.
- [84] L. McCuller et al. “Frequency-Dependent Squeezing for Advanced LIGO”. In: *Physical Review Letters* 124 (17 2020), p. 171102. DOI: 10.1103/PhysRevLett.124.171102. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.124.171102>.
- [85] B. P. Abbott et al. “Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA”. In: *Living Reviews in Relativity* 23.1 (2020), p. 3. ISSN: 1433-8351. DOI: 10.1007/s41114-020-00026-9. URL: <https://doi.org/10.1007/s41114-020-00026-9>.

- [86] David Shoemaker, Alessio Rocchi and Shinji Miyok. “*Observing Scenario timeline graphic, post-O3*”. Tech. rep. The LIGO Scientific Collaboration, 2024. URL: <https://dcc.ligo.org/LIGO-G2002127>.
- [87] T. Akutsu et al. “KAGRA: 2.5 generation interferometric gravitational wave detector”. In: *Nature Astronomy* 3.1 (Jan. 2019), pp. 35–40. ISSN: 2397-3366. DOI: 10.1038/s41550-018-0658-y. URL: <http://dx.doi.org/10.1038/s41550-018-0658-y>.
- [88] The LIGO Scientific Collaboration. “*Advanced LIGO, Advanced Virgo and KAGRA observing run plans*”. Tech. rep. The LIGO Scientific Collaboration, 2019. URL: <https://dcc.ligo.org/public/0161/P1900218/002/SummaryForObservers.pdf>.
- [89] John Miller et al. “Prospects for doubling the range of Advanced LIGO”. In: *Physical Review D* 91.6 (Mar. 2015). ISSN: 1550-2368. DOI: 10.1103/physrevd.91.062005. URL: <http://dx.doi.org/10.1103/PhysRevD.91.062005>.
- [90] Peter Fritschel et al. “*Instrument Science White Paper 2022*”. Tech. rep. The LIGO Scientific Collaboration, 2022. URL: <https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?.submit=Identifier&docid=T2200384&version=>.
- [91] R X Adhikari et al. “*LIGO Voyager Upgrade: Design Concept*”. Tech. rep. The LIGO Scientific Collaboration, 2024. URL: <https://docs.ligo.org/voyager/voyagerwhitepaper/main.pdf>.
- [92] G. D. Hammond et al. “Reducing the suspension thermal noise of advanced gravitational wave detectors”. In: *Classical and Quantum Gravity* 29 (2012). Ed. by Mark Hannam et al., p. 124009. DOI: 10.1088/0264-9381/29/12/124009.
- [93] Peter Fritschel et al. “*Report from the LSC Post-O5 Study Group*”. Tech. rep. The LIGO Scientific Collaboration, 2022. URL: <https://dcc.ligo.org/LIGO-T2200287/public>.
- [94] The LIGO Scientific Collaboration. “*Instrument Science White Paper 2015*”. Tech. rep. The LIGO Scientific Collaboration, 2015. URL: <https://dcc.ligo.org/LIGO-T1500290/public>.
- [95] M Saleem et al. “The science case for LIGO-India”. In: *Classical and Quantum Gravity* 39.2 (Dec. 2021), p. 025004. ISSN: 1361-6382. DOI: 10.1088/1361-6382/ac3b99. URL: <http://dx.doi.org/10.1088/1361-6382/ac3b99>.

- [96] K. Ackley et al. “Neutron Star Extreme Matter Observatory: A kilohertz-band gravitational-wave detector in the global network”. In: *Publications of the Astronomical Society of Australia* 37 (2020). ISSN: 1448-6083. DOI: 10.1017/pasa.2020.39. URL: <http://dx.doi.org/10.1017/pasa.2020.39>.
- [97] Matthew Evans et al. *A Horizon Study for Cosmic Explorer: Science, Observatories, and Community*. 2021. arXiv: 2109.09882 [astro-ph.IM].
- [98] Salvatore Vitale and Matthew Evans. “Parameter estimation for binary black holes with networks of third-generation gravitational-wave detectors”. In: *Physical Review D* 95 (6 2017), p. 064052. DOI: 10.1103/PhysRevD.95.064052. URL: <https://link.aps.org/doi/10.1103/PhysRevD.95.064052>.
- [99] Salvatore Vitale et al. “Measuring the Star Formation Rate with Gravitational Waves from Binary Black Holes”. In: *The Astrophysical Journal Letters* 886.1 (Nov. 2019), p. L1. ISSN: 2041-8213. DOI: 10.3847/2041-8213/ab50c0. URL: <http://dx.doi.org/10.3847/2041-8213/ab50c0>.
- [100] Ken K. Y. Ng et al. “Probing Multiple Populations of Compact Binaries with Third-generation Gravitational-wave Detectors”. In: *The Astrophysical Journal Letters* 913.1 (May 2021), p. L5. ISSN: 2041-8213. DOI: 10.3847/2041-8213/abf8be. URL: <http://dx.doi.org/10.3847/2041-8213/abf8be>.
- [101] Samuel Rowlinson et al. “Feasibility study of beam-expanding telescopes in the interferometer arms for the Einstein Telescope”. In: *Physical Review D* 103.2 (Jan. 2021). ISSN: 2470-0029. DOI: 10.1103/physrevd.103.023004. URL: <http://dx.doi.org/10.1103/PhysRevD.103.023004>.
- [102] Michele Maggiore et al. “Science case for the Einstein telescope”. In: *Journal of Cosmology and Astroparticle Physics* 2020.03 (Mar. 2020), pp. 050–050. ISSN: 1475-7516. DOI: 10.1088/1475-7516/2020/03/050. URL: <http://dx.doi.org/10.1088/1475-7516/2020/03/050>.
- [103] Aniello Grado. “Einstein Telescope, the future generation of ground based gravitational wave detectors”. In: *Journal of Physics: Conference Series* 2429.1 (2023), p. 012041. DOI: 10.1088/1742-6596/2429/1/012041. URL: <https://dx.doi.org/10.1088/1742-6596/2429/1/012041>.

- [104] David Shoemaker. “*Prospects for the future of gravitational-wave observation*”. The LIGO Scientific Collaboration. 1st Apr. 2021. URL: <https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?.submit=Identifier&docid=G2100697&version=> (visited on 04/02/2024).
- [105] Daniel Sigg. “*Gravitational Waves: From Idea to Discovery*”. The LIGO Scientific Collaboration. 4th June 2018. URL: <https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?.submit=Identifier&docid=G1801078&version=> (visited on 04/02/2024).
- [106] B. P. Abbott et al. “Multi-messenger Observations of a Binary Neutron Star Merger*”. In: *The Astrophysical Journal Letters* 848.2 (Oct. 2017), p. L12. ISSN: 2041-8213. DOI: 10.3847/2041-8213/aa91c9. URL: <http://dx.doi.org/10.3847/2041-8213/aa91c9>.
- [107] Pau Amaro-Seoane et al. *Laser Interferometer Space Antenna*. 2017. arXiv: 1702.00786 [astro-ph.IM].
- [108] Henrique Araujo et al. “LISA and LISA Pathfinder, the endeavour to detect low frequency GWs”. In: *Journal of Physics: Conference Series* 66 (Jan. 2007). DOI: 10.1088/1742-6596/66/1/012003.
- [109] The LISA Consortium. “*Laser Interferometer Space Antenna: A Cornerstone Mission for the Observation of Gravitational Waves*”. Tech. rep. Washington, DC, USA: National Aeronautics and Space Administration, 2000. URL: https://lisa.nasa.gov/archive2011/Documentation/sts_1.04.pdf.
- [110] European Space Agency. *Capturing the ripples of spacetime: LISA gets go-ahead*. Accessed on February 1, 2024. 2024. URL: https://www.esa.int/Science_Exploration/Space_Science/Capturing_the_ripples_of_spacetime_LISA_gets_go-ahead.
- [111] Shuichi Sato et al. “DECIGO: The Japanese space gravitational wave antenna”. In: *Journal of Physics: Conference Series* 154.1 (2009), p. 012040. DOI: 10.1088/1742-6596/154/1/012040. URL: <https://dx.doi.org/10.1088/1742-6596/154/1/012040>.
- [112] H.A. MacLeod. *Thin-Film Optical Filters*. Series in Optics and Optoelectronics. CRC Press, 2001. ISBN: 9781420033236. URL: <https://books.google.co.uk/books?id=D0S9hxxzPJq8C>.

- [113] O.S. Heavens. *Thin Film Physics*. Science paperbacks. Methuen, 1970. URL: <https://books.google.co.uk/books?id=heINAQAIAAJ>.
- [114] Matt Evans et al. “Advanced LIGO Arm Length Stabilisation System Design”. Tech. rep. The LIGO Scientific Collaboration, 2018. URL: https://dcc.ligo.org/public/0001/T0900144/007/T0900144_v7_ALS_Design.pdf.
- [115] Jessica Steinlechner et al. “Coatings for Gravitational Wave Detectors”. In: *Optical Interference Coatings Conference (OIC) 2022*. Optica Publishing Group, 2022, WA.4. DOI: 10.1364/OIC.2022.WA.4. URL: <https://opg.optica.org/abstract.cfm?URI=OIC-2022-WA.4>.
- [116] The Virgo Collaboration et al. “The Virgo large mirrors: a challenge for low loss coatings”. In: *Classical and Quantum Gravity* 21.5 (2004), S935. DOI: 10.1088/0264-9381/21/5/083. URL: <https://dx.doi.org/10.1088/0264-9381/21/5/083>.
- [117] LMA. *Laboratoire des Matériaux Avancés*. Université de Lyon, F-69622 Villeurbanne, Lyon, France.
- [118] N.A. Robertson, K.A. Strain and J. Hough. “Measurements of losses in high reflectance mirrors coated for $\lambda = 514.5$ nm”. In: *Optics Communications* 69.5 (1989), pp. 345–348. ISSN: 0030-4018. DOI: [https://doi.org/10.1016/0030-4018\(89\)90016-3](https://doi.org/10.1016/0030-4018(89)90016-3). URL: <https://www.sciencedirect.com/science/article/pii/0030401889900163>.
- [119] O Kienzle, J Staub and T Tschudi. “Description of an integrated scatter instrument for measuring scatter losses of ‘superpolished’ optical surfaces”. In: *Measurement Science and Technology* 5.6 (1994), p. 747. DOI: 10.1088/0957-0233/5/6/017. URL: <https://dx.doi.org/10.1088/0957-0233/5/6/017>.
- [120] K. Seshan. *Handbook of Thin-film Deposition Processes and Techniques: Principles, Methods, Equipment and Applications*. Materials science and process technology series: Electronic materials and process technology. Noyes Publications/William Andrew Pub., 2002. ISBN: 9786612253195. URL: <https://books.google.co.uk/books?id=72BEzQEACAAJ>.
- [121] K. Arai P. Sekhar G. Venugopalan. “Optical Loss Characterization”. Tech. rep. The LIGO Scientific Collaboration, 2018. URL: <https://dcc.ligo.org/LIGO-T1800245/public>.

- [122] R Flaminio et al. “A study of coating mechanical and optical losses in view of reducing mirror thermal noise in gravitational wave detectors”. In: *Classical and Quantum Gravity* 27.8 (2010), p. 084030. DOI: 10.1088/0264-9381/27/8/084030. URL: <https://dx.doi.org/10.1088/0264-9381/27/8/084030>.
- [123] G Billingsley, Hiroaki Yamamoto and Liyuan Zhang. “Characterization of Advanced LIGO core optics”. In: vol. 66. American Society for Precision Engineering (ASPE), 2017, pp. 78–83. ISBN: 9781510853553. URL: <https://www.proceedings.com/37625.html>.
- [124] L. Pinard et al. “Mirrors used in the LIGO interferometers for first detection of gravitational waves”. In: *Applied Optics* 56.4 (2017), pp. C11–C15. DOI: 10.1364/AO.56.000C11. URL: <https://opg.optica.org/ao/abstract.cfm?URI=ao-56-4-C11>.
- [125] F Acernese et al. “Advanced Virgo Plus: Future Perspectives”. In: *Journal of Physics: Conference Series* 2429.1 (2023), p. 012040. DOI: 10.1088/1742-6596/2429/1/012040. URL: <https://dx.doi.org/10.1088/1742-6596/2429/1/012040>.
- [126] M. Gouy. “Note sur le mouvement brownien”. In: *Journal of Physics: Theories and Applications* 7.1 (1888), pp. 561–566. DOI: 10.1051/jphystap:018880070056101. eprint: <https://hal.science/jpa-00238904>. URL: <https://hal.science/jpa-00238904>.
- [127] F. M. Exner. “Notiz. zu. Brown’s molecularbewegung”. In: *Annalen der Physik* 2 843 (1900).
- [128] A. Einstein. “On the Movement of Small Particles Suspended in Stationary Liquids Required by the Molecular-Kinetic Theory of Heat”. In: *Annalen der Physik* 17 (1905), pp. 549–560. URL: <https://spaces-cdn.owlstown.com/blobs/o0qsihzbhzhv1kubrbimzndwvghh1>.
- [129] Herbert B. Callen and Theodore A. Welton. “Irreversibility and Generalized Noise”. In: *Physical Review* 83 (1 1951), pp. 34–40. DOI: 10.1103/PhysRev.83.34. URL: <https://link.aps.org/doi/10.1103/PhysRev.83.34>.
- [130] Richard F. Greene and Herbert B. Callen. “On the Formalism of Thermodynamic Fluctuation Theory”. In: *Physical Review* 83 (6 1951), pp. 1231–1235. DOI: 10.1103/PhysRev.83.1231. URL: <https://link.aps.org/doi/10.1103/PhysRev.83.1231>.

- [131] Herbert B. Callen and Richard F. Greene. “On a Theorem of Irreversible Thermodynamics”. In: *Physical Review* 86 (5 1952), pp. 702–710. DOI: 10.1103/PhysRev.86.702. URL: <https://link.aps.org/doi/10.1103/PhysRev.86.702>.
- [132] A.S. Nowick. *Anelastic Relaxation In Crystalline Solids*. Materials science and technology. Elsevier Science, 2012. ISBN: 9780323143318. URL: <https://books.google.co.uk/books?id=w02eVGpR0nQC>.
- [133] K.G. McConnell. *Vibration Testing: Theory and Practice*. A Wiley-Interscience publication. Wiley, 1995. ISBN: 9780471304357. URL: <https://books.google.co.uk/books?id=tv8MiE2By74C>.
- [134] Peter R. Saulson. “Thermal noise in mechanical experiments”. In: *Physical Review D* 42 (8 1990), pp. 2437–2445. DOI: 10.1103/PhysRevD.42.2437. URL: <https://link.aps.org/doi/10.1103/PhysRevD.42.2437>.
- [135] Yu. Levin. “Internal thermal noise in the LIGO test masses: A direct approach”. In: *Physical Review D* 57 (2 1998), pp. 659–663. DOI: 10.1103/PhysRevD.57.659. URL: <https://link.aps.org/doi/10.1103/PhysRevD.57.659>.
- [136] François Bondu, Patrice Hello and Jean-Yves Vinet. “Thermal noise in mirrors of interferometric gravitational wave antennas”. In: *Physics Letters A* 246.3 (1998), pp. 227–236. ISSN: 0375-9601. DOI: [https://doi.org/10.1016/S0375-9601\(98\)00450-2](https://doi.org/10.1016/S0375-9601(98)00450-2). URL: <https://www.sciencedirect.com/science/article/pii/S0375960198004502>.
- [137] Yuk Tung Liu and Kip S. Thorne. “Thermoelastic noise and homogeneous thermal noise in finite sized gravitational-wave test masses”. In: *Physical Review D* 62 (12 2000), p. 122002. DOI: 10.1103/PhysRevD.62.122002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.62.122002>.
- [138] Yukino Mori et al. “Coating thermal noise investigation for KAGRA”. In: *PoS ICRC2023* (2023), p. 1577. DOI: 10.22323/1.444.1577.
- [139] N. Nakagawa et al. “Thermal noise in half-infinite mirrors with nonuniform loss: A slab of excess loss in a half-infinite mirror”. In: *Physical Review D* 65 (10 2002), p. 102001. DOI: 10.1103/PhysRevD.65.102001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.65.102001>.

- [140] Gregory M Harry et al. “Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings”. In: *Classical and Quantum Gravity* 19.5 (2002), p. 897. DOI: 10.1088/0264-9381/19/5/305. URL: <https://dx.doi.org/10.1088/0264-9381/19/5/305>.
- [141] Ting Hong et al. “Brownian thermal noise in multilayer coated mirrors”. In: *Physical Review D* 87 (8 2013), p. 082001. DOI: 10.1103/PhysRevD.87.082001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.87.082001>.
- [142] John D. Ferry and Henry S. Myers. “Viscoelastic Properties of Polymers”. In: *Journal of The Electrochemical Society* 108.7 (1961), p. 142C. DOI: 10.1149/1.2428174. URL: <https://dx.doi.org/10.1149/1.2428174>.
- [143] William Yam, Slawek Gras and Matthew Evans. “Multimaterial coatings with reduced thermal noise”. In: *Physical Review D* 91 (4 2015), p. 042002. DOI: 10.1103/PhysRevD.91.042002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.91.042002>.
- [144] Jessica Steinlechner and Iain W. Martin. “How can amorphous silicon improve current gravitational-wave detectors?” In: *Physical Review D* 103 (4 2021), p. 042001. DOI: 10.1103/PhysRevD.103.042001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.103.042001>.
- [145] Martin M. Fejer. “*Effective Medium Description of Multilayer Coatings*”. Tech. rep. Pasadena, CA, USA: The LIGO Scientific Collaboration, 2021. URL: <https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?.submit=Identifier&docid=T2100186&version=>.
- [146] George E. Backus. “Long-wave elastic anisotropy produced by horizontal layering”. In: *Journal of Geophysical Research (1896-1977)* 67.11 (1962), pp. 4427–4440. DOI: <https://doi.org/10.1029/JZ067i011p04427>. eprint: <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/JZ067i011p04427>. URL: <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/JZ067i011p04427>.
- [147] Gabriele Vajente et al. *Low Mechanical Loss TiO₂:GeO₂ Coatings for Reduced Thermal Noise in Gravitational Wave Interferometers Supplemental Material*. American Physical Society. 10th Aug. 2021. URL: <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.127.071101> (visited on 11/01/2023).

- [148] Gabriele Vajente et al. “Low Mechanical Loss $\text{TiO}_2 : \text{GeO}_2$ Coatings for Reduced Thermal Noise in Gravitational Wave Interferometers”. In: *Physical Review Letters* 127 (7 2021), p. 071101. DOI: 10.1103/PhysRevLett.127.071101. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.127.071101>.
- [149] Massimo Granata et al. “Mechanical loss in state-of-the-art amorphous optical coatings”. In: *Physical Review D* 93 (1 2016), p. 012007. DOI: 10.1103/PhysRevD.93.012007. URL: <https://link.aps.org/doi/10.1103/PhysRevD.93.012007>.
- [150] K. Prasai et al. “High Precision Detection of Change in Intermediate Range Order of Amorphous Zirconia-Doped Tantalum Thin Films Due to Annealing”. In: *Physical Review Letters* 123 (4 2019), p. 045501. DOI: 10.1103/PhysRevLett.123.045501. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.045501>.
- [151] Le Yang et al. “Enhanced medium-range order in vapor-deposited germania glasses at elevated temperatures”. In: *Science Advances* 7.37 (Sept. 2021). ISSN: 2375-2548. DOI: 10.1126/sciadv.abh1117. URL: <http://dx.doi.org/10.1126/sciadv.abh1117>.
- [152] Gregory M Harry et al. “Titania-doped tantalum/silica coatings for gravitational-wave detection”. In: *Classical and Quantum Gravity* 24.2 (2006), p. 405. DOI: 10.1088/0264-9381/24/2/008. URL: <https://dx.doi.org/10.1088/0264-9381/24/2/008>.
- [153] Mikhail N. Polyanskiy. “Refractiveindex.info database of optical constants”. In: *Scientific Data* 11.1 (2024), p. 94. ISSN: 2052-4463. DOI: 10.1038/s41597-023-02898-2. URL: <https://doi.org/10.1038/s41597-023-02898-2>.
- [154] Carmen S. Menoni et al. “Survey of metal oxides for coatings of ultra-stable optical cavities”. In: *Conference on Lasers and Electro-Optics*. Optica Publishing Group, 2021, STu1C.7. DOI: 10.1364/CLEO_SI.2021.STu1C.7. URL: https://opg.optica.org/abstract.cfm?URI=CLEO_SI-2021-STu1C.7.
- [155] Alexandr Ageev et al. “Very high quality factor measured in annealed fused silica”. In: *Classical and Quantum Gravity* 21.16 (2004), p. 3887. DOI: 10.1088/0264-9381/21/16/004. URL: <https://dx.doi.org/10.1088/0264-9381/21/16/004>.

- [156] Steven D. Penn et al. “Frequency and surface dependence of the mechanical loss in fused silica”. In: *Physics Letters A* 352.1 (2006), pp. 3–6. ISSN: 0375-9601. DOI: <https://doi.org/10.1016/j.physleta.2005.11.046>. URL: <https://www.sciencedirect.com/science/article/pii/S0375960105017846>.
- [157] G. Vajente. private communication. 2021.
- [158] Shiuh Chao, Wen-Hsiang Wang and Cheng-Chung Lee. “Low-loss dielectric mirror with ion-beam-sputtered TiO₂–SiO₂ mixed films”. In: *Applied Optics* 40.13 (2001), pp. 2177–2182. DOI: 10.1364/AO.40.002177. URL: <https://opg.optica.org/ao/abstract.cfm?URI=ao-40-13-2177>.
- [159] Jessica Steinlechner et al. “Optical absorption of ion-beam sputtered amorphous silicon coatings”. In: *Physical Review D* 93 (6 2016), p. 062005. DOI: 10.1103/PhysRevD.93.062005. URL: <https://link.aps.org/doi/10.1103/PhysRevD.93.062005>.
- [160] J. E. Fredrickson et al. “Effects of thermal annealing on the refractive index of amorphous silicon produced by ion implantation”. In: *Applied Physics Letters* 40.2 (Jan. 1982), pp. 172–174. ISSN: 0003-6951. DOI: 10.1063/1.93032. eprint: https://pubs.aip.org/aip/apl/article-pdf/40/2/172/18445347/172_1_online.pdf. URL: <https://doi.org/10.1063/1.93032>.
- [161] R. Birney et al. “Amorphous Silicon with Extremely Low Absorption: Beating Thermal Noise in Gravitational Astronomy”. In: *Physical Review Letters* 121 (19 2018), p. 191101. DOI: 10.1103/PhysRevLett.121.191101. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.121.191101>.
- [162] Carl Lévesque et al. *Internal mechanical dissipation mechanisms in amorphous silicon*. 2022. arXiv: 2209.02342 [cond-mat.dis-nn].
- [163] M. Molina-Ruiz et al. “Hydrogen-Induced Ultralow Optical Absorption and Mechanical Loss in Amorphous Silicon for Gravitational-Wave Detectors”. In: *Physical Review Letters* 131 (25 2023), p. 256902. DOI: 10.1103/PhysRevLett.131.256902. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.131.256902>.
- [164] Simon C. Tait et al. “Demonstration of the Multimaterial Coating Concept to Reduce Thermal Noise in Gravitational-Wave Detectors”. In: *Physical Review Letters* 125 (1 2020), p. 011102. DOI: 10.1103/PhysRevLett.125.011102. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.125.011102>.

- [165] M. Granata et al. “Progress in the measurement and reduction of thermal noise in optical coatings for gravitational-wave detectors”. In: *Applied Optics* 59.5 (Feb. 2020), A229. ISSN: 2155-3165. DOI: 10.1364/ao.377293. URL: <http://dx.doi.org/10.1364/AO.377293>.
- [166] Huang-Wei Pan et al. “Silicon nitride and silica quarter-wave stacks for low-thermal-noise mirror coatings”. In: *Physical Review D* 98 (10 2018), p. 102001. DOI: 10.1103/PhysRevD.98.102001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.98.102001>.
- [167] Dung-Sheng Tsai et al. “Amorphous silicon nitride deposited by an NH₃-free plasma enhanced chemical vapor deposition method for the coatings of the next generation laser interferometer gravitational waves detector”. In: *Classical and Quantum Gravity* 39.15 (2022), 15LT01. DOI: 10.1088/1361-6382/ac79f6. URL: <https://dx.doi.org/10.1088/1361-6382/ac79f6>.
- [168] V. Pierro et al. “Ternary quarter wavelength coatings for gravitational wave detector mirrors: Design optimization via exhaustive search”. In: *Physical Review Res.* 3 (2 2021), p. 023172. DOI: 10.1103/PhysRevResearch.3.023172. URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.3.023172>.
- [169] G. D. Cole et al. “Substrate-transferred GaAs/AlGaAs crystalline coatings for gravitational-wave detectors”. In: *Applied Physics Letters* 122.11 (Mar. 2023). ISSN: 1077-3118. DOI: 10.1063/5.0140663. URL: <http://dx.doi.org/10.1063/5.0140663>.
- [170] Dhruv Kedar et al. *Frequency stability of cryogenic silicon cavities with semiconductor crystalline coatings*. 2022. arXiv: 2210.14881 [physics.optics].
- [171] Graeme McGhee et al. “Titania Mixed with Silica: A Low Thermal-Noise Coating Material for Gravitational-Wave Detectors”. In: *Physical Review Letters* 131 (Oct. 2023). DOI: 10.1103/PhysRevLett.131.171401.
- [172] Steven D Penn et al. “Mechanical loss in tantala/silica dielectric mirror coatings”. In: *Classical and Quantum Gravity* 20.13 (2003), p. 2917. DOI: 10.1088/0264-9381/20/13/334. URL: <https://dx.doi.org/10.1088/0264-9381/20/13/334>.
- [173] *FiveNine Optics*. <https://www.fivenineoptics.com/>. Accessed on: 2024-02-21. 2016.

- [174] Wei-Kan Chu, James W. Mayer and Marc-A. Nicolet. *Backscattering Spectrometry*. New York: Academic Press, 1978. ISBN: 978-0-12-173850-1. DOI: <https://doi.org/10.1016/B978-0-12-173850-1.50002-8>. URL: <https://www.sciencedirect.com/science/article/pii/B9780121738501500028>.
- [175] Agilent Technologies. “*Cary 100/300/4000/5000/6000i/7000 Spectrophotometers User’s Guide*”. Agilent Technologies. 2021. URL: https://www.agilent.com/cs/library/usermanuals/public/1972_7000.pdf.
- [176] *WTheiss Hardware and Software*. URL: <http://www.wtheiss.com/>.
- [177] Juan Alcalde Cartel. *Accelerator Mass Spectrometry - AMS*. Presented at the OPAC 2014 conference. 2014. URL: https://indico.cern.ch/event/297045/sessions/146678/attachments/557272/767791/Juan_Alcalde_CARTEL_OPAC_AMS_GROUP.pdf.
- [178] M. Mayer. “SIMNRA, a simulation program for the analysis of NRA, RBS and ERDA”. In: *AIP Conference Proceedings* 475.1 (1999), pp. 541–544. DOI: 10.1063/1.59188. eprint: <https://aip.scitation.org/doi/pdf/10.1063/1.59188>. URL: <https://aip.scitation.org/doi/abs/10.1063/1.59188>.
- [179] R Birney et al. “Exploring the fundamental limits of NIR absorption in amorphous silicon”. In: *Physical Review Letters* 1 (), pp. 1–6. URL: <https://dcc.ligo.org/LIGO-P1800148>.
- [180] R. Groleau, S.C. Gujrathi and J.P. Martin. “Time-of-flight system for profiling recoiled light elements”. In: *Nuclear Instruments and Methods in Physics Research* 218.1 (1983), pp. 11–15. ISSN: 0167-5087. DOI: [https://doi.org/10.1016/0167-5087\(83\)90946-8](https://doi.org/10.1016/0167-5087(83)90946-8). URL: <https://www.sciencedirect.com/science/article/pii/0167508783909468>.
- [181] M. Chicoine et al. “Oxy-nitrides characterization with a new ERD-TOF system”. In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 406 (Sept. 2017), pp. 112–114. DOI: 10.1016/j.nimb.2017.02.041. URL: <https://doi.org/10.1016/j.nimb.2017.02.041>.
- [182] Simon C. Tait. “Studies of mechanical and optical properties of thin film coatings for future gravitational wave detectors”. Available at <https://theses.gla.ac.uk/82514/>. PhD thesis. University of Glasgow, 2021.

- [183] Maya M. A. Kinley-Hanlon. “Characterization of new materials and designs for reflective optical coatings of future gravitational wave detectors”. Available at <https://theses.gla.ac.uk/83760/>. PhD thesis. University of Glasgow, 2023.
- [184] E. Cesarini et al. “A “gentle” nodal suspension for measurements of the acoustic attenuation in materials”. In: *Review of Scientific Instruments* 80.5 (2009), p. 053904. DOI: 10.1063/1.3124800. eprint: <https://doi.org/10.1063/1.3124800>. URL: <https://doi.org/10.1063/1.3124800>.
- [185] Alex Amato. “Low Thermal Noise Coating for New Generation Gravitational-Wave Detectors”. PhD thesis. Lyon U., 2019.
- [186] SIOS GmbH. “*Precision in Measurement*”. Tech. rep. SIOS, 2020. URL: <https://sios.de/>.
- [187] Sharon M. Twyford. “Developments Towards Low Loss Suspensions for Laser Interferometric Gravitational Wave Detectors”. Available at <https://theses.gla.ac.uk/75385/>. PhD thesis. University of Glasgow, 1998.
- [188] G.E. Roberts. *From Music to Mathematics: Exploring the Connections*. Johns Hopkins University Press, 2016, p. 112. ISBN: 9781421419190. URL: <https://books.google.co.uk/books?id=tv6TCwAAQBAJ>.
- [189] Richard Phillips Feynman, Robert Benjamin Leighton and Matthew Sands. *The Feynman lectures on physics; New millennium ed.* Originally published 1963-1965. New York, NY: Basic Books, 2010. Chap. 48: Beats. URL: <https://cds.cern.ch/record/1494701>.
- [190] Diana Lumaca et al. “Stability of samples in coating research: From edge effect to ageing”. In: *Journal of Alloys and Compounds* 930 (2023), p. 167320. ISSN: 0925-8388. DOI: <https://doi.org/10.1016/j.jallcom.2022.167320>. URL: <https://www.sciencedirect.com/science/article/pii/S0925838822037112>.
- [191] COMSOL. “*COMSOL Multiphysics User’s Guide*”. Tech. rep. COMSOL, 2012. URL: <https://doc.comsol.com/5.5/doc/com.comsol.help.comsol/COMSOLReferenceManual.pdf>.

- [192] M. M. Fejer et al. “Thermoelastic dissipation in inhomogeneous media: loss measurements and displacement noise in coated test masses for interferometric gravitational wave detectors”. In: *Phys. Rev. D* 70 (8 2004), p. 082003. DOI: 10.1103/PhysRevD.70.082003. URL: <https://link.aps.org/doi/10.1103/PhysRevD.70.082003>.
- [193] D.L. Logan. *A First Course in the Finite Element Method*. Thomson, 2007. ISBN: 9780534552985. URL: https://kntu.ac.ir/DorsaPax/userfiles/file/Mechanical/OstadFile/dr_nakhodchi/DarylL.LoganAFirstCourse.pdf.
- [194] Štefan Barta. “Effective Young’s modulus and Poisson’s ratio for the particulate composite”. In: *Journal of Applied Physics* 75.7 (Apr. 1994), pp. 3258–3263. ISSN: 0021-8979. DOI: 10.1063/1.356132. eprint: https://pubs.aip.org/aip/jap/article-pdf/75/7/3258/7445878/3258_1_online.pdf. URL: <https://doi.org/10.1063/1.356132>.
- [195] R. M. Christensen and K. H. Lo. “Solutions for effective shear properties in three phase sphere and cylinder models”. In: *Journal of Mechanics Physics of Solids* 27.4 (Aug. 1979), pp. 315–330. DOI: 10.1016/0022-5096(79)90032-2.
- [196] J. D. Eshelby. “The Determination of the Elastic Field of an Ellipsoidal Inclusion, and Related Problems”. In: *Proceedings of the Royal Society of London Series A* 241.1226 (Aug. 1957), pp. 376–396. DOI: 10.1098/rspa.1957.0133.
- [197] E. H. Kerner. “The Elastic and Thermo-elastic Properties of Composite Media”. In: *Proceedings of the Physical Society B* 69.8 (Aug. 1956), pp. 808–813. DOI: 10.1088/0370-1301/69/8/305.
- [198] The MathWorks Inc. *Symbolic Math Toolbox*. Natick, Massachusetts, United States, 2023. URL: <https://www.mathworks.com>.
- [199] Tianjun Li et al. “Measurements of mechanical thermal noise and energy dissipation in optical dielectric coatings”. In: *Physical Review D* 89 (9 2014), p. 092004. DOI: 10.1103/PhysRevD.89.092004. URL: <https://link.aps.org/doi/10.1103/PhysRevD.89.092004>.
- [200] Stuart Reid and Iain W. Martin. “Development of Mirror Coatings for Gravitational Wave Detectors”. In: *Coatings* 6.4 (2016). ISSN: 2079-6412. DOI: 10.3390/coatings6040061. URL: <https://www.mdpi.com/2079-6412/6/4/61>.

- [201] Kenji Numata et al. “Systematic measurement of the intrinsic losses in various kinds of bulk fused silica”. In: *Physics Letters A* 327.4 (2004), pp. 263–271. ISSN: 0375-9601. DOI: <https://doi.org/10.1016/j.physleta.2004.05.028>. URL: <https://www.sciencedirect.com/science/article/pii/S0375960104007005>.
- [202] K. S. Gilroy and W. A. Phillips. “An asymmetric double-well potential model for structural relaxation processes in amorphous materials”. In: *Philosophical Magazine B* 43.5 (1981), pp. 735–746. DOI: 10.1080/01418638108222343. eprint: <https://doi.org/10.1080/01418638108222343>. URL: <https://doi.org/10.1080/01418638108222343>.
- [203] I W Martin et al. “Effect of heat treatment on mechanical dissipation in Ta₂O₅ coatings”. In: *Classical and Quantum Gravity* 27.22 (2010), p. 225020. DOI: 10.1088/0264-9381/27/22/225020. URL: <https://dx.doi.org/10.1088/0264-9381/27/22/225020>.
- [204] Roger P. Netterfield et al. “Low mechanical loss coatings for LIGO optics: progress report”. In: *Advances in Thin-Film Coatings for Optical Applications II*. Ed. by Michael L. Fulton and Jennifer D. T. Kruschwitz. Vol. 5870. International Society for Optics and Photonics. SPIE, 2005, 58700H. DOI: 10.1117/12.618288. URL: <https://doi.org/10.1117/12.618288>.
- [205] M.-W. Moon et al. “An experimental study of the influence of imperfections on the buckling of compressed thin films”. In: *Acta Materialia* 50.5 (2002), pp. 1219–1227. ISSN: 1359-6454. DOI: [https://doi.org/10.1016/S1359-6454\(01\)00423-2](https://doi.org/10.1016/S1359-6454(01)00423-2). URL: <https://www.sciencedirect.com/science/article/pii/S1359645401004232>.
- [206] Grégory Abadias et al. “Review Article: Stress in thin films and coatings: Current status, challenges, and prospects”. In: *Journal of Vacuum Science Technology A* 36.2 (Mar. 2018), p. 020801. ISSN: 0734-2101. DOI: 10.1116/1.5011790. eprint: https://pubs.aip.org/avs/jva/article-pdf/doi/10.1116/1.5011790/15931577/020801_1_online.pdf. URL: <https://doi.org/10.1116/1.5011790>.
- [207] MPO Imaging Systems. “*Optical Coating Stress*”. Tech. rep. Isle-of-Man, UK: Manx Precision Optics Ltd., 2020. URL: <https://mpo.im/wp-content/uploads/2020/06/TN-2020-02-Optical-Stress-1.pdf>.

- [208] Émile Lalande. “Ar transport and blister growth kinetics in titania doped germania-based optical coatings”. Manuscript in preparation. URL: <https://dcc.ligo.org/DocDB/0190/P2300328/002/ArDesorptionV4.pdf>.
- [209] S. Gras and M. Evans. “Direct measurement of coating thermal noise in optical resonators”. In: *Physical Review D* 98 (12 2018), p. 122001. DOI: 10.1103/PhysRevD.98.122001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.98.122001>.
- [210] S. Gras et al. “Audio-band coating thermal noise measurement for Advanced LIGO with a multimode optical resonator”. In: *Physical Review D* 95 (2 2017), p. 022001. DOI: 10.1103/PhysRevD.95.022001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.95.022001>.
- [211] S. Gras. private communication. 2022.
- [212] A. Amato et al. “Optical and mechanical properties of ion-beam-sputtered Nb_2O_5 and $\text{TiO}_2\text{--Nb}_2\text{O}_5$ thin films for gravitational-wave interferometers and an improved measurement of coating thermal noise in Advanced LIGO”. In: *Physical Review D* 103 (7 2021), p. 072001. DOI: 10.1103/PhysRevD.103.072001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.103.072001>.
- [213] Raymond Robie. “Characterisation of the mechanical properties of thin-film mirror coating materials for use in future interferometric gravitational wave detectors”. PhD thesis. University of Glasgow, 2018. URL: <http://theses.gla.ac.uk/id/eprint/30645>.
- [214] David R M Crooks. “Mechanical Loss and its Significance in the Test Mass Mirrors of Gravitational Wave Detectors”. PhD thesis. University of Glasgow, 2002. eprint: <http://theses.gla.ac.uk/id/eprint/2893>. URL: <http://theses.gla.ac.uk/id/eprint/2893>.
- [215] A. Alexandrovski et al. “Photothermal common-path interferometry (PCI): new developments”. In: *Solid State Lasers XVIII: Technology and Devices*. Ed. by W. Andrew Clarkson, Norman Hodgson and Ramesh K. Shori. Vol. 7193. International Society for Optics and Photonics. SPIE, 2009, p. 71930D. DOI: 10.1117/12.814813. URL: <https://doi.org/10.1117/12.814813>.

- [216] M. Leidinger et al. “Comparative study on three highly sensitive absorption measurement techniques characterizing lithium niobate over its entire transparent spectral range”. In: *Opt. Express* 23.17 (2015), pp. 21690–21705. DOI: 10.1364/OE.23.021690. URL: <https://opg.optica.org/oe/abstract.cfm?URI=oe-23-17-21690>.
- [217] Iain W. Martin et al. “Time-evolution of NIR absorption in hydroxide-catalysis bonds”. In: *Materialia* 6 (2019), p. 100331. ISSN: 2589-1529. DOI: <https://doi.org/10.1016/j.mtla.2019.100331>. URL: <https://www.sciencedirect.com/science/article/pii/S2589152919301279>.
- [218] Corning. “*Corning® HPFS® 7979, 7980, 8655 Fused Silica Optical Materials Product Information*”. Tech. rep. Corning, NY, USA: Corning, Inc., 2015. URL: https://www.corning.com/media/worldwide/csm/documents/HPFS_Product_Brochure_All_Grades_2015_07_21.pdf.
- [219] Jessica Steinlechner et al. “Thermal noise reduction and absorption optimization via multimaterial coatings”. In: *Physical Review D* 91.4 (2015). DOI: 10.1103/PhysRevD.91.042001. URL: <https://doi.org/10.1103/PhysRevD.91.042001>.
- [220] Mariana A. Fazio et al. “Structure and morphology of low mechanical loss TiO₂-doped Ta₂O₅”. In: *Optical Materials Express* 10.7 (2020), pp. 1687–1703. DOI: 10.1364/OME.395503. URL: <https://opg.optica.org/ome/abstract.cfm?URI=ome-10-7-1687>.
- [221] Lin Gu et al. “Temperature-dependent oxygen annealing effect on the properties of Ga₂O₃ thin film deposited by atomic layer deposition”. In: *Journal of Alloys and Compounds* 925 (2022), p. 166727. ISSN: 0925-8388. DOI: <https://doi.org/10.1016/j.jallcom.2022.166727>. URL: <https://www.sciencedirect.com/science/article/pii/S0925838822031188>.
- [222] Namjun Kim and Jonathan F. Stebbins. “Effects of annealing on the structure of ion beam sputtered amorphous tantalum oxide: Oxygen-17 NMR spectra and relaxation times”. In: *Journal of Non-Crystalline Solids* 378 (2013), pp. 158–162. ISSN: 0022-3093. DOI: <https://doi.org/10.1016/j.jnoncrysol.2013.07.003>. URL: <https://www.sciencedirect.com/science/article/pii/S0022309313003815>.

- [223] R. Bassiri et al. “Correlations between the mechanical loss and atomic structure of amorphous TiO₂-doped Ta₂O₅ coatings”. In: *Acta Materialia* 61.4 (2013), pp. 1070–1077. ISSN: 1359-6454. DOI: <https://doi.org/10.1016/j.actamat.2012.10.009>. URL: <https://www.sciencedirect.com/science/article/pii/S1359645412007343>.
- [224] C. Michel. private communication. 2023.
- [225] S Soni et al. “Reducing scattered light in LIGO’s third observing run”. In: *Classical and Quantum Gravity* 38.2 (2020), p. 025016. DOI: 10.1088/1361-6382/abc906. URL: <https://dx.doi.org/10.1088/1361-6382/abc906>.
- [226] P. Kwee et al. “Decoherence and degradation of squeezed states in quantum filter cavities”. In: *Physical Review D* 90 (6 2014), p. 062006. DOI: 10.1103/PhysRevD.90.062006. URL: <https://link.aps.org/doi/10.1103/PhysRevD.90.062006>.
- [227] K. S. Thorne E. Flanagan. “Noise due to backscatter off baffles, the nearby wall, and objects at the far end of the beam tube; and recommended actions”. Tech. rep. Pasadena, CA, USA: California Institute of Technology, 1994. URL: <https://dcc.ligo.org/public/0028/T940063/000/T940063-00.pdf>.
- [228] T Accadia et al. “Noise from scattered light in Virgo’s second science run data”. In: *Classical and Quantum Gravity* 27.19 (2010), p. 194011. DOI: 10.1088/0264-9381/27/19/194011. URL: <https://dx.doi.org/10.1088/0264-9381/27/19/194011>.
- [229] David Halliday, Robert Resnick and Jearl Walker. *Fundamentals of Physics*. 10th. Wiley, 2013.
- [230] P. Bousquet, F. Flory and P. Roche. “Scattering from multilayer thin films: theory and experiment”. In: *J. Optics Soc. Am.* 71.9 (1981), pp. 1115–1123. DOI: 10.1364/JOSA.71.001115. URL: <https://opg.optica.org/abstract.cfm?URI=josa-71-9-1115>.
- [231] J. M. Elson, J. P. Rahn and J. M. Bennett. “Light scattering from multilayer optics: comparison of theory and experiment”. In: *Applied Optics* 19.5 (1980), pp. 669–679. DOI: 10.1364/AO.19.000669. URL: <https://opg.optica.org/ao/abstract.cfm?URI=ao-19-5-669>.

- [232] Labsphere. “*Integrating Sphere Theory and Applications*”. Tech. rep. North Sutton, NH, USA: Labsphere, Inc., 2017. URL: <https://www.labsphere.com/wp-content/uploads/2021/09/Integrating-Sphere-Theory-and-Applications.pdf>.
- [233] J. L. Taylor. “*Integrating Sphere Functionality: The Scatter Transmission Measurement*”. Tech. rep. Shelton, CT, USA: PerkinElmer, Inc., 2013. URL: https://resources.perkinelmer.com/corporate/cmsresources/images/44-156124tch_011486_01_integratingspherefunctionalitythescattertransmissionmeasurement.pdf.
- [234] Shigeo Kubota, T. Matsumoto and T. Shimura. “An integrating sphere system to realize very-low-luminance reference light sources”. In: Dec. 2008. DOI: 10.13140/2.1.4398.1129.
- [235] K. Gindele, M. Köhl and M. Mast. “Spectral reflectance measurements using an integrating sphere in the infrared”. In: *Applied Optics* 24.12 (1985), pp. 1757–1760. DOI: 10.1364/AO.24.001757. URL: <https://opg.optica.org/ao/abstract.cfm?URI=ao-24-12-1757>.
- [236] Labsphere. “*Spectralon® Diffuse Reflectance Material Technical Guide*”. Tech. rep. North Sutton, NH, USA: Labsphere, Inc., 2017. URL: <https://www.labsphere.com/wp-content/uploads/2023/01/Spectralon-Diffuse-Reflectance-Material-Tech-Guide.pdf>.
- [237] Labsphere. “*Spectralon® Diffuse Reflectance Material Data Sheet*”. Tech. rep. North Sutton, NH, USA: Labsphere, Inc., 2017. URL: <https://www.labsphere.com/wp-content/uploads/2023/01/Spectralon.pdf>.
- [238] Carol Bruegge, Nadine Chrien and David Haner. “A Spectralon BRF data base for MISR calibration applications”. In: *Remote Sensing of Environment* 77.3 (2001), pp. 354–366. ISSN: 0034-4257. DOI: [https://doi.org/10.1016/S0034-4257\(01\)00214-0](https://doi.org/10.1016/S0034-4257(01)00214-0). URL: <https://www.sciencedirect.com/science/article/pii/S0034425701002140>.
- [239] J.C. Stover. “*Optical Scattering: Measurement and Analysis*”. SPIE Press, 2012. ISBN: 9780819492524. URL: <https://www.spiedigitallibrary.org/ebooks/PM/Optical-Scattering-Measurement-and-Analysis-Third-Edition/eISBN-9780819492524/10.1117/3.975276>.

- [240] Photonic Cleaning Technologies. “*MIRRORS, METALS AND FIRST CONTACT®*”. Tech. rep. Platteville, WI, USA: Photonic Cleaning Technologies, LLC, 2010. URL: <http://www.photoniccleaning.com/v/vspfiles/downloadables/Documentation/Mirrors%20Datasheet%20PCT%203.08.pdf>.
- [241] Photonic Cleaning Technologies. “*FIRST CONTACT® Polymer Brochure: How Clean?*”. Tech. rep. Platteville, WI, USA: Photonic Cleaning Technologies, LLC, 2019. URL: <http://www.photoniccleaning.com/v/vspfiles/downloadables/Documentation/Brochure%20Final%202019.pdf>.
- [242] Photonic Cleaning Technologies. “*FIRST CONTACT® Polymer: Safety Data Sheet*”. Tech. rep. Platteville, WI, USA: Photonic Cleaning Technologies, LLC, 2019. URL: http://www.photoniccleaning.com/v/vspfiles/downloadables/MSDS/PCT%20First%20Contact%20Polymer%20Solutions%20SDS%20All%20Sizes%20All%20Colors%20ver%201.19%20SDS_EU-USA%205-12-2019.pdf.
- [243] “*SIMCO ION® Top Gun*”. <https://www.simco-ion.co.uk/products/static-eliminators/ionising-guns/top-gun/>. Accessed: 2023-10-05.
- [244] SIMCO. “*User Manual: Top Gun III*”. Tech. rep. Lochem, Netherlands: SIMCO (Nederland) B.V., 2023. URL: https://www.simco-ion.co.uk/wp-content/uploads/Top_Gun_UM_9752050420_NL_D_GB_F.pdf.
- [245] Jérôme Degallaix et al. “Large and extremely low loss: the unique challenges of gravitational wave mirrors”. In: *Journal of the Optical Society of America A* 36.11 (2019), pp. C85–C94. DOI: 10.1364/JOSAA.36.000C85. URL: <https://opg.optica.org/josaa/abstract.cfm?URI=josaa-36-11-C85>.
- [246] Elenna M. Capote et al. “In-vacuum measurements of optical scatter versus annealing temperature for amorphous Ta₂O₅ and TiO₂:Ta₂O₅ thin films”. In: *Journal of the Optical Society of America A* 38.4 (2021), pp. 534–541. DOI: 10.1364/JOSAA.415665. URL: <https://opg.optica.org/josaa/abstract.cfm?URI=josaa-38-4-534>.
- [247] Fabian Magaña-Sandoval et al. “Large-angle scattered light measurements for quantum-noise filter cavity design studies”. In: *Journal of the Optical Society of America A* 29.8 (2012), pp. 1722–1727. DOI: 10.1364/JOSAA.29.001722. URL: <https://opg.optica.org/josaa/abstract.cfm?URI=josaa-29-8-1722>.

- [248] P.Y. Bely. “*The design and construction of large optical telescopes*”. Springer New York, 2003. ISBN: 978-0-387-95512-4. DOI: 10.1007/b97612. URL: https://archive.org/details/springer_10.1007-b97612.
- [249] Oli Patane. *CSUF-GWPAC Angle Resolved Scatterometer setup image*. 2023.
- [250] Oleg Mazonka. “*Solid Angle of Conical Surfaces, Polyhedral Cones, and Intersecting Spherical Caps*”. 2013. arXiv: 1205.1396 [math.MG].
- [251] CSUF-GWPAC: Scatterometer. “*TS3 w BRDF*”. 2023. URL: <https://www.youtube.com/watch?v=VJ5JBpHPMGk>.
- [252] CSUF-GWPAC: Scatterometer. “*TS2 w BRDF*”. 2023. URL: <https://www.youtube.com/watch?v=JW2M2zyIwEg>.
- [253] L. Pinard. private communication. 2023.
- [254] Liyuan Zhang et al. “*The Coating Scattering and Absorption Measurements of LIGO I mirrors at Caltech*”. California Institute of Technology. URL: <https://dcc.ligo.org/LIGO-G080162/public>.
- [255] G. Billingsley. private communication. 2022.
- [256] M R Abernathy et al. “Cryogenic mechanical loss measurements of heat-treated hafnium dioxide”. In: *Classical and Quantum Gravity* 28.19 (2011), p. 195017. DOI: 10.1088/0264-9381/28/19/195017. URL: <https://dx.doi.org/10.1088/0264-9381/28/19/195017>.
- [257] C.Y. She. “Raman characterization of optical thin film coatings”. In: *Thin Solid Films* 154.1 (1987), pp. 239–247. ISSN: 0040-6090. DOI: [https://doi.org/10.1016/0040-6090\(87\)90368-3](https://doi.org/10.1016/0040-6090(87)90368-3). URL: <https://www.sciencedirect.com/science/article/pii/0040609087903683>.
- [258] Hon. J.W. Strutt. “XV. On the light from the sky, its polarization and colour”. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 41.271 (1871), pp. 107–120. DOI: 10.1080/14786447108640452. eprint: <https://doi.org/10.1080/14786447108640452>. URL: <https://doi.org/10.1080/14786447108640452>.

- [259] Hon. J.W. Strutt. “LVIII. On the scattering of light by small particles”. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 41.275 (1871), pp. 447–454. DOI: 10.1080/14786447108640507. eprint: <https://doi.org/10.1080/14786447108640507>. URL: <https://doi.org/10.1080/14786447108640507>.
- [260] Adolf Smekal. “Zur Quantentheorie der Dispersion”. In: *Naturwissenschaften* 11.43 (Oct. 1923), pp. 873–875. DOI: 10.1007/BF01576902.
- [261] C. V. Raman and K. S. Krishnan. “The Negative Absorption of Radiation”. In: 122.3062 (July 1928), pp. 12–13. DOI: 10.1038/122012b0.
- [262] Nobel Media AB. *Nobel Prize Announcement*. <https://www.nobelprize.org/prizes/physics/1930/summary/>. Accessed on: 2024-02-28. 1930.
- [263] C. Candler. *Practical Spectroscopy*. Prentice-Hall physics series. London, 1949. URL: <https://books.google.co.uk/books?id=Vio6AAAAAAAJ>.
- [264] J T Yates Jr and T E Madey. “Vibrational spectroscopy of molecules on surfaces”. In: (). URL: <https://www.osti.gov/biblio/6479927>.
- [265] DJ Gardiner, PR Graves and HJ Bowley. *Practical Raman spectroscopy. 1989*.
- [266] S. Svanberg. *Atomic and Molecular Spectroscopy: Basic Aspects and Practical Applications*. Advanced Texts in Physics. Springer Berlin Heidelberg, 2012. ISBN: 9783642185205. URL: <https://books.google.co.uk/books?id=v1T9CAAAQBAJ>.
- [267] W. Demtröder. *Laser Spectroscopy: Basic Concepts and Instrumentation*. Advanced Texts in Physics. Springer Berlin Heidelberg, 2002. ISBN: 9783540652250. URL: <https://books.google.co.uk/books?id=dNx10Lgn1xcC>.
- [268] E.U. Condon and G.H. Shortley. *The Theory of Atomic Spectra*. Cambridge Univ.Pr.209. Cambridge University Press, 1951. ISBN: 9780521092098. URL: https://books.google.co.uk/books?id=hPyD-Nc_YmgC.
- [269] L. Bergmann, C. Schaefer and H. Niedrig. *Bergmann/schaefer:optics (niedrig)*. W. de Gruyter, 1999. ISBN: 9783110143188. URL: <https://books.google.co.uk/books?id=sJxv-yCSNEAC>.
- [270] G. Placzek. *Rayleigh-Streuung und Raman-Effekt*. 1961. URL: <https://books.google.co.uk/books?id=s00pmgEACAAJ>.

- [271] W. A. P. Luck. “D. J. Gardiner, P. R. Graves (Eds.): Practical Raman Spectroscopy, mit Beiträgen von H. J. Bowley, D. J. Gardiner, D. L. Gerrard, P. R. Graves, J. D. Loudon, and G. Turrell, Springer-Verlag, Berlin, Heidelberg, New York 1989. 157 Seiten, brosch., Preis: DM 86,—.” In: *Berichte der Bunsengesellschaft für physikalische Chemie* 94.9 (1990), pp. 1047–1047. DOI: <https://doi.org/10.1002/bbpc.19900940938>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/bbpc.19900940938>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/bbpc.19900940938>.
- [272] David Tuschel. “Why are the Raman spectra of crystalline and amorphous solids different?” In: 32 (Mar. 2017), pp. 26–33.
- [273] L. S. Hsu and C. Y. She. “Real-time monitoring of crystallization and structural transformation of titania films with Raman spectroscopy”. In: *Optics Letters* 10.12 (1985), pp. 638–640. DOI: 10.1364/OL.10.000638. URL: <https://opg.optica.org/ol/abstract.cfm?URI=ol-10-12-638>.
- [274] Long-Sheng Hsu, Chiao-Yao She and Gregory J. Exarhos. “Reduction of substrate interference in Raman spectroscopy of submicron titania coatings”. In: *Applied Optics* 23.18 (1984), pp. 3049–3051. DOI: 10.1364/AO.23.003049. URL: <https://opg.optica.org/ao/abstract.cfm?URI=ao-23-18-3049>.
- [275] L. S. Hsu et al. “Thermally induced crystallization of amorphous-titania films”. In: *Journal of Applied Physics* 59.10 (May 1986), pp. 3475–3480. ISSN: 0021-8979. DOI: 10.1063/1.336817. eprint: https://pubs.aip.org/aip/jap/article-pdf/59/10/3475/9364434/3475_1_online.pdf. URL: <https://doi.org/10.1063/1.336817>.
- [276] Massimo Granata et al. “Correlated evolution of structure and mechanical loss of a sputtered silica film”. In: *Physical Review Materials* 2.5 (2018). DOI: 10.1103/physrevmaterials.2.053607. URL: <https://doi.org/10.1103/physrevmaterials.2.053607>.
- [277] V. Martinez. private communication. 2023.
- [278] B Lafuente et al. “The power of databases: the RRUFF project. In: Highlights in Mineralogical Crystallography - Anatase R060277.4 Raman Spectrum depolarised”. In: (2015). URL: <https://rruff.info/Anatase/R060277>.

- [279] L. Kernazhitsky et al. “Laser-excited excitonic luminescence of nanocrystalline TiO₂ powder”. In: *Ukrainian Journal of Physics* 59 (Mar. 2014), pp. 246–253. DOI: 10.15407/ujpe59.03.0246.
- [280] U. Balachandran and N.G. Erer. “Raman spectra of titanium dioxide”. In: *Journal of Solid State Chemistry* 42.3 (1982), pp. 276–282. ISSN: 0022-4596. DOI: [https://doi.org/10.1016/0022-4596\(82\)90006-8](https://doi.org/10.1016/0022-4596(82)90006-8). URL: <https://www.sciencedirect.com/science/article/pii/0022459682900068>.
- [281] Bhanu Mantri and Rajni Verma. “Shape Control Synthesis, Characterizations, Mechanisms And Optical Properties Of Larg Scaled Metal oxide Nanostructures Of ZnO And TiO₂”. In: *advanced materials letters 0976-3976* 6 (Dec. 2014), pp. 324–333. DOI: 10.5185/amlett.2015.5661.
- [282] Shuibo Xie, Enrique Iglesia and Alexis T. Bell. “Effects of Temperature on the Raman Spectra and Dispersed Oxides”. In: *The Journal of Physical Chemistry B* 105.22 (2001), pp. 5144–5152. DOI: 10.1021/jp004434s. eprint: <https://doi.org/10.1021/jp004434s>. URL: <https://doi.org/10.1021/jp004434s>.
- [283] Nathan Van Velson, Hamidreza Zobeiri and Xinwei Wang. “Thickness-Dependent Raman Scattering from Thin-Film Systems”. In: *The Journal of Physical Chemistry C* 127.6 (2023), pp. 2995–3004. DOI: 10.1021/acs.jpcc.2c06353. eprint: <https://doi.org/10.1021/acs.jpcc.2c06353>. URL: <https://doi.org/10.1021/acs.jpcc.2c06353>.
- [284] Dong Liu and Peter Flewitt. “Raman Measurements of Stress in Films and Coatings”. In: vol. 45. Jan. 2014, pp. 141–177. ISBN: ISBN (online): 978-1-78262-148-5 ISBN (print): 978-1-84973-919-1. DOI: 10.1039/9781782621485-00141.
- [285] Ingrid De Wolf. “Micro-Raman spectroscopy to study local mechanical stress in silicon integrated circuits”. In: *Semiconductor Science and Technology* 11.2 (1996), p. 139. DOI: 10.1088/0268-1242/11/2/001. URL: <https://dx.doi.org/10.1088/0268-1242/11/2/001>.
- [286] P E J Flewitt and R K Wild. *Physical Methods for Materials Characterisation (3rd ed.)* CRC Press, 2017. URL: <https://doi.org/10.1201/9781315382012>.

- [287] Garth Scannell, Akio Koike and Liping Huang. "Structure and thermo-mechanical response of TiO₂-SiO₂ glasses to temperature". In: *Journal of Non-Crystalline Solids* 447 (2016), pp. 238–247. ISSN: 0022-3093. DOI: <https://doi.org/10.1016/j.jnoncrysol.2016.06.018>. URL: <https://www.sciencedirect.com/science/article/pii/S0022309316302332>.
- [288] PETER C. SCHULTZ. "Binary Titania-Silica Glasses Containing 10 to 20 Wt% TiO₂". In: *Journal of the American Ceramic Society* 59.5-6 (1976), pp. 214–219. DOI: <https://doi.org/10.1111/j.1151-2916.1976.tb10936.x>. eprint: <https://ceramics.onlinelibrary.wiley.com/doi/pdf/10.1111/j.1151-2916.1976.tb10936.x>. URL: <https://ceramics.onlinelibrary.wiley.com/doi/abs/10.1111/j.1151-2916.1976.tb10936.x>.
- [289] Donald R. Sandstrom et al. "Coordination of Ti in TiO₂ SiO₂ glass by X-ray absorption spectroscopy". In: *Journal of Non-Crystalline Solids* 41.2 (1980), pp. 201–207. ISSN: 0022-3093. DOI: [https://doi.org/10.1016/0022-3093\(80\)90165-9](https://doi.org/10.1016/0022-3093(80)90165-9). URL: <https://www.sciencedirect.com/science/article/pii/0022309380901659>.
- [290] Ian Lewis and H. Edwards. *Handbook of Raman Spectroscopy*. Aug. 2001. ISBN: 0-8247-0557-2. DOI: 10.1201/9781420029253.
- [291] J. C. Parker and R. W. Siegel. "Raman microprobe study of nanophase TiO₂ and oxidation-induced spectral changes". In: *Journal of Materials Research* 5.6 (1990), pp. 1246–1252. DOI: 10.1557/JMR.1990.1246.
- [292] J. C. Parker and R. W. Siegel. "Calibration of the Raman spectrum to the oxygen stoichiometry of nanophase TiO₂". In: *Applied Physics Letters* 57.9 (Aug. 1990), pp. 943–945. ISSN: 0003-6951. DOI: 10.1063/1.104274. eprint: <https://pubs.aip.org/aip/apl/article-pdf/57/9/943/7778553/943\1\online.pdf>. URL: <https://doi.org/10.1063/1.104274>.
- [293] A. Krost et al. "Chemical reaction at the ZnSe/GaAs interface detected by Raman spectroscopy". In: *Applied Physics Letters* 57.19 (Nov. 1990), pp. 1981–1982. ISSN: 0003-6951. DOI: 10.1063/1.104149. eprint: <https://pubs.aip.org/aip/apl/article-pdf/57/19/1981/7777239/1981\1\online.pdf>. URL: <https://doi.org/10.1063/1.104149>.

- [294] David Tuschel. “Practical Group Theory and Raman Spectroscopy, Part I: Normal Vibrational Modes”. In: *Spectroscopy* 29 (Feb. 2014), p. 14.
- [295] Daniel T. Larson and Efthimios Kaxiras. “Raman spectrum of CrI₃: An *ab initio* study”. In: *Physical Review B* 98 (8 2018), p. 085406. DOI: 10.1103/PhysRevB.98.085406. URL: <https://link.aps.org/doi/10.1103/PhysRevB.98.085406>.
- [296] Maria O’Brien et al. “Low wavenumber Raman spectroscopy of highly crystalline MoSe₂ grown by chemical vapor deposition”. In: *physica status solidi (b)* 252.11 (2015), pp. 2385–2389. DOI: <https://doi.org/10.1002/pssb.201552225>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/pssb.201552225>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/pssb.201552225>.
- [297] Chia-Ti Wu et al. “Anisotropic effects in the Raman scattering of Re-doped 2H-MoSe₂ Layered Semiconductors”. In: *Results in Physics* 7 (Oct. 2017). DOI: 10.1016/j.rinp.2017.10.033.
- [298] Sawanta S. Mali et al. “Hydrothermal synthesis of rutile TiO₂ with hierarchical microspheres and their characterization”. In: *CrystEngComm* 13 (21 2011), pp. 6349–6351. DOI: 10.1039/C1CE05928A. URL: <http://dx.doi.org/10.1039/C1CE05928A>.
- [299] Changgu Lee et al. “Anomalous Lattice Vibrations of Single- and Few-Layer MoS₂”. In: *ACS Nano* 4.5 (2010). PMID: 20392077, pp. 2695–2700. DOI: 10.1021/nn1003937. eprint: <https://doi.org/10.1021/nn1003937>. URL: <https://doi.org/10.1021/nn1003937>.
- [300] “TFCalc”. “*TFCalc*”. Tech. rep. HU LInks, 2024. URL: <https://www.hulinks.co.jp/en/tfcalc-e>.
- [301] G.A. Ascoli, J. Bezhanskaya and V. Tsytsarev. “Microscopy”. In: *Encyclopedia of the Neurological Sciences (Second Edition)*. Ed. by Michael J. Aminoff and Robert B. Daroff. Second Edition. Oxford: Academic Press, 2014, pp. 16–20. ISBN: 978-0-12-385158-1. DOI: <https://doi.org/10.1016/B978-0-12-385157-4.00587-X>. URL: <https://www.sciencedirect.com/science/article/pii/B978012385157400587X>.
- [302] David Tuschel. “The effect of microscope objectives on the raman spectra of crystals”. In: *Spectroscopy (Santa Monica)* 32 (Sept. 2017), pp. 14–23.

- [303] Sheng Liu and Hong Hua. “Extended depth-of-field microscopic imaging with a variable focus microscope objective”. In: *Optics Express* 19.1 (2011), pp. 353–362. DOI: 10.1364/OE.19.000353. URL: <https://opg.optica.org/oe/abstract.cfm?URI=oe-19-1-353>.
- [304] Biswajit Choudhury and Amarjyoti Choudhury. “Local structure modification and phase transformation of TiO₂ nanoparticles initiated by oxygen defects, grain size, and annealing temperature”. In: *International Nano Letters* 3 (Nov. 2013), p. 55. DOI: 10.1186/2228-5326-3-55.
- [305] Kunlun Ding et al. “Study on the Anatase to Rutile Phase Transformation and Controlled Synthesis of Rutile Nanocrystals with the Assistance of Ionic Liquid”. In: *Langmuir* 26.12 (2010). PMID: 20426393, pp. 10294–10302. DOI: 10.1021/la100468e. eprint: <https://doi.org/10.1021/la100468e>. URL: <https://doi.org/10.1021/la100468e>.
- [306] Sonja Rau et al. “Acoustic properties of oxide glasses at low temperatures”. In: *Physical Review B* 52 (10 1995), pp. 7179–7194. DOI: 10.1103/PhysRevB.52.7179. URL: <https://link.aps.org/doi/10.1103/PhysRevB.52.7179>.
- [307] I W Martin et al. “Low temperature mechanical dissipation of an ion-beam sputtered silica film”. In: *Classical and Quantum Gravity* 31.3 (2014), p. 035019. DOI: 10.1088/0264-9381/31/3/035019. URL: <https://dx.doi.org/10.1088/0264-9381/31/3/035019>.
- [308] Sonja Rau et al. “Acoustic properties of oxide glasses at low temperatures”. In: *Physical Review B* 52 (10 1995), pp. 7179–7194. DOI: 10.1103/PhysRevB.52.7179. URL: <https://link.aps.org/doi/10.1103/PhysRevB.52.7179>.
- [309] K. A. Topp and David G. Cahill. “Elastic properties of several amorphous solids and disordered crystals below 100 K”. English (US). In: *Zeitschrift fur Physik B-Condensed Matter* 101.2 (Oct. 1996), pp. 235–245. ISSN: 0722-3277. DOI: 10.1007/s002570050205.
- [310] J. Classen et al. “Low frequency acoustic and dielectric measurements on glasses”. In: *Annalen der Physik* 506.5 (1994), pp. 315–335. DOI: <https://doi.org/10.1002/andp.19945060502>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/andp.19945060502>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/andp.19945060502>.

- [311] Heraeus. “*Quartz Glass for Optics Data and Properties*”. Tech. rep. Buford, GA, USA: Heraeus, 2007. URL: <https://markoptics.com/wp-content/uploads/2019/03/Heraeus-Quartz-Fused-Silica.pdf>.
- [312] G. Vajente. private communication. 2024.
- [313] S. Tait. private communication. 2023.
- [314] A.. Davenport. “*Progress on TiO₂:GeO₂/SiO₂ coatings*”. The LIGO Scientific Collaboration. 23rd Sept. 2023. URL: <https://dcc.ligo.org/LIGO-G2301806public> (visited on 28/02/2024).
- [315] G Vajente et al. “Effect of elevated substrate temperature deposition on the mechanical losses in tantala thin film coatings”. In: *Classical and Quantum Gravity* 35.7 (2018), p. 075001. DOI: 10.1088/1361-6382/aaad7c. URL: <https://dx.doi.org/10.1088/1361-6382/aaad7c>.
- [316] Xiao Liu et al. “Hydrogen-Free Amorphous Silicon with No Tunneling States”. In: *Phys. Rev. Lett.* 113 (2 2014), p. 025503. DOI: 10.1103/PhysRevLett.113.025503. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.113.025503>.
- [317] A. Markosyan and M. Fejer. private communication. 2024.
- [318] N. Demos and S. Gras. private communication. 2024.
- [319] K. Hono and S.S. Babu. “15 - Atom-Probe Field Ion Microscopy”. In: *Physical Metallurgy (Fifth Edition)*. Ed. by David E. Laughlin and Kazuhiro Hono. Fifth Edition. Oxford: Elsevier, 2014, pp. 1453–1589. ISBN: 978-0-444-53770-6. DOI: <https://doi.org/10.1016/B978-0-444-53770-6.00015-0>. URL: <https://www.sciencedirect.com/science/article/pii/B9780444537706000150>.
- [320] Zygo Corporation. “*Verifire Specifications*”. Tech. rep. Middlefield, CT, USA: Zygo Corporation, 2018. URL: https://www.lambdaphoto.co.uk/pdfs/Zygo/LAMBDA_Verifire-Specs.pdf.
- [321] Aidan F. Brooks et al. “Point absorbers in Advanced LIGO”. In: *Appl. Opt.* 60.13 (2021), pp. 4047–4063. DOI: 10.1364/AO.419689. URL: <https://opg.optica.org/ao/abstract.cfm?URI=ao-60-13-4047>.
- [322] A. Brooks P. Willems M. Smith K. Mailand. “*Advanced LIGO Thermal Compensation System Preliminary Design*”. Tech. rep. The LIGO Scientific Collaboration, 2009. URL: <https://dcc.ligo.org/LIGO-T0900304/public>.

- [323] W. Kells G. Billingsley G. Harry. “*Core Optics Components Design Requirements Document*”. Tech. rep. The LIGO Scientific Collaboration, 2008. URL: <https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?.submit=Identifier&docid=T080026&version=>.
- [324] Aidan F. Brooks et al. “Overview of Advanced LIGO adaptive optics”. In: *Appl. Opt.* 55.29 (2016), pp. 8256–8265. DOI: 10.1364/AO.55.008256. URL: <https://opg.optica.org/ao/abstract.cfm?URI=ao-55-29-8256>.
- [325] G. Vajente. “In situ correction of mirror surface to reduce round-trip losses in Fabry–Perot cavities”. In: *Appl. Opt.* 53.7 (2014), pp. 1459–1465. DOI: 10.1364/AO.53.001459. URL: <https://opg.optica.org/ao/abstract.cfm?URI=ao-53-7-1459>.
- [326] Wenxuan Jia et al. “Point Absorber Limits to Future Gravitational-Wave Detectors”. In: *Phys. Rev. Lett.* 127 (24 2021), p. 241102. DOI: 10.1103/PhysRevLett.127.241102. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.127.241102>.
- [327] W. Winkler et al. “Heating by optical absorption and the performance of interferometric gravitational-wave detectors”. In: *Phys. Rev. A* 44 (11 1991), pp. 7022–7036. DOI: 10.1103/PhysRevA.44.7022. URL: <https://link.aps.org/doi/10.1103/PhysRevA.44.7022>.
- [328] P. Fritschel and D. Coyne. “*Advanced LIGO Systems Design*”. Tech. rep. Pasadena, CA, USA: The LIGO Scientific Collaboration, 2015. URL: <https://dcc.ligo.org/LIGO-T010075/public>.
- [329] G. Billingsley. private communication. 2024.
- [330] Zygo Corporation. “*ZYGOS guide to Typical Interferometer Setups*”. Tech. rep. Middlefield, CT, USA: Zygo Corporation, 2001. URL: <https://www.zygo.com/-/media/project/ameteksxa/zygo/ametekzygo/downloadables/brochures/interferometers/typical-interferometer-setups.pdf>.
- [331] Zygo Corporation. “*Transmission Sphere Selection*”. Tech. rep. Middlefield, CT, USA: Zygo Corporation, 2014. URL: <https://www.zygo.com/-/media/project/ameteksxa/zygo/ametekzygo/downloadables/brochures/interferometers/transmission-sphere-selection-guide.pdf?dmc=1&revision=50654516-f92e-4837-8972-11dcc708f318>.

- [332] von F. Zernike. “Beugungstheorie des schneidenverfahrens und seiner verbesserten form, der phasenkontrastmethode”. In: *Physica* 1.7 (1934), pp. 689–704. ISSN: 0031-8914. DOI: [https://doi.org/10.1016/S0031-8914\(34\)80259-5](https://doi.org/10.1016/S0031-8914(34)80259-5). URL: <https://www.sciencedirect.com/science/article/pii/S0031891434802595>.
- [333] Matthe T. D. Cashmore. “Interferometric Metrology Using Reprogrammable Binary Holograms”. Available at <https://core.ac.uk/download/pdf/16281492.pdf>. PhD thesis. Durham University, 2013.
- [334] James Wyant and Katherine Creath. “Basic Wavefront Aberration Theory for Optical Metrology”. In: *Appl Optics Optical Eng* 11 (Jan. 1992).
- [335] J. Y. Wang and D. E. Silva. “Wave-front interpretation with Zernike polynomials”. In: *Applied Optics* 19.9 (1980), pp. 1510–1518. DOI: 10.1364/AO.19.001510. URL: <https://opg.optica.org/ao/abstract.cfm?URI=ao-19-9-1510>.
- [336] E.P. Goodwin and J.C. Wyant. *Field Guide to Interferometric Optical Testing*. Field Guides. SPIE, 2006. ISBN: 9780819465108. URL: <https://books.google.co.uk/books?id=x-FzKsL0rm8C>.
- [337] Zygo Corporation. “*MetroPro Reference Guide OMP-0347J*”. Tech. rep. Middlefield, CT, USA: Zygo Corporation, 2012. URL: <https://www.zygo.com/-/media/project/ameteksxa/zygo/ametekzygo/downloadables/brochures/interferometers/transmission-sphere-selection-guide.pdf?dmc=1&revision=50654516-f92e-4837-8972-11dcc708f318>.
- [338] G. Billingsley. private communication. 2022.
- [339] G. Billingsley. “*Advanced LIGO End Test Mass (ETM) A+ version*”. Tech. rep. Pasadena, CA, USA: The LIGO Scientific Collaboration, 2020. URL: <https://dcc.ligo.org/LIGO-E080512/public>.
- [340] The LIGO Scientific Collaboration. “Advanced LIGO”. In: *Classical and Quantum Gravity* 32.7 (Mar. 2015), p. 074001. ISSN: 1361-6382. DOI: 10.1088/0264-9381/32/7/074001. URL: <http://dx.doi.org/10.1088/0264-9381/32/7/074001>.
- [341] Craig Cahillane and Georgia Mansell. “Review of the Advanced LIGO Gravitational Wave Observatories Leading to Observing Run Four”. In: *Galaxies* 10.1 (Feb. 2022), p. 36. ISSN: 2075-4434. DOI: 10.3390/galaxies10010036. URL: <http://dx.doi.org/10.3390/galaxies10010036>.

-
- [342] GariLynn Billingsley. “*ITM03 Figure Measurement Report*”. Tech. rep. Pasadena, CA, USA: The LIGO Laboratory Caltech, 2017. URL: <https://dcc.ligo.org/LIGO-E1300943>.
- [343] GariLynn Billingsley. “*ITM11 Figure Measurement Report*”. Tech. rep. Pasadena, CA, USA: The LIGO Laboratory Caltech, 2014. URL: <https://dcc.ligo.org/LIGO-E1400037/public>.
- [344] G. Billingsley. “*Characterizing the figure change expected after coating at LMA*”. Tech. rep. The LIGO Scientific Collaboration, 2020. URL: <https://dcc.ligo.org/public/0171/T2000643/001/Characterizing%20the%20coating%20plume%20at%20LMA.pdf>.
- [345] G. Harry. “*Advanced LIGO Input Test Mass Coating Specification*”. Tech. rep. Pasadena, CA, USA: The LIGO Scientific Collaboration, 2014. URL: <https://dcc.ligo.org/LIGO-E0900041/public>.