Using Electron Energy Loss Spectroscopy to Investigate the Structure of the Mirror Coatings of Gravitational Wave Detectors

Rebecca Beatrice Cummings
MPhys (Hons)

Submitted in partial fufilment of the requirements for the Degree

of

Doctor of Philosophy

Materials and Condensed Matter Physics Group

School of Physics and Astronomy

College of Science and Engineering

October 2022

Abstract

Reducing the noise in gravitational wave (GW) detectors is of upmost importance for improving their sensitivity, which enables them to detect smaller GWs. There are many sources of noise in GW detectors, but the dominating source at their most sensitive frequency is the thermal noise of the coatings of their mirrors. While the relationship between the optical properties of the coatings and their manufacturing process is reasonably well understood, the link between these two things and their structure is not. Fully understanding their structure will improve the search for new coatings with lower noise contributions.

In order to study the coatings, scanning transmission electron microscopy was used, specifically dark-field imaging and electron energy-loss spectroscopy (EELS). Two different techniques were used, one on each of the two coating types (tantala and silica, which have a low and high refractive index, respectively).

The first technique was performing spatially resolved quantification using high angle annular dark field imaging and DualEELS (a form of electron energy loss spectroscopy (EELS) that takes two spectra in quick succession from the high- and low-loss region). It has been established for some time that there are small spherical voids, or bubbles, in the coatings and it has been hypothesised that they are contributing to laser damage in the coatings. It was also known that there was argon in the coatings, which is left over from the ion-beam assisted deposition process used in preparing the coatings. Semi-empirical standards were created for the quantification using argon data from the EELS

Atlas and experimental data scaled using a Hartree Slater cross-section. The densities and pressures of the bubbles were then calculated. It was found that the argon coalesced into bubbles post-annealing in all samples, and that greater annealing temperatures caused the bubbles to coarsen. The bubbles' interior pressures decreased with diameter. The technique was also applied to xenon and krypton bubbles in Zircaloy-4. While these are not GW mirror coatings, they corroborate the findings in the GW samples and show the wide-ranging applicability of the technique.

The second technique was investigating the extended electron energy-loss fine structure (EXELFS) of silica, to see if there are visible differences in the EXELFS of samples made using different techniques. While EXELFS has been around for many years, it is used much less commonly than its X-ray absorption spectroscopy counterpart, EXAFS. This is because, depsite EXELFS' many advantages over EXAFS, the signal to noise ratio in EXELFS has traditionally been much lower than in EXAFS. This chapter investigated whether recent advances in EELS data acquisition meant it would a viable analysis technique. It was determined that EXELFS could be used, and that different silica samples did show subtle differences in their EXELFS.

The two techniques used and developed in this thesis have applications far beyond GW mirror coatings and will be of use to anyone wishing to perform characterisation and quantification of thin films using EELS.

"...only as you gasp your dying breath shall you understand, your life amounted to no more than one drop in a limitless ocean!"

Yet what is any ocean but a multitude of drops?

— David Mitchell, *Cloud Atlas*.

Declaration

This thesis was written by me, Rebecca Beatrice Cummings, and is a record of the work I carried out in the Materials and Condensed Matter group of the School of Physics and Astronomy at the University of Glasgow between 2018 and 2022 under the supervision of Ian MacLaren and Iain Martin.

Some of the work described in this thesis can be found in the following publications:

1. *Advances in Imaging and Electron Physics: Chapter Twelve - Performing EELS at higher energy losses at both 80 and 200 kV*, I. MacLaren, **R.B. Cummings**, F. Gordon, F. Frutos-Myro, S. McFadzean, A.P. Brown, A.J. Craven, Elsevier, Vol. 210, Pages 299-355 [1].

2. *Direct Detection EELS at High Energy: Elemental Mapping and EXELFS*, J. Hart, A. Lang, **R.B. Cummings**, I. MacLaren, M. Taheri, Microscopy and Microanalysis, Vol. 25, Supplement S2, Pages 584-585 [2].

3. *Argon bubble formation in tantalum oxide-based films for gravitational wave interferometer mirrors,*, **R.B. Cummings**, R. Bassiri, I.W. Martin, and I. MacLaren, Optical Materials Express, 2021, Vol. 11, Pages 707-718 [3].

4. *Xenon bubbles formed by ion implantation in zirconium alloy films*, **R.B. Cummings**, Matthew S. Blackmur, Mateusz Grunwald, Andrew Minty, Paul Styman, and Ian MacLaren, Journal of Nuclear Materials, 2022, Volume 560 [4].

5. *EELS Primer*, M. Taheri, I. MacLaren, **R.B. Cummings**, et al., Nature Methods and Primers, in preparation [5].

Paper five is relevant in Chapter 3, papers three, four and five in Chapter 4, and one and two in Chapter 5.

This thesis has not been submitted elsewhere for any other degree or qualification.

Acknowledgements

Firstly my thanks go to my supervisor, Ian MacLaren. This thesis would never have been completed without his help and guidance. My thanks also to Iain Martin, my second supervisor, whose help with the gravitational wave side of my work has been invaluable. Special thanks to Sam McFadzean for all his help with technical questions over the course of my PhD. I also wish to thank the rest of MCMP, staff and students alike, for ensuring my time in the group has been (mostly!) enjoyable.

To my friends in Glasgow (Zara, Andrew, Jack, Aidan, Sandaigh, Euan, George), thank for you for making my time in this city so wonderful, despite the pandemic. Particularly Emily, for dragging me out of the city when our schedules aligned. Eilidh, I'm so grateful we got to live together for four years, I already miss our film nights. Thanks especially to Adam, for letting me spend most of my free time in the last few months being stressed on your sofa, and making sure I ate. I'm so grateful for your kindness and I appreciate all the pep talks you gave me, even if you had to repeat them dozens of times. I'm less appreciative that you made me into a coffee snob but I can't even really complain about that either because it led to being friends with everyone at Black Pine. Thanks to them, not only for providing the best coffee in Glasgow, but also for the many times you all let me stay past close so I could finish just one more paragraph.

Thank you also to all my other friends further afield, Imogen, David and Charlie. Caitlin, who's always there to share joys and grumbles alike; waking up to a dozen messages from 2 a.m., about I show don't watch is a source of constant amusement. Meg, who made me a better writer. Emma, my partner in ~~crime~~ baking and whose friendship was the best thing to come out of undergrad. Isobel, my oldest friend, two decades in and I can't wait to see what nonsense the rest will bring. Gavin, thank you for, well, everything – most especially your love, patience and encouragement, even from 800 miles away – but also for introducing me to the West Wing, and being the best possible rubber duck.

I'm so lucky to have so many friends whom I love, and who love me. You've all, in your own ways, made my life and PhD so much better.

Finally to my family, Grandad, Diana, and Mum and Dad: thank you for everything you've done for me, all the years of love, support, and believing in me. I couldn't be more grateful.

—

AMDG

—

List of Abbreviations and Acronyms

ADF: annular dark field

aLIGO: Advanced LIGO

ARM: atomic resolution microscope

BF: bright field

BH: black hole(s)

CCD: charged coupled device

CMB: cosmic microwave background

DF: dark field

EELS: electron energy-loss spectroscopy

ELFS: energy-loss fine structure (combination of ELNES and EXELFS)

ELNES: energy-loss near-edge structure

EoS: equation of state

EXAFS: extended X-ray absorption fine structure

EXELFS: extended electron energy-loss fine structure

FEG: field emission gun

FT: Fourier transform

GW: gravitational wave(s)

HAADF: high angle annular dark field

IBAD: ion beam assisted deposition

IBSD: ion beam sputtering deposition

LIGO: laser interferometer gravitational wave observatory

LRO: long range order

MFP: mean free path

RDF: radial distribution function

SRO: short range order

STEM: transmission eelctron microscope/microscopy

TEM: transmission eelctron microscope/microscopy

XAFS: X-ray absorption (combination of EXAFS and XANES)

XAS: X-ray absorption spectroscopy

XANES: X-ray absorption near edge structure

ZLP: zero-loss peak

Contents

List of Figures

List of Tables

Thesis Overview

The goal of this thesis is twofold:

1. To provide insight into the atomic structure of gravitational wave detector mirror coatings, with a view to helping reduce the noise that they contribute, which will make the detectors more sensitive.

2. To improve the electron energy loss spectroscopy (EELS) characterisation and quantification techniques used for analysing the coatings.

Chapter 1 contains a brief history of gravitational waves and their detection, as well as the underlying theory. The design of gravitational wave detectors is discussed, with a particular focus on the detectors' mirrors and their coatings.

Chapter 2 provides a similar function as the previous chapter, this time on electron microscopy, with a particular focus on scanning transmission electron microscopy. It discusses the structure of the electron microscope used for almost all of the data collection in this thesis

Chapter 3 is the last background chapter and focusses on EELS. It goes over the different areas of an EELS spectrum; DualEELS; absolute quantification; and the steps of processing EELS data. It also discusses extended x-ray absorption fine structure (EXAFS) and its similarities to extended energy loss fine structure (EXELFS), and how the same techniques used for EXAFS analysis can be used for EXELFS.

Chapter 4 is the first results chapter and uses a combination of DualEELS and high angle annular dark field (HAADF) iamging to perform absolute spatial quantification on a variety of samples to examine their noble gas content.

Chapter 5 is the second results chapter. It investigates whether EXELFS is a viable analysis technique following recent advancements in the signal to noise ratio of EELS; and if there are visible differences in the EXELFS of silica samples deposited in different environments.

Chapter 6 is the conclusion, and summarises the whole thesis, before discussing potential further work. It finishes with some final thoughts on the results of the two experimental chapters.

A Note About the Impact of the Pandemic

The work in section 4.4 would perhaps, in a pre-pandemic era, not make up such a large portion of this chapter, as the film material is Zircaloy-4 – used for the cladding of nuclear reactors, not GW detector mirror coatings. However, even before the pandemic, there was no access to the microscope from September 2019 due to a fault with the spectrometer. Subsequent delays caused by the pandemic meant that it was unavailable until late 2020. There were also considerable periods of downtime in 2021, exacerbated by the pandemic. This combination of factors meant that old data containing the same type of gas-filled voids were analysed. While they do not shed light themselves on potential improvements and noise sources in the coatings, they corrobate the findings in the GW samples; help to identify what would happen in similar fabrication methods; and provide more general insight on thin film properties.

1

Gravitational Waves

1.1 History of Gravitational Wave Detection

Gravitational waves (GW) were first directly detected in 2015 by the LIGO and Virgo scientific collaboration [6], just under a century after they were first predicted by Einstein in his 'General Theory of Relativity' [7]. As the name suggests, GW are waves, or ripples, in space-time. They travel at the speed of light and are caused by the acceleration of a non-spherically-symmetric mass distribution. They remained unobserved for such a long time partially because of how weak they are when compared with the interactions of other forces. Only very large masses, such as black holes (BH) merging and supernovae, give off GW large enough that they can be detected by current detectors. GW exert strain on space – changing the length of the objects that they pass through. The strain is on the order of $10^{-23}\text{Hz}^{-\frac{1}{2}}$ [8], which translates

to changing the 4 km arm of LIGO by 1 am (10^{-18} m). Today, one of the main goals of the GW community is to make detectors even more sensitive, so that smaller GW can be identified.

The reason GW are so important is that they can give us information about our universe and its origins that is not available by other means. The early universe, for example, was opaque to electromagnetic radiation [9], so GW provide a way to investigate this period. BH mergers also cannot be observed using traditional telescopes.

The first indirect proof for GW was published by Hulse, Taylor and Weisberg [10, 11]. In 1975, Hulse and Taylor detected a binary pulsar system and following work in 1982 by Taylor and Weisberg showed that the decay of the pulsar's orbit was consistent with the energy that would be lost by GW emission. The first proposal of a GW detector came before this indirect proof and was designed by Weber – a collection of aluminium cylinders designed such that their resonant frequencies would be excited by passing GW [12]. Weber claimed to have detected GW on several occasions [13, 14, 15] but his results could never be duplicated. Gertsenshtein put forward interferometers [16] and they remain the detectors of choice for GW observation, with advancements being made throughout the rest of the twentieth century [17, 18, 19, 20]. Long-baseline interferometers were built by the early 2000s in multiple places around the world; GEO600 in Germany [21], LIGO in the United States [19], TAMA in Japan (now superceeded by KAGRA [22]) and Virgo in Italy [23]. By 2011 all these facilities had joined together to work collaboratively. LIGO was taken offline in 2010 so it could be upgraded to Advanced LIGO (aLIGO), which made it able to detect signals 10 times smaller than its predecessors [24]. It came back online in 2015 and made the first direct detection of GW on the second day of its initial operating run [25]. In total, 90 GW events have now

been catalogued with many more possible events also recorded [26]. Further updates to aLIGO are planned so that smaller GW can be detected.

1.2 Describing Gravitational Waves

The space-time interval (ds) is one of the most important concepts in GW detection. The difference between two points in a co-ordinate system (t, x, y, z) can be described as;

$$ds^2 = -c^2dt^2 + dx^2 + dy^2 + dz^2. \tag{1.1}$$

Equation (1.1) can be generalised to;

$$ds^2 = g_{\alpha\beta}dx^\alpha dx^\beta, \tag{1.2}$$

where $g_{\alpha\beta}$ is the metric tensor and is symmetric in any coordinate system. It describes the local curvature of spacetime and α and β can be (0, 1, 2, 3), representing (t, x, y, z). The metric tensor is a combination of the Minkowski tensor (η), which represents flat space-time and another term (h) that accounts for small changes in curvature;

$$g_{\alpha\beta} = \eta_{\alpha\beta} + h_{\alpha\beta}. \tag{1.3}$$

This tensor can be chosen such that Einstein's field equations take the form of a wave equation;

$$(\nabla^2 - \frac{1}{c^2}\frac{\delta^2}{\delta x^2}) = h_{\alpha\beta}. \tag{1.4}$$

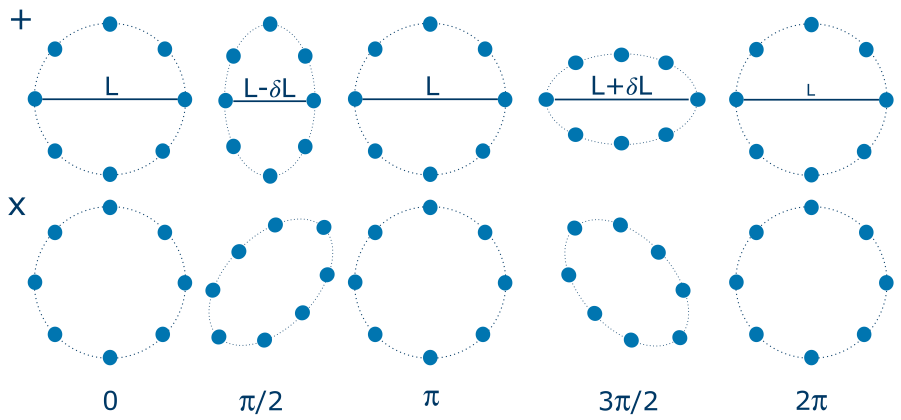

Figure 1.1: As the GW passes through the ring of test masses, it increases the distance between them in one plane while decreasing it in the orthogonal one, demonstrating the quadropolar nature of gravitational radiation.

In this case, we take $h_{\alpha\beta}$ to be a pertubation in length, describing a GW moving along the z-axis of a system;

$$h_{\alpha\beta} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_{xx} & h_{xy} & 0 \\ 0 & h_{yx} & h_{yy} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \tag{1.5}$$

where;

$$h_{xx} = -h_{yy},$$

$$h_{xy} = h_{yx}.$$

This results in two polarisations of the GW (at a 45° angle to one another), known as the cross ($\times$) and plus ($+$) polarisations. The effect they have on a ring of test masses, is illustrated in fig. 1.1. The strain from the change in the distance between two particles on opposite sides of the ring is therefore;

$$h = \frac{2\delta L}{L}. \tag{1.6}$$

It is this strain[1] that is measured by GW detectors. Instead of one ring of test masses there are two; one at each end of the interferometer arms. The strain is calculated from the phase difference of the lasers between the two interferometer arms.

1.3 Gravitational Wave Sources

Although GW are produced by any masses accelerating asymmetrically, astronomical sources of them are generally sorted into three categories; transient sources, continuous sources and the stochastic background.

1.3.1 Transient Sources

Transient sources are sometimes also called burst sources and come from big astronomical events, such as supernovae and compact binary mergers. Compact binary coalescences (CBC) are the best understood of all GW sources, and it is therefore easier to isolate smaller GW from them than is possible from other sources. This is because their waveforms can be calculated with good precision, meaning specific analysis techniques can be used, rather than more general techniques needed for other sources. This results in greater sensitivity [27].

Compact binary systems are made up of two massive objects such as BH and neutron stars. There are three stages to a CBC, the inspiral, merger and ringdown. The inspiral is the decay of the orbital radius of the system, the merger when they coalesce and the ringdown is the system settling to a stationary state [28].

As the two masses grow closer together the frequency of the GW they give out increases. The signal given out is known as a chirp and gives rise to a strain;

[1]Note this is not the same h as in previous equations, their most common symbols just happen to be the same.

$$h = 10^{-23} {m_{tot}}^{\frac{2}{3}} \mu \left(\frac{f}{100 \text{ Hz}} \right)^{\frac{2}{3}} \left(\frac{100 \text{ Mpc}}{r} \right), \tag{1.7}$$

where m_{tot} is the combined mass of the two objects, μ is the reduced mass of the system, f is the frequency and r is the distance from the earth.

Type II supernovae are another burst source; they can occur when massive stars reach the end of their lifespan. The stars can no longer sustain sufficient nuclear fusion in their cores to counteract their gravity and they collapse in on themselves. Exactly how the core collapses depends on the mass of the star and its composition [29]. No GW have been observed from a supernova as of yet; calculations suggest with current detector sensitivity it would only be possible to detect GW from a supernova within the Milky Way and Magellanic Clouds, which happen very infrequently [30, 31, 32, 33]. Additionally, the collapse must be asymmetric for a GW to be produced, so the star must still have angular momentum [34]. It is expected that the strain given out would be on the order of 10^{-21} at a distance of 10 kpc [35].

1.3.2 Continuous Sources

The first indirect proof of GW was a continous source – the pulsar PSR B1913+16 (see section 1.1). Pulsars are neutron stars with a very high rotational frequency that emit beams of radio waves from their magnetic poles. This results in them giving off a lighthouse-like signal. They and all other types of neutron stars are formed post-supernovae.

The asymmetric component necessary for GW production occurs in pulsars with magnetic poles not aligned to their rotation [36]. The strain for a typical neutron star would be;

$$h \approx 6 \times 10^{-25} \left(\frac{f_{rot}}{500\text{Hz}} \right)^2 \left(\frac{1\text{kpc}}{r} \right) \left(\frac{\epsilon}{10^{-6}} \right), \tag{1.8}$$

where f_{rot} is the rotation frequency, r is the distance to Earth and ϵ is the pulsar's equatorial ellipticity [37].

The other type of continous source is low-mass X-ray binary systems, which also include neutron stars. The neutron star accretes mass from its companion; the source of the necessary assymetry and also the eponymous X-rays. They are currently undetectable, for two reasons: the expected strain is much lower than from other sources and it is in a strain range where ground based detectors have much lower sensitivity (see section 1.5.2.5). The strain;

$$h \approx 8.5 \times 10^{-27} \left(\frac{r}{10\text{kpc}} \right)^{-1} \left(\frac{f_{rot}}{500\text{Hz}} \right)^{-\frac{1}{2}} \left(\frac{\dot{M}}{10^{-8} M_{\odot} \text{yr}^{-1}} \right)^{\frac{1}{2}}, \tag{1.9}$$

where r is the distance from Earth, f_{rot} is the rotation frequency, $M_{\odot}$ is the mass of the sun, and $\dot{M}$ is the accretion rate [38].

Stochastic Background

The stochastic background is analagous to the cosmic microwave background. It is made up of a large number of random and unresolved sources and contains information about the very early universe [25]. The strain is predicted to be;

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

where Ω_{GW} is the energy density required for a closed universe. Finding the stochastic background is one of the main goals of GW detection as it would shed light on the structure of the early universe. There are plans for a space based detector, Laser Interferometer Space Antenna (LISA), that is expected to launch in the 2030s [39]. One of its aims is to search for the stochastic background between 0.1 and 100 mHz [40], and it is guaranteed to detect the portion of the stochastic background caused by the galactic population of compact white dwarf binaries [41]. There are also suggested earth-based detectors searching for GW between 0.3 and 3 Hz, which will fill the frequency gap between LISA and the LIGO-Virgo detectors [42]. There are also proposals for a Japanese space-based detector, DECi-hertz Interferometer Gravitational wave Observatory (DECIGO), to search between 0.1 and 10 Hz [43].

1.4 Gravitational Wave Detectors

1.4.1 Interferometer design

As stated (see section 1.1), interferometers are the detector of choice for GW astronomy. They are essentially Michelson interferometers, made up of two perpendicular arms. The laser beam is split into two using a beam-splitter and one goes down each beam. The separate beams are reflected back at the mirrors and then recombined and collected by a photodetector (see fig. 1.2). The GW change the length of the interferometer arms, which changes the phase of the reflected light and therefore the interference pattern.

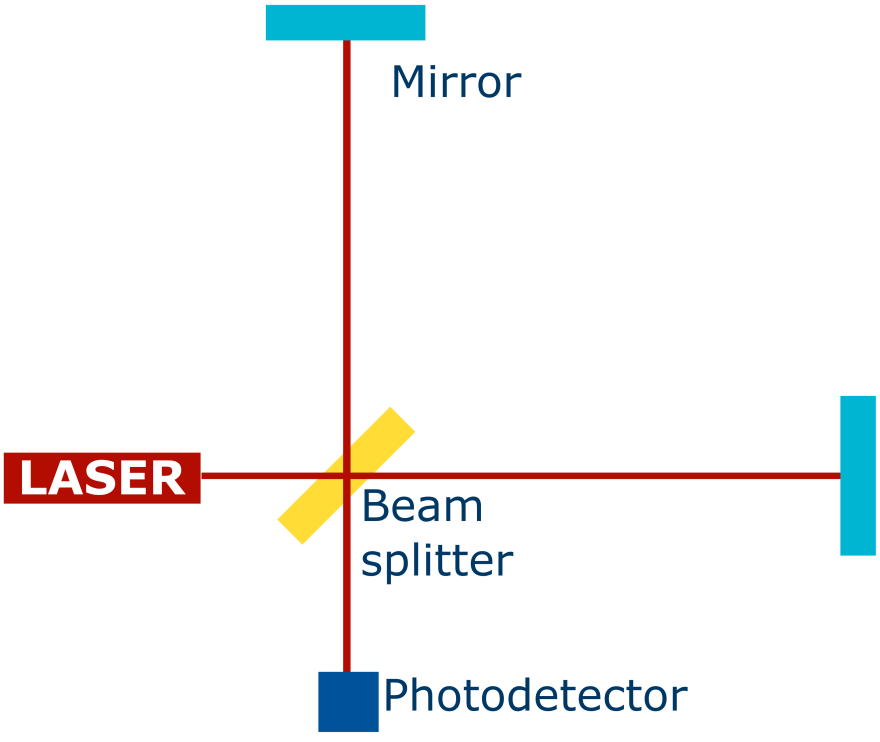

Figure 1.2: Diagram of a simple interferometer. There is a single laser with its beam split by the beam splitter. The two beams are reflected back by the mirrors at the end of the arms and then recombined, with the interference pattern measured by the photodetector.

1.4.2 Noise Sources

Due to their high sensitivity interferometers are also very susceptible to noise. In order to increase the sensitivity of the detector you either need to be able to eliminate these sources or accurately quantify them.

Seismic noise

As all current detectors are placed on the Earth's surface, they are subject to terrestrial vibrations. This is known as seismic noise and is the biggest noise source in the detector. It can range from earthquakes to trucks driving past the facility, and even down to birds pecking at ice built up on the pipes outside [44]. This noise is accounted for by complex damping systems, which fall into two categories; passive and active.

Active damping systems respond to local vibrations and neutralise them. They sit atop the passive systems, which attempt to isolate the systems from any noise sources by keeping them as still as possible. The suspension of the test masses from several[2] vertically connected pendula is part of this noise cancelling system. See section 1.5 for more information on mirror systems.

[2]LIGO has four pendulum stages, whereas Virgo has seven.

Thermal noise

Thermal noise is the second biggest noise contributor and is, therefore, one of the biggest limits on sensitivity. This is because thermal noise occurs in the frequency band in which the interferometer is most sensitive. There are two types of thermal noise; Brownian and thermoelastic. They both occur in the mirror coatings of the test masses [45]. Thermal noise also occurs in the suspension fibres and is a significant noise source at low frequencies [46].

Brownian noise is due to the random thermal motion of atoms and thermoelastic noise is due to the random statistical fluctuations in temperature. Brownian noise is difficult to measure directly but relates to a much easier to measure property called mechanical loss, which is caused by internal friction of the material. The lower the mechanical loss of the material is, the lower the Brownian noise contribution will be. As such, a significant amount of research goes into developing coating materials with very low internal friction [47].

Thermal noise will be discussed further in section 1.5.2

Quantum noise

Sometimes also called photon noise, quantum noise occurs due to the quantum nature of light. There are statistical fluctuations in the pressure exerted by the photons on the mirrors and in the number of detected photons at the detector [48], the latter being known as shot noise.

Quantum noise can mitigated in two main ways; alternating the set up of the detector and using a technique called squeezing [49].

Gravitational gradient noise

Gravitational gradient noise arises due to fluctuations in the local gravitational field, which can be caused by all manner of things; for example an aeroplane flying overhead that exerts some gravitational pull on the test mass. Its biggest source is actually seismic noise – seismic waves are pressure waves, causing density fluctuations in the Earth's crust. As the density of the Earth's crust below the detector changes, so does the gravitational force it exerts on the test mass. This form of the noise is possible to isolate, as described above, but it can persist through the gravitational coupling between the local surface of the Earth and the test mass. Other forms of gravitational gradient noise are not currently possible to isolate, but it can be minimised by carefully choosing detector location. KAGRA was placed underground, in an old mine, for this reason [22]. Placing a detector in space, as is planned with the LISA telescope [50], would eliminate this noise source.

1.4.3 LIGO

LIGO is made up of two interferometers: one in Hanford, Washington and the other in Livingston, Lousiana. They both have arms of 4 km, which is so long that the curvature of the Earth had to be accounted for during construction. Despite this, their length is not sufficiently long to detect GW. To overcome this, LIGO makes use of Fabry-Perot cavities. These are mirrors near the beam-splitter that reflect the beam back and forth approximately 280 times, so the actual distance travelled is over 1000 km.

The initial power of the laser is also too low, and alongside the Fabry-Perot cavities, more mirrors are used to increase the laser's power from 40 W up to 750 kW. These mirrors are known as power-recycling mirrors and they are one way mirrors that allow the light that has been transmitted back out through the beam splitter to be recycled. Signal recycling mirrors operate in a similar way once the beam is en-route to the photodetector [51].

The two LIGO locations were chosen to be isolated as possible from seismic noise sources, such as motorways and towns – much like traditional observatories are normally placed in areas with little light pollution. The two detectors are also on opposite sides of the United States, meaning that the seismic noise each detects will be independent of the other's. As the speed of GW and the separation of the detectors are known, the time difference of a GW passing through the two detectors is calculable. This difference can contain information about where the GW originated.

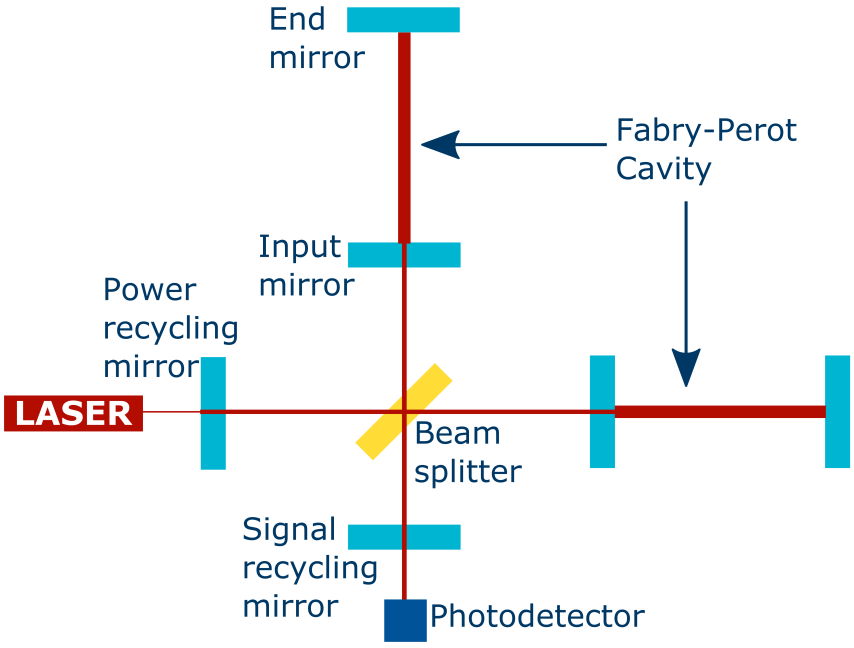

Figure 1.3: Simplified diagram of LIGO's additional mirrors, which increase both the laser's power and its sensitivity. The vast majority of the reflected beam passes back through the beam splitter to the power recycling mirror and is then 'recycled' back into the interferometer proper. There are actually multiple power recycling mirrors but for ease only one is shown. The Fabry-Perot cavities increase the effective length of the interferometer arms and build up the laser light within the interferometer, both of which increase LIGO's sensitivity. The signal recycling mirror boosts the signal received by the photodetector.

1.5 Mirrors and their Coatings

The LIGO mirrors are often refered to as test masses. They are made of silica, weigh approximately 40 kg and are hung at the bottom of a complicated suspension system. A simplified diagram of the mirror suspension system can be seen in fig. 1.4, showing the wires and multi-stage pendulum, although the scaffolding that envelops the system is not included in the diagram. The suspension system is a key component in reducing the seismic noise in the detector section 1.4.2. The mirrors' design means that they only absorb one in three million photons, to minimise the heating of the mirrors, which warps their shape and negatively affects the data quality by introducing thermal

noise. There are also other systems in place that counteract the heating effect of the beam.

In order to make the test masses into the highly reflective mirrors required for the interferometer to function, mirror coatings are applied to the front face. Like the test masses, the coatings have very low photon absorption and scatter losses; however, they also have a much higher mechanical loss than the test masses. Therefore, they contribute much more to thermal noise than the test masses [49].

The coatings are made up of multilayers of materials with a low refractive index and low dielectric constant, alternating with materials with a high refractive index (an image of the coatings taken by an electron microscope can be see in fig. 1.5). It is the difference in refractive index between the coatings that leads to their reflectivity. The first coatings for LIGO were undoped tantala and amorphous silica. Tantala had a history of being used as a coating in the optical industry due to its insulating and low-loss properties, and silica is frequently used for mechanical resonators with small dissipation [53].

Currently the mirrors are made up of 16 bilayers of amorphous silica (low refractive index) and titania-doped tantalum pentoxide (the high refractive index coating, often referred to as tantala) [54, 55, 56]. Doping with titania reduced the loss of the tantala layers by 40% [47] but despite this, the tantala layers are still the dominant source of loss [49].

Fully understanding the structure of these materials will allow us to lower the mechanical loss as much as possible. Research is also ongoing to find new materials that might be even better suited to mirror coatings. One material being considered as a possible replacement for the high refractive index layer is amorphous silicon, due to its very low mechanical loss properties. It would, however, require a change in the operating wavelength of the detector – at

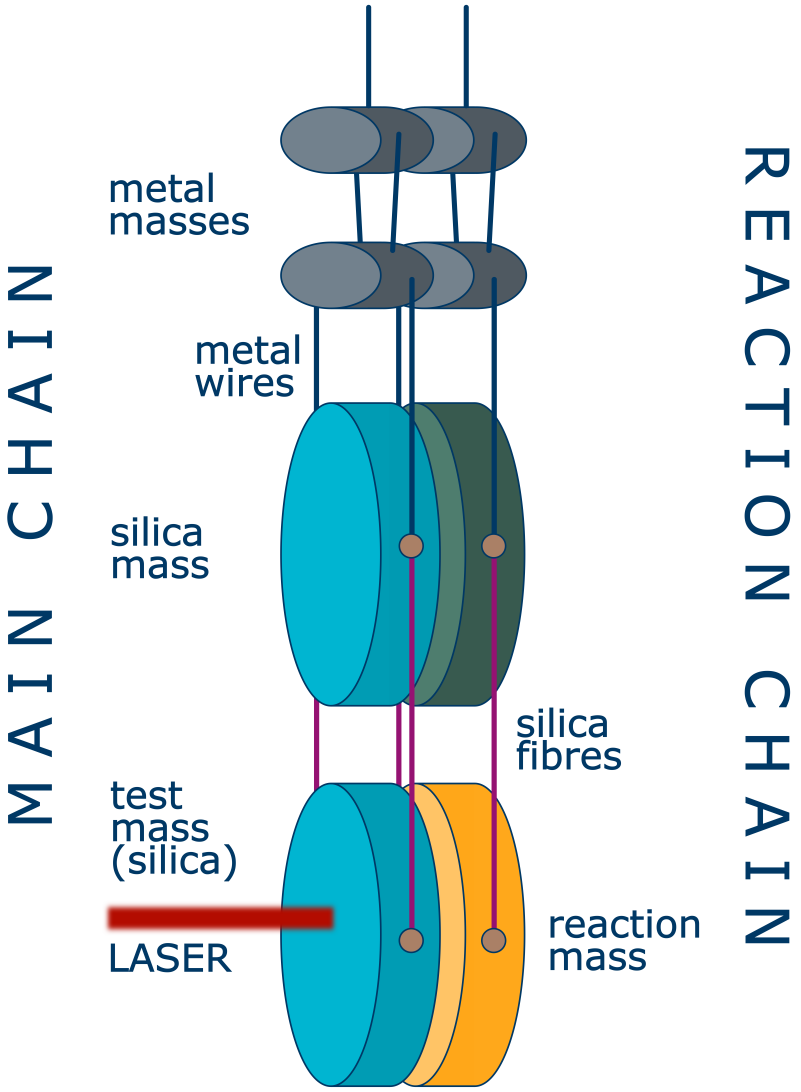

Figure 1.4: Simplified diagram of LIGO's mirror suspension system. The LIGO mirrors are the final bob of a quadruple pendulum. The first two pendula are made of metal with the second two being made of fused silica. The first three pendula are connected by metal wires, with the final test mass being suspended by silica fibres. There are springs within the upper levels of the pendulum that help to isolate the system from seismic vibrations, with the silica of the fibres and the test mass chosen to reduce thermal noise (see section 1.4.2). The first three pendula also contain magnets to assist with both passive and active damping. Note that this diagram is not anywhere close to scale; the silica test masses have a diameter of 340 mm and the silica fibres' varies between 400 and 800 μm [46]. The test mass weighs 40kg. Additionally, the mirror suspension system is surrounded by metal scaffolding not depicted here [52].

the currently used 1064 nm, amorphous silicon absorbs around 100% of the laser photons. Suggested alternate wavelengths include 1550 nm or 2000 nm, with the former having the added benefit of being the commonly used wavelength for telecommunications and so many highpower LASERs and optical components are readily available [57]. Amorphous silicon is already planned for use in cryogenic detectors [58]. Another candidate is a mixture of titania and germania because of the high degree of corner sharing between structural units, which results in low loss [59].

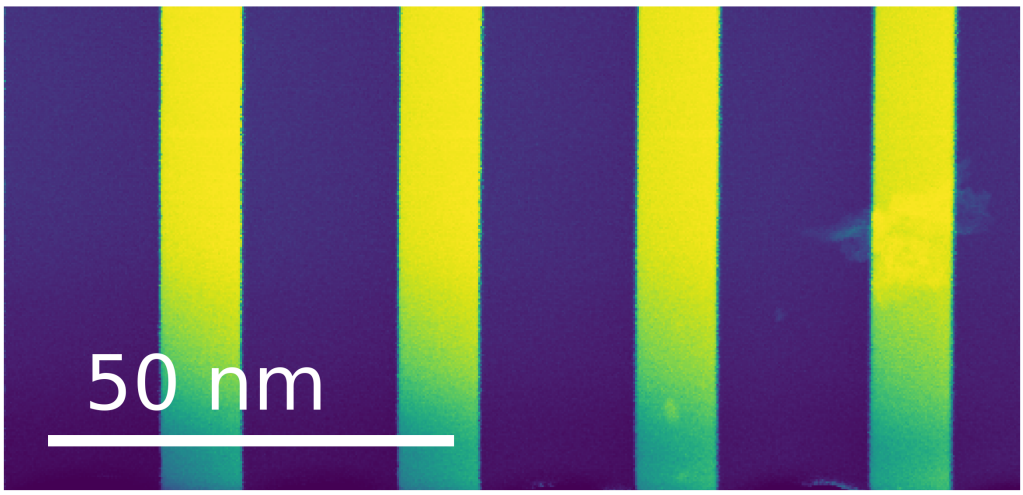

Figure 1.5: A HAADF-STEM (see section 2.2.2) image of a coating sample, showing the silica (dark, thicker) and tantala (brighter and thinner) layers.

1.5.1 Ion Beam Sputtering Deposition

Most GW mirror coatings are made using ion beam sputtering deposition (IBSD), including both the current aLIGO coatings and all of the GW coating samples analysed in this thesis. IBSD had its origins as a spacecraft thruster before it became widely used in thin film fabrication. Its first use for optical coatings was for a ring laser gyroscope [60]. The unparalleled quality of the coatings was such that the technique became the predominant optical coating fabrication method [61].

IBSD works by creating a high density plasma from which small ion beams are extracted and then accelerated through a grid to create a broad ion beam that strikes the target. The target atoms are then sputtered away and condense onto a substrate to form the thin films. A diagram of the setup can be seen in fig. 1.6 [62, 63]. In ion beam assisted deposition (IBAD) a second ion source of low energy is used. This is done to assist growth of the film without causing resputtering. The plasma gas is typically argon, although in IBAD oxygen (or a mixture of argon and oxygen) can be used to enhance oxidation in oxide films. Argon is used as it has a high number of sputters per incident ion and is relatively cheap compared with other noble gases, which have similarly high sputter yields [61]. It is possible for the plasma gas atoms to become trapped in the films during IBAD and form bubbles on annealing of the films. Chapter 4 investigates such argon bubbles in GW mirror coatings.

The biggest disadvantage to IBSD is that the filament (visible in fig. 1.6) requires frequent maintenance. This can be overcome by generating plasma in other ways but they are often prohibitively expensive. Another disadvantage is how slow the technique is – on average only a few tens of Ångstroms are deposited per minute. There are multiple advantages to IBSD, most notably

the high degree of control over the deposition parameters and how simple it is to create multilayer coatings, which can be done simply by spinning the target. Another benefit that is particularly relevant to GW coatings is tht the technique provides films with much lower optical scattering than any other deposition method thanks to the high energy of the deposition process. The argon ions give the target atoms high kinetic energy, resulting in a dense deposited film. Other techniques such as electron beam evaporation are lower energy and the deposited films are consequently much more likely to have voids, which can cause optical scattering.

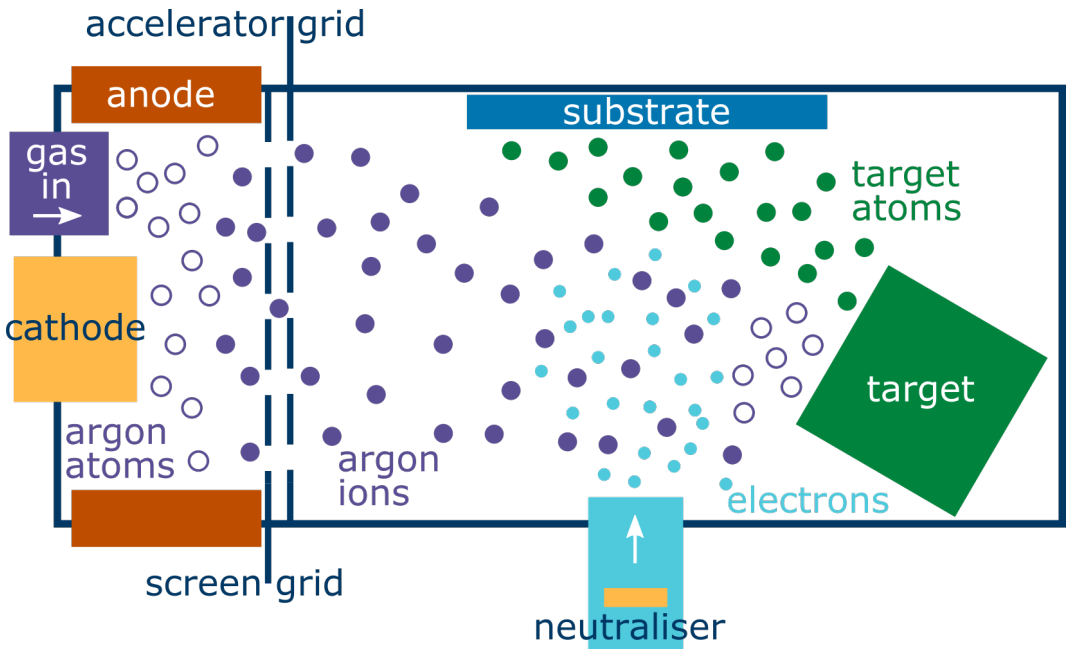

Figure 1.6: A ion beam sputtering set-up, showing the hot filament (cathode) and the extraction grids through which the argon ions are accelerated towards the target. The target can be attached to the target holder, which can then be rotated to create multilayer samples. The neutraliser emits electrons in order to neutralise the argon ions.

1.5.2 Noise Sources in the Coatings

One of the biggest sources of noise within the detectors is thermal noise, which can be split into two components; Brownian and thermoelastic noise. Thermoelastic noise includes thermorefractive noise. The thermal noise in any given system corresponds to the number of degrees of freedom the system has, where the noise is equal to

$$\frac{n}{2}k_B T, \tag{1.11}$$

where, n is the number of degrees of freedom, k_B is the Boltzmann constant and T is temperature. The magnitude of this noise for a gravitational wave detector is likely to be the same magnitude as a gravitational wave and so, the thermal noise of the detectors is a severe limiting factor on their sensitivity. In aLIGO, the thermal noise restricts the detector between ∼10-100 Hz. Thermal noise most commonly arises in the mirror substrates, coatings and suspension systems.

1.5.2.1 Brownian Motion

Brownian motion was first observed in 1828 by Robert Brown [64]. He observed the random motion of dust and pollen in water and it was later proven by Einstein that their random motion was due to the stochastic collisions between the debris and water molecules. For atoms or molecules, the thermal energy of the system can also be described by eq. (1.11) and their thermal motion can be more fully described by the fluctuation dissipation theorem.

1.5.2.2 The Fluctuation Dissipation Theorem

The fluctuation dissipation theorem was built on Einstein's initial work and was developed by Callen and Welton in 1951 [65]. It relates the fluctuations – or excitations – of a system with how these fluctuations are dissipated – or dampened – by friction. The mobility ($M(\omega)$) of the system relates to the force applied to the system ($F(\omega)$) and its resulting velocity ($v(\omega)$), which is dependent on the angular frequency (ω)

$$M(\omega) = \frac{v(\omega)}{F(\omega)}. \tag{1.12}$$

This can also be shown using the power spectral density $S_x(\omega)$ of the thermal noise

$$S_x(\omega) = \frac{4k_B T}{\omega^2} \Re\{M(\omega)\}. \tag{1.13}$$

Therefore, the thermal noise can be calculated if the mobility of the system is known; in GW detectors it arises due to the mechanical loss of the mirror coatings. Fortunately, the mechanical loss is a lot easier to measure than Brownian noise.

1.5.2.3 Mechanical Loss

Mechanical loss is also known as the mechanical loss angle, and arises due to a delayed response to stress. When materials are subject to stress they do not react immediately, and the strain builds up over time. In an ideal material, stress θ and strain ϵ are related by Young's modulus (Y) and Hooke's law ($\theta = \epsilon Y$). If we apply a more complex, periodic strain to a material, it is possible to derive an equation for the strain that involves mechanical loss

$$\epsilon = \epsilon_0 e^{i(\omega - \phi)}, \tag{1.14}$$

where ω is the angular frequency of the oscillation caused by the stress and ϕ is the angle by which the strain lags the stress. It is for this reason that mechanical loss is sometimes called the (mechanical) loss angle. The detector mirrors can be modelled as harmonic oscillators with internal damping. Combining eq. (1.13) and eq. (1.14), it is possible to derive an equation for the power spectral density of thermal noise in terms of the mechanical loss ϕ,

$$S_x(\omega) = \frac{4k_B T}{\omega} \frac{{\omega_0}^2 \phi(\omega)}{m(\phi^2(\omega){\omega_0}^4 + ({\omega_0}^2 - \omega^2)^2)}. \tag{1.15}$$

Noise in the mirrors comes from both the coatings and the substrate. It is possible to use the preceeding equations, and create an equation that describes the total mechanical loss in terms of the mechanical properties of the materials that make up the mirrors. We start with an alternative equations for the mechanical loss

$$\phi = \frac{1}{2\pi} \frac{E_{\text{dissipated}}}{E_{\text{stored}}}, \tag{1.16}$$

where ϕ is the mechanical loss and E is the energy stored and lost in each oscillation. After a derivation that can be found in full in Harry et al. [66] the resulting equation is

$$S_x^{total}(f) = \frac{2k_BT}{\sqrt{\pi^3}} \frac{1-\theta^2}{\omega_0 Y} (\phi_{substrate} + \frac{2}{\sqrt{\pi}} \frac{(1-2\theta)}{1-\theta} \frac{d}{\omega_0} \phi_{coating}). \tag{1.17}$$

This shows that if a material with a low mechanical loss angle is chosen then the thermal noise contribution will be lowered. What it does not explain, however, is where the mechanical loss arises. The lag of the strain can have many causes, including point defects, dislocations and grain boundaries [67].

In the mirrors, the substrate noise contribution is lower than the coating noise by around five orders of magnitude [68, 69], and the coatings are also intrinsically more important due to their proximity to the beam [70]. However, as we reach the limits of reducing the noise in the coatings, reducing the contribution from the substrate becomes increasingly important.

1.5.2.4 Scatter Loss

LIGO's coatings had many point defects after being installed in the detector. These resulted in scatter and localised absorption of the laser photons [71], which was a problem as scattered photons can then reflect from random surfaces within the detector that are not seismically isolated. In doing so, they pick up a lot of vibration noise, potentially introducing that back into the laser beam [72]. Additionally, too much scatter leads to a loss of light from the arm cavities, which means it will not be possible to reach the desired power buildup (section 1.4.3). Manufacturing coatings without these defects is a major challenge for the future.

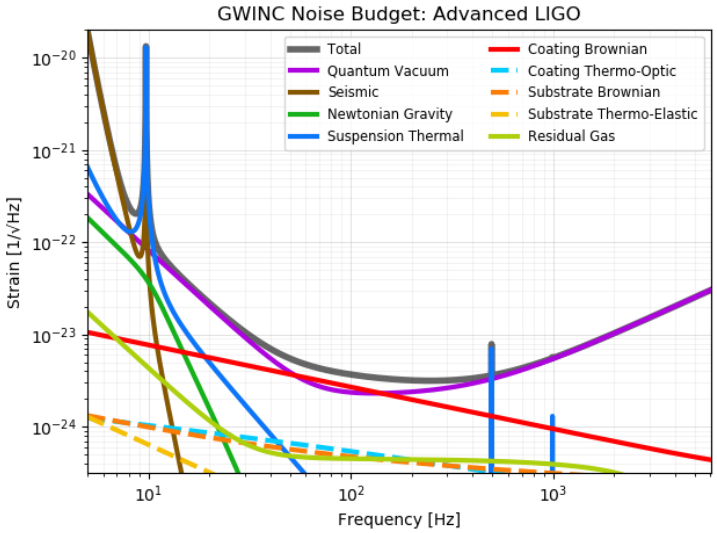

Figure 1.7: The calculated noise contributions for various noise sources in aLIGO. The coating thermal noise is largest at low frequencies but is dwarfed by all the other contributions. It is only at the peak sensitivity that it is a crucial contributor. This figure was taken from the Gravitational Wave Interferometer Noise Calculator [73].

1.5.2.5 Sensitivity Limits

The frequency at which the detector is most sensitive is a function of all of the contributing noise sources, a breakdown of which can be seen in fig. 1.7. In LIGO, most sensitive frequency is around 100 Hz, between the shot noise that dominates at higher frequencies and the radiation pressure noise that dominates the lower ones. It is at this point that the quantum noise and the coating thermal noise are approximately the same amplitude, and are the two biggest contributors to the total noise. It is in this frequency range that most signals have been observed – the first ever detection swept between 30 and 150 Hz. It was an inspiral, where the GW increased in frequency during the last few co-orbits of the BH, lasting about 0.1 s [6].

If future detectors are to reach their sensitivity goals then both quantum and coating thermal noise must be reduced. While there are many routes for

reducing quantum noise, the coating thermal noise remains a serious problem, as it has for many years.

1.6 Summary

This chapter has provided a brief overview of GW, starting with the history and theory of their detection, followed by how they are produced. It then discussed GW detectors, looking at their design before discussing the sources of noise in GW detectors. The final section was on the mirrors of GW detectors, focussing specifically on the mirror coatings. Sources of noise within the coatings, as well as their significance, were discussed.

Reducing the thermal noise of the coatings is crucial for future GW detectors. Currently, this noise is one of the biggest limits on sensitivity in the detector's most sensitive region. Fully understanding the structure of the coatings is one step in the process of reducing their thermal noise. Characterising the coatings is the focus of this thesis.

2

Microscopy

As stated in the previous chapter, understanding the structure of the materials used in LIGO's mirrors is crucial to fully characterising and eliminating their noise contribution. Electron microscopy is one way to examine the atomic structure of the mirror coatings. This chapter describes the make-up of a transmission electron microscope, in addition to covering the fundamentals of transmission electron microscopy (TEM), scanning transmission electron microscopy (STEM)[1] and electron-matter interactions.

Microscopy traces its origins back to the 16th century[2]. The inventor of the first microscope is in dispute, although it is sometimes attributed to Galileo Galilei, who adjusted the lenses in his famous telescope to magnify

[1] TEM and STEM can be used to interchangeably to refer both the technique (microscopy) and the instrument (microscope).

[2] Primitive forms of magnification, such as using water and gemstones, date back much further and the first lenses are believed to have been created some time after the 11th century [74]

terrestial, rather than celestial objects [75]. These early microscopes were optical microscopes, that is, they used visible light to magnify their subject. Unfortunately, this means they are limited by the wavelength of visible light (∼400nm to ∼700nm). This limitation was first recognised by Ernst Abbe in 1873 [76], with an equation now known as the Abbe diffraction limit

$$d = \frac{\lambda}{2\ NA} = \frac{\lambda}{2\ n \sin(\theta)}, \tag{2.1}$$

where d is the minimum resolvable distance, λ is the wavelength of light and NA is the numerical aperture – a dimensionless number that represents the range of angles over which the lens can absorb/emit light. NA can be replaced with $n \sin(\theta)$, where n is refraction index and θ is the half-angle of the convergence spot. Even at the limit of modern optics, using visible light results in a limit on the order of hundreds of nanometres [77] for conventional microscopes. This obviously poses problems for any attempts at atomic-resolution imaging, and even for imaging smaller biological samples such as viruses [78]. Using non-conventional microscopes, however, it is possible to overcome many of these limitations using methods such as scanning near-field optical microscopy [79] and super-resolved fluorescence microscopy [80, 81]. Near-field optical microscopy gives a resolution of 50-100 nm [79]. In fluorescence microscopy fluorescent tags are placed on molecules and can give a resolution on the order of tens of nanometres [80].

The exact 'beginning' of any discipline or technique is often difficult to definitively pinpoint; usually there are a collection of simultaneous or sequential but unrelated discoveries that are all later realised to be crucial. For electron microscopy the first of these discoveries is arguably de Broglie's proposal of

wave-particle duality [82, 83, 84]. The de Broglie wavelength of an electron (λ_e) relates to the voltage (V) through which it is accelerated

$$\lambda_e = \frac{h}{\sqrt{2\, m_e\, e\, V}}, \tag{2.2}$$

where h is the Planck constant, m_e is the rest mass of an electron and e is the charge of an electron. It holds for cases where the accelerating voltage is low enough that the electrons are not accelerated beyond relativistic speeds (approximately $0.1c$). Because modern microscopes often operate at accelerating voltages of hundreds of thousands of volts, relativistic effects need to be considered when calculating the wavelength of the electrons

$$\lambda_{rel.\,e} = \frac{h}{\sqrt{2\, m_e\, e\, V\left(1 + \frac{e\, V}{2\, m_e\, c^2}\right)}}, \tag{2.3}$$

where c is the speed of light. Assuming the microscope is operating at 200kV the wavelength of the electrons is 2.511pm, which would give greater than atomic resolution. Due to other limitations of the microscope however, the resolution of the microscope is not this low. Nonetheless, with aberration correction (see section 2.2.1), it is possible to get resolution of 0.5Å.

The first electromagnetic lens was created in 1926 by Busch [85] and five years later, Knoll and Ruska built the first electron microscope. Their invention would win Ruska the Nobel Prize in 1986 (Knoll died in 1969 and so was unable to share the prize), alongside Binnig and Rohrer for designing the first scanning tunnelling microscope [86]. While the design of the electron-microscope was initially an exercise in proof of concept, by the end of the 1930s electron microscopes were being manufactured by Siemens [87].

2.1 Transmission Electron Microscopy

2.1.1 The transmission electron microscope

All of the data collection in this thesis was done using a JEOL ARM 200F, unless stated otherwise. Henceforth, it will be referred to as the ARM (atomic resolution microscope) and it is a TEM with STEM capability. A simplified diagram of the ARM can be seen in fig. 2.1, showing the condenser, objective and projector systems, as well as the STEM detector. It has ray diagrams for both STEM and TEM, with the former being the set-up used for the work in this thesis. The following sections describe the different components of the microscope.

Emission source

An electron source is arguably the most crucial component of of electron microscope. The ARM uses a cold[3] field emission gun, which is made of a sharp tungsten tip followed by an extraction anode, with an electric field applied between the two that causes the electrons to tunnel across the vacuum. The tip of the filament is very small, which is one of the reasons for using tungsten – a strong material that is easily sharpened to a fine point. Tungsten also has a low work function, meaning it is easy to liberate electrons from it.

Despite operating in a vacuum to avoid contamination some gas molecules still finds their way into the microscope column. The gun's applied field is more than strong enough to ionise these gas molecules, which are then drawn towards the gun, coating the tip and thereby increasing its work function and suppressing the extraction of electrons. To help prevent this build-up the tip

[3]The cold moniker comes from the fact it is operated at room temperature, instead of the high temperatures for the other most common type of electron source, known as a thermionic emitter or Schottky field emission gun.

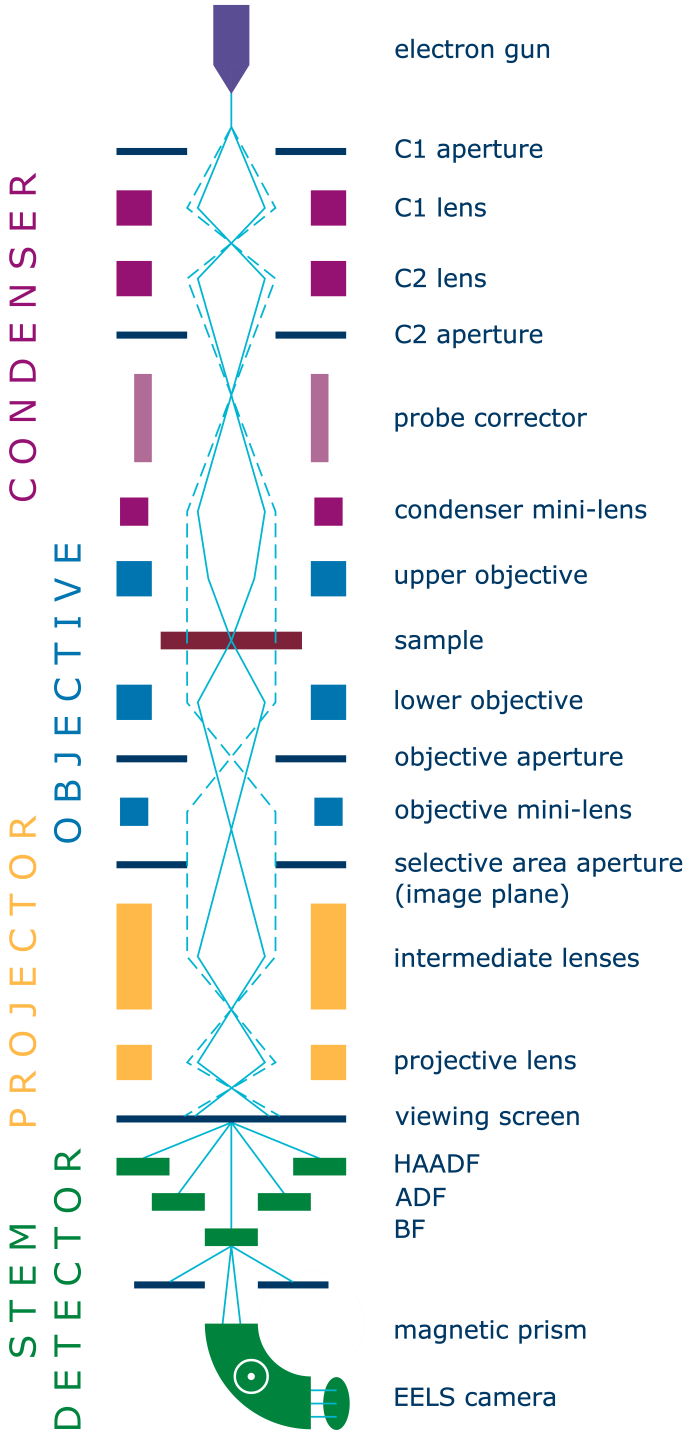

Figure 2.1: A simplified schematic of an electron microscope showing both a TEM in imaging mode (dashed line) and a STEM (solid line) set up. Note that the tilt and shift controls and the aberration correction (section 2.2.1) are not included in this diagram.

must be 'flashed' – briefly heated to a high temperature by increasing the current – a few times a day. A high flash is performed at the beginning of the day, with low flashes performed a few times throughout the day as necessary.

The energy of the beam E is dependent on the radius of the tip r and the applied voltage V through

$$E = \frac{V}{r}. \tag{2.4}$$

The ARM's normal operating voltage is 200 kV, with the tip being around 100 nm, although the electric field at the tip is dependent on the more modest extraction voltage. The narrow, focussed beam is of particular importance for electron energy loss spectroscopy (EELS, see Chapter 3). The ARM can also be set up with a voltage of 60, 80 and 200 kV. Other set-ups are also capable of 30 and 120 kV.

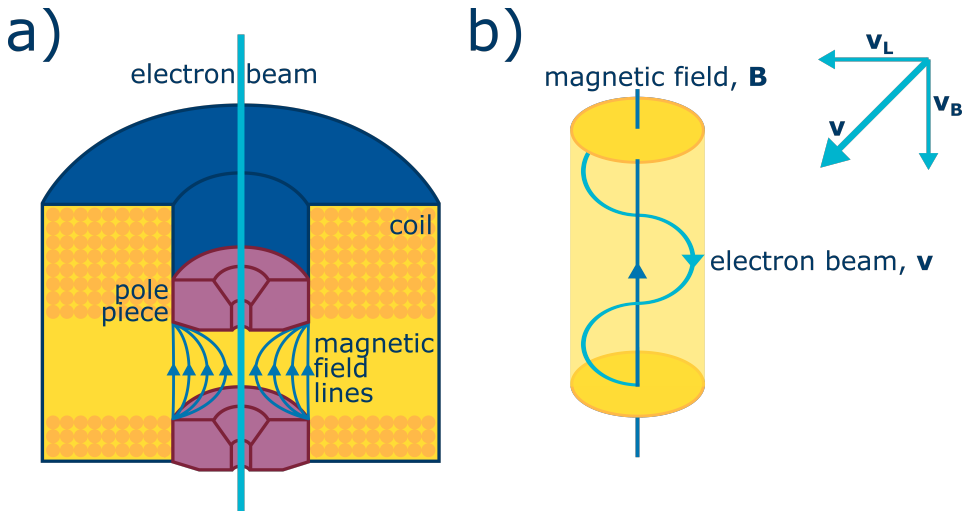

Figure 2.2: a) A simplified diagram of a magnetic lens, showing the pole pieces, the coil and the magnetic field created when a current is passed through the coils. It is important to note at this stage that although in this figure, and indeed throughout the figures in this chapter, the electron beam is drawn as a straight line, while going through the magnetic lenses it will take on a spiral trajectory, which can be seen in; b) the spiral trajectory of an electron, with velocity **v** through the gap between the pole pieces, with a panel showing the velocity vector.

Magnetic lenses

Like a conventional microscope, the TEM uses lenses to manipulate and focus the beam, the difference being that it uses magnetic lenses rather than physical ones. Magnetic lenses work because the electrons have charge and can therefore be deflected by a magnetic field. While analogous to conventional glass convex lenses, they are made up of copper coils and two pole pieces. A diagram of a magnetic lens can be seen in fig. 2.2. The pole pieces are generally made from a soft (in the magnetic not mechanical sense) iron. Sometimes they are connected and sometimes they are two separate pieces. When a current is passed through the copper coils surrounding the pole pieces a magnetic field in between the pole pieces, is generated, which can be described using Ampere's law. The resulting force experienced by the electrons in the beam can be described using

$$\mathbf{F} = -e\mathbf{v} \times \mathbf{B}. \tag{2.5}$$

The beam experiences a magnetic force that can be split into two vector components; one parallel to its path ($\mathbf{H}_B$) and the other in the plane of the lens ($\mathbf{H}_L$). $\mathbf{H}_L$ first causes the path of the beam to spiral, which exposes it to $\mathbf{H}_B$, which also affects its path, thereby focussing it. The path of the beam and its focal point can therefore be controlled by the strength of the current passed through coils. Note that this is a classical explanation of the Lorentz force, a quantum mechanical treatment can also be given, which gives a more complete picture [88].

Condenser lens

The condenser lenses control the illumination of the microscope. There are two condenser lenses in the ARM; C1 and C2.

C1 controls the 'cross-over' point and therefore controls the size of the beam after C2. The 'cross-over' point can be seen in fig. 2.1, between the C1 and C2 lenses and is the point where the beam electrons converge. The stronger C1 is, the more it will demagnify the beam; increasing its strength increases the beam current that can pass through. The C2 lens determines how convergent the beam will be; a stronger C2 will result in parallel or broad-beam, illumination; and a weaker C2 in a convergent beam or probe illumination. Only a weak C2 is required for STEM, to maintain the focus for rastering across the sample. The ARM also utilises another pole piece that effectively acts a third, 'mini' condenser lens, C3, which enables better resolution as it further reduces the probe size by increasing the convergence angle [76].

Above and below the condenser systems there are shift and tilt controls, which are two $x - y$ pairs of coils with magnetic fields perpendicular to the beam. The upper set aligns the beam with C1 and the lower set aligns it with the objective lens. Between C2 and C3, there is an aberration corrector. Aberration correction is described in more detail in section 2.2.1.

Objective lens

The sample sits within the objective lens system, which magnifies the image once the beam has passed through the sample. It refocuses the beam after it has been deflected by the sample. To produce an image, the lenses must be adjusted such that the object plane is the image plane of the objective lens. Diffraction mode will not be covered as it is not within the scope of this thesis, except in that the diffraction pattern of the sample is projected onto the entrance aperture of the spectrometer for EELS.

Sample stage

The sample stage allows four degrees of freedom when moving the sample; x, y, z and tilt. The sample itself is held within a rod that is inserted into the microscope. Another degree of freedom can be obtained by using a double tilt rod. There are also rods that allow more control of the sample itself, such as heating rods. The sample holder used for the data collection of this thesis was the beryllium[4] double tilt holder, although the double tilt capability was not extensively utilised as the default position was sufficient for STEM work.

Keeping the ARM stage steady is imperative when working at such great magnification; a drift of 10 nm could be disastrous for some experiments, and the ability to return to the previous position is also crucial when rastering across a sample. Fortunately, the ARM is very resistant to drift due to the design of its holder and goniometer, as well as the acoustic shielding of both the column and the sample rod. The ARM room is also carefully controlled; the temperature is kept constant (to 0.1 K) and there is no turbulent airflow. Properly set up, the drift can be brought down to a few nanometres per minute.

[4]Beryllium is used, despite being expensive, toxic and relatively soft, because its highest energy transition is relatively low in energy, making it useful for X-ray analysis

In the ARM, the stage is controlled electronically but in older designs the control was mechanical and depended on gears.

Post-sample lenses

After the sample and the remainder of the objective system there is the projector system, which is made up of several intermediate lenses and the projective lens. These further magnify the image and are what couple the main column to the detectors at the bottom of the microscope in STEM. There are several intermediate lenses in order to reach the required magnification, which can be approximately 500,000 times for high-resolution images. The projector system can be used for either creating a real-space image or for diffraction.

2.2 Scanning transmission electron microscopy

Scanning transmission electron microscopy is a technique that allows for sub-Ångstrom resolution. It is a subset of transmission electron microscopy that does not illuminate the entire sample simultaneously. Rather, a fine electron beam (0.5 to 2.5 nm) is rastered over the electron-transparent sample in a grid pattern. The image is collected by a dark field detector, which will be discussed in section 2.2.2. In this thesis, the images were collected pixel by pixel by a charge-coupled device (CCD), although this is not typical. The rastering is enabled by deflector coils that move the beam in a grid pattern. In typical STEM, the beam is a few nanometres when uncorrected. The corrected beam can be below 1 Å.

There are advantages and disadvantages to STEM. The biggest advantage is that it allows us to collect data from scattering mechanisms that cannot be recorded in conventional TEM. The main disadvantage comes from how

much slower STEM image acquisition is; because the beam rasters over the image and is collected pixel by pixel, the images are much more susceptible to drift. Additionally, since the beam is focussed on a sub-nanometre area, sample damage from the beam is more common when the beam is allowed to rest in a single location. This can, however, be avoided by using a low beam current and scanning quickly.

2.2.1 Aberration correction

Resolution within the microscope can be severly limited by aberrations of different kinds and for most of the 20th century this was the case. Fortunately, resolution can be improved by aberration correction. '*In-situ*' aberration correction techniques had been suggested as far back as the 1930s but were beyond the technical capabilites of the time [89, 90]; early solutions relied solely on post-collection processing. In the 1990s, advancements were made by Haider et al. that allowed for spherical aberration correction. Since then, further improvements have allowed for resolution down to 0.5 nm, although instruments operated at 200 kV, such as the ARM, are still only around 0.7 Å at best. Aberrations affect not only resolution but can also distort the image.

Aberrations in magnetic lenses are comparable to those in glass lenses and come in three categories; geometric, chromatic or a combination of the two. In glass lenses geometric aberrations can be corrected with better grinding of the glass, however this is not the case for magnetic lenses; there are no diverging lenses for electrons. This problem is why it took 60 years for '*in-situ*' aberration correction to be implemented in the electron microscope.

Aberrations can be thought of in terms of two angles – θ, the angle away from the optical axis and ϕ. the azimuthal angle around the axis (i.e. up to 2π).

Different aberration effects go with the power of one or both of these angles, with higher-order aberrations having more complex symmetries. Astigmatism, for example, goes with θ^2 and has two-fold symmetry. It is therefore corrected using an octopole lens. Lower-order aberrations are generally corrected using a single multipole lense but higher-order aberrations may require multiple multipole lenses. It is the lower-order aberrations that have the greatest effect on the beam.

A secondary effect of aberration correction is the increase of the convergence angle without a loss of spatial resolution in both image and spectroscopy mode. This also shortens data acquistion times. The convergence semi-angle is increased from 10 mrad pre-aberration correction to 30 mrad through aberration correction. The semi-angle for the data collected in this thesis was 29 mrad throughout.

There are three main types of aberrations in STEM: spherical aberration, chromatic aberration and astigmatism. They are are detected using a Ronchigram, which is the convergent beam diffraction pattern. Its contrast shows defocus, astigmatism and other higher-order aberrations. Of the three main types, spherical aberrations and astigmatism are corrected in the ARM. All three will be discussed further below.

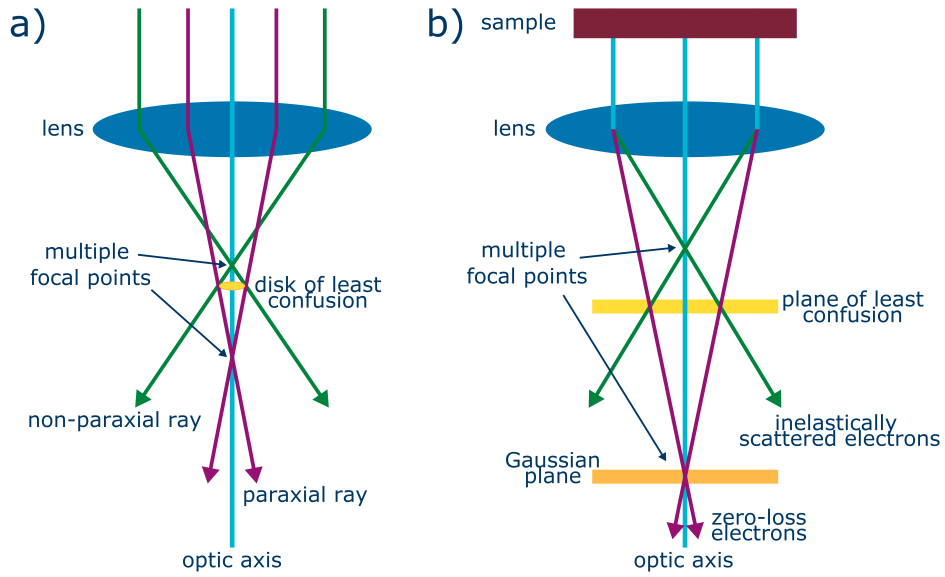

Figure 2.3: Aberration in STEM: a) spherical aberation showing how rays further from the optical axis are bent more strongly by the lens; b) chromatic aberration, showing how electrons with more energy are bent less strongly by the lens.

Spherical aberration

Spherical aberration is a gemotric abberation and occurs when the lens field affects the off-axis rays more strongly than the on-axis rays. The further off-axis the rays are, the more they are affected, i.e. they are bent more strongly back towards the axis. If the rays are too far from the axis, this effect leads to multiple focal points and decreases the resolution of the image. In that event, the image becomes a disk rather than a point. A ray-diagram of spherical aberration can be seen in fig. 2.3. Any spherical aberrations higher up the microscope column are only magnified by the following lenses. Spherical aberrations are lower order and therefore give rise to the largest effects in the ARM. As all aberration types are cascaded through subsequent lenses, spherical aberrations are the most problematic in the objective lens system.

All spherical aberrations of the cylindrical lenses in STEM (i.e. the main ones) are positive and so are corrected using a method that creates negative

spherical aberrations. In the ARM, this is done by a computer-controlled set of hexapoles and several thin lenses, but can also be achieved using a quadrupole-octupole design. The first set of hexapoles corrects for the primary fifth-order aberrations, with the second hexapoles correcting for secondary fifth-order aberrations. The two hexapoles sets also create a cascade of higher-order spherical aberrations that must then also be corrected. In the ARM, the CEOS C_s aberration corrector is placed in the upper column in order to correct the lenses used in probe formation and to produce the smallest and most aberration-free probe possible.

Chromatic aberration

Chromatic aberration occurs because not all of the electrons in the beam have exactly the same energy, which means they will have different focal points, with lower energy electrons focused more strongly. This results in all points being slightly blurred and is of particular concern for fine image features at high spatial frequency. A ray-diagram of chromatic aberration can be seen in fig. 2.3.

Chromatic aberration is a smaller issue in the ARM than spherical aberrations is not corrected. It is a small effect because the FEG has a very low energy spread relative to its acceleration voltage ($\leq$ 1 eV versus 200 keV) but it can cause problems when operating the ARM at a lower voltage as the energy spread is proportionally greater. Chromatic effects post-specimen have not been studied extensively but they are more of a problem post-sample as the beam has undergone inelastic scattering, leading to a polychromatic beam [91]. This also means they are a bigger issue in the post-specimen lenses for thicker samples.

Astigmatism

Astigmatism is a geometric aberration and is caused in two ways, the most common being when the lenses' magnetic fields are not perfectly cylindrically symmetric with the beam down the axis. The second cause is contamination on apertures. These non-uniformities cause both quadropolar distortions at 45°, and change the beam from circular to elliptical, distorting the image. Astigmatism is detected by 'streaking' in the image in one direction that changes to the perpendicular direction as the image passes from under-focus through to over-focus. It is, however, relatively easy to correct using the stigmators. These are octopole lenses that produce a field that counteracts the imperfections in the beam caused by the earlier magnetic lenses and in the beam itself. They trim along two quadropole axes at 45 ° to one another. There are three sets of stigmators; one for the objective lenses, one for condenser lenses and one for diffraction.

2.2.2 Bright and dark field imaging

In STEM-BF, the beam rasters across the sample, rather than illuminating the whole sample simultaneously as is the case in TEM, as can be seen in fig. 2.1. The BF detector lies at the bottom of the optical axis, although when collecting EELS data the BF detector is replaced by the spectrometer.

Dark field (DF) images are those formed when the beam has been scattered off the optical axis. DF imaging is different between STEM and TEM; in TEM the beam can be brought back to parallel with the optical axis using pre-specimen deflector coils. The objective aperture can then allow for only the deflected beam to carry on down the column. In STEM, no aperture is used, both the direct beam and the scattered electrons pass down to the detectors and can instead be differentiated using annular detectors[5] (see fig. 2.4). There are two types of DF detectors; annular dark field (ADF) and high angle annular dark field (HAADF). Their exact collection angles (β) vary depending on the set-up of the microscope (see table 2.1) but as suggested from the name, the ADF collects less strongly deflected rays and the HAADF the more strongly deflected. It is possible for the HAADF detector not to collect truly 'high' angles as it so dependent on the camera length (the distance from the specimen to the recording plane), the physical position of the detectors, and their inner and outer radii.

Detector	Range	$\beta = 29$ mrad
EELS Signal	0-β	0-29
ADF	β - 2β	29 - 58
HAADF	2β-7β	58-203

Table 2.1: Collection angles for the detectors based on the set-up used for the data collection in this thesis, where $\beta = 29$ mrad.

[5]Instead of the annular detectors described here, 4DSTEM (i.e. picking up the diffraction pattern with a pixelated direct electron detector) is becoming increasingly popular [92, 93].

The ADF detectors themselves are made up of scintillator-photomultipliers (p-n diodes can also be used), converting the captured electrons into a readout signal. In the ARM there are four possible DF detectors, although only two can be inserted at any time. When performing EELS, the Gatan detectors are used but when doing pure imaging, the ARM's own detectors (ADF1/2) can be used. These allow for longer camera lengths and a greater dynamic range (i.e. increasing the difference in intensity the camera can record simultaneously), resulting in better images.

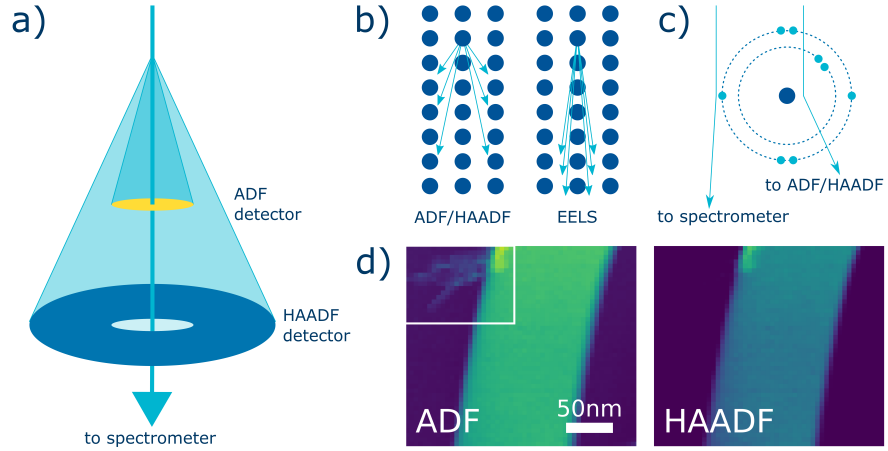

Figure 2.4: a) A simplified representation of the ADF and HAADF detectors, showing their ring shape around the optical axis; b) scattering interactions showing paths taken by the rays used for DF imaging and for EELS Chapter 3; c) a closer look at (b) using an individual atom; d) an example of the differences between ADF and HAADF images for the same sample, showing contamination detail in the upper left-hand corner on the ADF image not visible on the HAADF image. This sample is an aLIGO multilayer coatings sample, the brighter area being titania-doped tantala and the darker areas either side are silica.

Contrast in DF images comes primarily from Z-contrast [94], i.e. the atomic number of the sample atom with which the beam interacted. The Z-contrast is roughly proportional to Z^2 and occurs due to rays that are incoherently scattered, primarily through Rutherford scattering [95]. At lower scattering angles, diffraction contrast dominates. This scattering is still Rutherford-like but due to its low angle does not have the ideal Rutherford Z^2 dependence.

2.2.3 The spectrometer

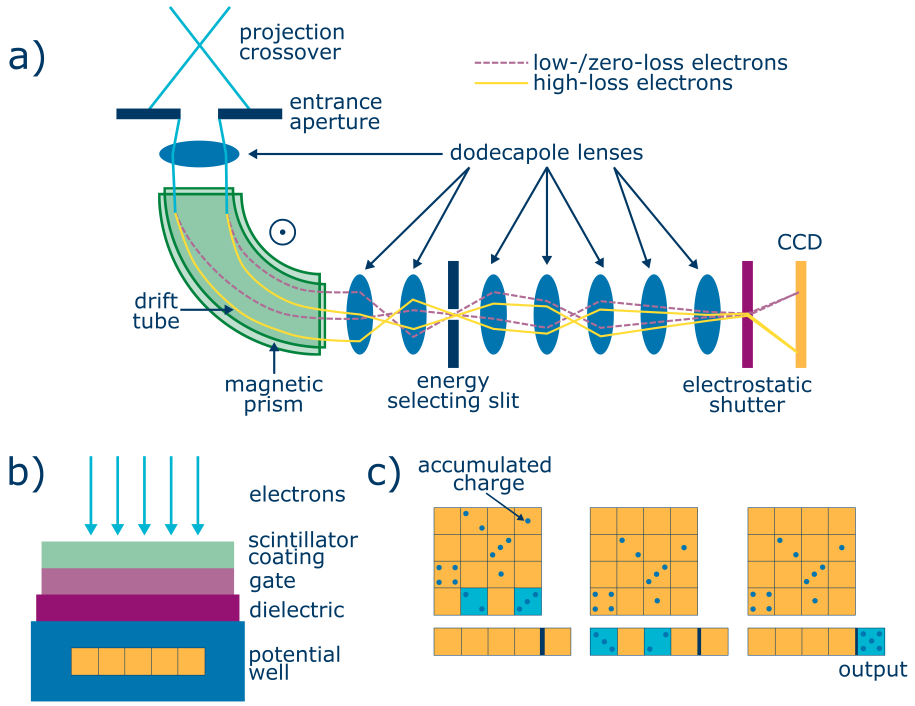

Figure 2.5: a) a diagram of the interior of the spectrometer, b) a single cell of the CCD, showing the arrangement of the layers and the potential well; c) an expansion of the potential well in (b), showing how the accumulated charge is read out.

The spectrometer lies at the very end of the optical axis in the microscope, replacing the BF detector when performing EELS. It has three main components; the magnetic prism, the internal lenses and the CCD.

The beam passes through a first dodecapole lens and then into the drift tube within the magnetic prism. The magnetic prism sorts the electrons based on their kinetic energy [96], splitting the beam into a spectrum. A magnetic field is applied perpendicular to the beam so that the beam is bent by approximately 90°. The slower the electrons are travelling (i.e. the more energy they have lost) the more strongly they will be affected by the Lorentz force

$$r = \left(1 - \sqrt{1 - \left(\frac{v}{c}\right)^2}\right) \frac{m_e\, v}{eB}, \tag{2.6}$$

where r is the radius of curvature, v is velocity, c is the speed of light, m_e is the rest mass of an electron, e is the charge of an electron and B is the strength of the applied field. The current inside the prism is fixed in order to avoid thermal loading of the electro-magnet, which would cause drift. As a result, the post-prism electron beam dispersion is also fixed.

The three dodecapole lenses between the entrance aperture and the energy-selecting slit are used for focusing and aberration correction of the beam/spectrum, bringing the spectrum into focus at the plane of the slit. The post-prism lenses exist to magnify the narrow dispersion caused by the fixed current. They also help correct chromatic aberration (see section 2.2.1), as well as focussing, magnifying and projecting the spectrum onto the CCD. The energy selecting slit is not often used for EELS. For EELS, it is more important to correct geometric aberrations so that the spectrum is sharp and well aligned.

The CCD represents perhaps one of the most significant advances in electron microscopy. Developed in the 1970s, they were first used in electron microscopes in 1982 [97] but widespread usage did not occur until the early 2000s, with film remaining the preferred collection method. Despite the slow uptake, CCDs offer many advantages over film, with the main advantage being the direct readout of data without degrading resolution. Additionally, their electronic nature also allows for the collection of gigabytes of data in a very short period of time.

CCDs are made up of an array of pixels that store charge of incident electrons until the integration period is over. At this point they offload into the output at the edge of the array, this vertical binning means that the readout is essentially one-dimensional, which increases the speed of the CCD. Once the offload is complete a new image can immediately be recorded. In a traditional CCD the signal is light, so it converts the signal into electrons and then stores their charge. In the microscope, as the beam is made up of

electrons, there is an intermediate stage where the electrons are converted to photons. This is done by a scintillator.

In DualEELS mode the CCD is halved; one section collects the high-loss signal and the other the low-loss section. This is done using an electrostatic shutter and a direct current offset switch (0-2 kV) in the drift tube to change the incoming energy range. DualEELS will be discussed in more detail in the next chapter (see section 3.2).

The CCD limits signal to noise ratio (SNR) in STEM and EELS through thermally generated counts, as well as noise that occurs in readout and amplification. When the signal is low, such as in the high-loss images, the SNR will be poor as the background noise is large proportionally to the real signal. The low count number also requires longer integration times, increasing likelihood of drift.

The specific spectrometer used for the work in this thesis is a Gatan Imaging Filter Quantum ER 965 (hereafter referred to as the GIF). Its high collection angles allow for the high resolution EELS analysis that makes up the work in this thesis.

2.3 Electron-matter interactions

Most of the electrons in the beam will pass through a thin sample without significantly interacting with it. Of those that do, there exist two main scattering processes; elastic scattering and inelastic scattering. These two types of scattering produce a variety of signals.

2.3.1 Elastic scattering

Elastic scattering is where the electrons lose no or a negligible amount of energy [98]. This happens in two ways; the beam electron passing by the edge of the atom, and Rutherford-like scattering. The first type is the most common and only causes weak deflection, between 10-100 mrad, as the electrostatic interaction with the shielded nucleus will be weakest at the edge of the atom. In gases the electrons are scattered by individual electrons, which can also be assumed for a first approximation for amorphous materials. In a solid however, the electrons act as waves and interference patterns characteristic of the atomic spacings appear, which is known as diffraction.

The second type of elastic scattering is less common but is a stronger effect as it occurs when the electron is closer to the nucleus of the atom. Here the electrostatic attraction is much greater and so the electron will be more strongly deflected and may even undergo backscattering, which is when the deflection angle is greater than 90°.

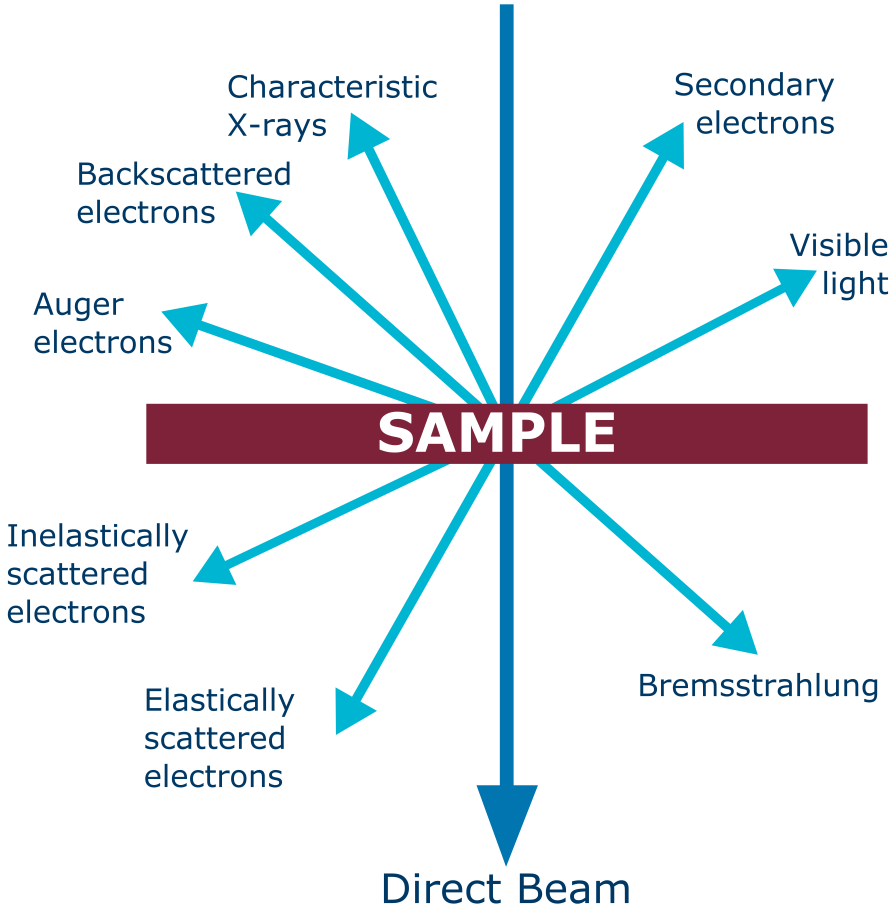

Figure 2.6: The different signals produced by the beam hitting the sample. A thin sample allows both forward- and backscattering. The different scattering types can be used for a variety of imaging techniques. Of particular interest for this thesis are the inelastically scattered electrons.

If the sample is crystalline this backscattering follows the Bragg diffraction law [99]

$$2dsin(\theta) = n\lambda, \tag{2.7}$$

where n is a positive integer, d is the distance between the crystal planes, 2θ is the scattering angle and λ is the wavelength of the electron (eq. (2.3)). If the microscope is in diffraction mode you will then see bright spots or discs that relate to the lattice constants.

If dealing with an amorphous sample, Bragg scattering can also be used to find information about the short-range order of the material, although in this case you will see diffuse rings rather than bright spots.

2.3.2 Inelastic scattering

Inelastic scattering is where the electrons lose energy. There are three main types of inelastic scattering, all occuring by the beam interacting with the sample electrons [98].

The first happens when the sample electrons in the inner shells of an atom absorb energy from the beam electrons and are excited to higher energy levels. This then deflects the beam electrons by a small angle, on the order of milli-radians, with an energy loss characteristic of the element of the atom with which they have interacted (see Chapter 3). The sample electrons will eventually de-excite, releasing the energy they gained in the form of radiation. This type of scattering is what allows for elemental analysis in EELS.

The second is the most common type of inelastic scattering and occurs when the beam electrons interact with the valence band electrons, or the free electron gas, causing collective, longitudinal oscillations of the free electrons. These collective oscillations are called plasmons. Plasmons are dampened in under a femtosecond and have a range of less than 10 nm. They occur primarily in metals but can be found in any material with free or weakly-bound electrons [98]. These two types of scattering, along with the unscattered beam make up the EELS spectrum, which will be described in more detail in section 3.1.

The third type happens when the beam electrons are decelerated due to an electrostatic interaction with the sample nuclei. The atom will release the energy it gains by characteristic radiation known as Bremmstrahlung.

There is also a fourth type of inelastic scattering known as phonon scattering. It involves only small changes in the energy of the beam electrons and happens due to the thermal oscillations of the sample atoms. The beam electrons can either gain or lose these small amounts of thermal energy. Although the energies involved are very small it can lead to large scattering angles as it involves a large momentum transfer as well as the energy transfer. As the energies are very small they are normally within the zero-loss peak (ZLP) on an EELS spectrum.

2.3.3 Plural Scattering

The thicker the sample is, the more likely a beam electron will be scattered more than once. If they are scattered too many times the information gained from the different types of scattering will be lost [96]. The mean free path (MFP) is the average distance between scattering events. Knowing the MFP for scattering in the material helps to know how thin our sample needs to be to prevent this; generally the sample needs to be less than half of the MFP.

The most common equation to calculate the thickness of the sample is the log-ratio method

$$\frac{t}{\lambda} = ln\left(\frac{I_t}{I_0}\right), \tag{2.8}$$

where t is the thickness of the sample at the beam location, λ is the MFP, I_t is area under the whole EELS spectrum and I_0 is the area under the ZLP. This technique requires already knowing the MFP so the relative thickness (t/λ) is more often used. In cases where the absolute thickness is required, $1/\lambda$ can be calculated using the semi-collection angle, beam energy and the density of the materials. This calculation was derived by Iakoubovskii et al. [100], with later adjustments found by Craven et al., and confirmed by Varambhia

et al. [101, 102], showing that the initial value found by Iakoubovskii needed to be reduced by 20%. As Digital Micrograph (DM) software calculates the relative thickness, it is trivial to multiply this data by the derived value of λ to get the actual thickness.

2.4 Summary

This chapter has provided a brief history of electron microscopy and the theoretical underpinnings of the field, as well as an outline of an electron microscope, specifically the JEOL ARM 200F used for almost all of the data collection in this thesis. It focussed in particular on STEM.

The ARM is made up of three main stages, plus the electron gun; the condenser, objective and projector systems. In STEM, there are also post-projector system detectors. The different STEM detectors are the annular dark-field detectors and the spectrometer. These detectors were discussed, as was aberration correction.

Understanding electron-matter interactions and the structure of the microscope – especially the spectrometer – is crucial for EELS, which is at the centre of all the work in this thesis. The next chapter will discuss EELS in detail, going into both technical detail on advances in the field how data is processed.

3

Electron Energy Loss Spectroscopy

Electron energy loss spectroscopy (EELS) is the measurement of the change in energy of the beam electrons after they have passed through the sample. It gives information on the chemical and physical structure of a material as the energy lost by the electrons is characteristic of the element with which they have interacted. EELS is the basis of all the work in this thesis.

Some of the work in reference [5] has been included in this chapter. Specifically; section 3.2, section 3.3 and section 3.4.1 are rewordings and expansions of work originally written for that paper.

3.1 The EELS spectrum

The EELS spectrum can be split up into two main areas; the low-loss region and the core- or high-loss region (core and high can be used interchangeably). The low-loss region has two main features; the ZLP and the plasmon peaks.

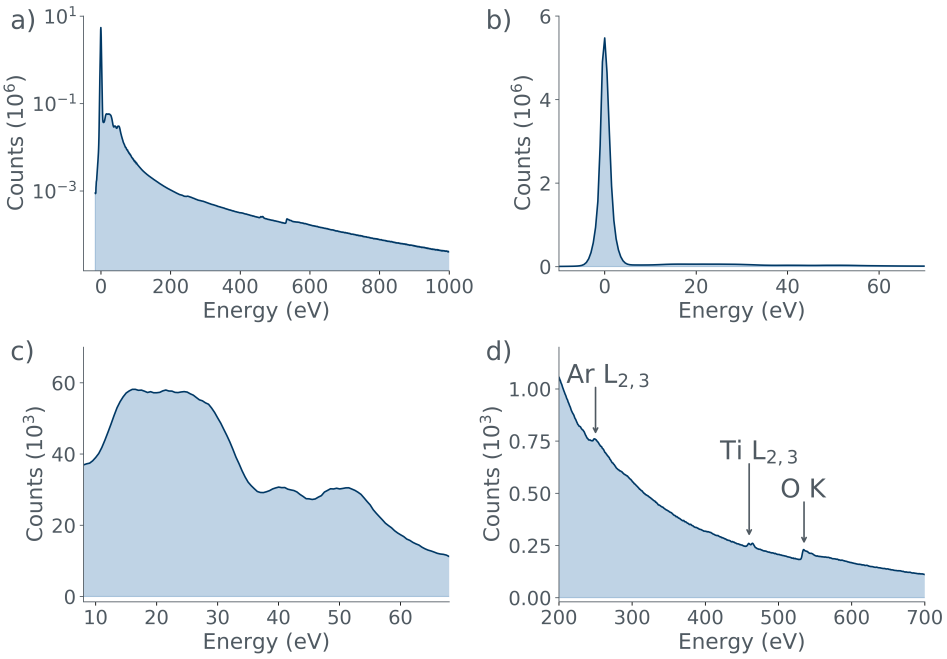

Figure 3.1: The different parts of an EELS spectrum; a) the whole EELS spectrum, dominated by the zero-loss peak; b) the low-loss region of the EELS spectrum c) the low-loss region sans the ZLP to show the shape of the plasmon peaks; d) the high-loss region, with the characteristic edges in this sample annotated. This sample is a LIGO titania doped tantala mirror coating. This and and all of the other graphs in this thesis were made using the Matplotlib library for python [103].

Most of the electrons passing through the sample do not interact with it and therefore do not lose energy. These zero-loss electrons make up the zero-loss peak (ZLP), alongside the multiply scattered electrons (section 2.3.3). Despite its name, the ZLP generally does not appear centred exactly around zero in the spectrum's energy scale. There are several causes for this; the most common are energy instability and drift in the microscope, not calibrating zero to the zeroth channel, scanning effects, and in very thick samples by multiply

scattered electrons. The ZLP being off-centre can be fixed via alignment but as it is normally only a few electron-Volts awry it is typically fixed during post-processing.

The ZLP is followed by the low loss region (generally up to 50 eV) in which the plasmon peaks are found. Plasmon peaks are caused by the beam interacting with the valence electrons of a material, i.e. electrons below the Fermi level. Interacting with multiple atoms' electrons simultaneously creates an oscillation of the valence electron density – a plasma resonance. This can be thought of as a pseudo-particle called a plasmon, the energy of which is proportional to the square root of the valence electron density. At any given moment the plasmon's energy is carried by one single electron, however when viewed over an extended period the energy is shared over many atoms, due to the collective nature of the phenomenon. Plasmons decay very quickly, into heat or secondary electrons. The strength of plasmon effects are dependent on the sample material, in metals they are strong as metallic electrons behave like free particles. While the plasmon peaks do carry information on the structure of the sample, study of plasmons is outwith the scope of this thesis. The intensity of both the plasmon peaks and the ZLP are dependent on the thickness of the sample.

The core or high-loss region is the area after the plasmon peaks. This is where the beam interacts with atoms in the material and excites the core electrons. It is perhaps better to use the term high-loss region as some core-loss features can appear in the low-loss region, although the majority are in the high-loss region. The electrons in the beam lose an amount of energy characteristic of the atom with which they have interacted and with its electron energy shells, producing the edges visible in fig. 3.1. After an initial gain in intensity the spectrum will decay back to the previous background levels of scattering plus the log term trend of the cross section for the new edge. The

intensity of an edge is proportional to how much of the atomic species there is within the sample; the more of it there is, the more electrons will interact with it and lose its characteristic energies.

The work in this thesis will focus almost solely on high-loss EELS, with the exception of using the ZLP for normalisation of the edges for quantitatisation.

3.2 DualEELS

One of the biggest problems in EELS is the large discrepancy in intensity between high- and low-loss regions. If recording on a single exposure, the much shorter shutter-time needed for the low-loss region, and its much higher intensity, will result in high-loss data with an extremely poor SNR. If using a more appropriate shutter-time for the high-loss region then the low-loss region will become over-exposed, as the ideal high-loss exposure time can sometimes be up to 4000 times the ideal low-loss exposure time. Collecting good data from the two regions does not require just a difference in exposure time but also a difference in beam conditions [104]. This is usually through a different spot size, although it can be done using condenser apertures and by changing the convergence angle. The different spot size method is generally preferred as the second method means deconvolution of the two spectra is no longer possible, and requires the microscope to be realigned between collections.

DualEELS overcomes this by taking two spectra simultaneously; one of the low-loss region and one of the high-loss. It was first described in 2008, having been developed in the decade prior [105]. DualEELS was made possible by the GIF Quantum, which was commercialised in 2010 [106]. There were three main advancements in the GIF from its predecessor that enabled DualEELS: changing the CCD shutter (which can be seen in fig. 2.5) from electromagnetic

to a much faster electrostatic mechanism; careful control of the timing of the drift tube; and introducing an electrostatic non-dispersive deflector. The increased shutter speed allows for a much increased acquisition of spectra; while previously it had only been 30 spectra per second with a 1 keV field-of-view, it became 1000 spectra per second with a 2 keV field of view with a 200 kV beam. As two simultaneous images are taken, the field-of-view can be up to 4 keV. In order to collect the two images simultaneously the CCD is split in half, with one side collecting the low- and the other the high-loss electrons. The high- and low-loss images are taken at different shutter speeds to account for the difference in intensities and the high-loss region is offset to the edge of the low-loss region (around 200eV). Another improvement that the GIF brought was an exposure time that was increased by seven orders of magnitude. The GIF also has a BF/DF detector that allows for spatial resolution through the combination of images and spectral information. A diagram of the GIF can be seen in section 2.2.3.

Having both the core- and low-loss regions means it is possible to calculate sample thickness, perform absolute quantification, account for and remove plural scattering, and combine the high-loss features with the low-loss ones. It also allows for spatial resolution when performing quantification; if the images are taken sequentially, as was the case before DualEELS then the beam or sample may have drifted and there was no guarantee that the low- and high-loss images had been taken from the exact same position.

3.3 Absolute Quantification

Quantification of core-loss EELS spectra has been pursued right from the earliest days of using EELS, all based on the following equation

$$I_{edge}(\Delta E) = I_0 \; N \; \sigma(\Delta E), \tag{3.1}$$

where $I_{edge}(\Delta E)$ is the intensity in an energy range ΔE for the edge, I_0 is the intensity in the zero loss peak, N is the number of atoms per unit area in the path of the beam, and $\sigma(\Delta E)$ is the cross section. The cross section describes the probability of the interactions that give rise to the edge.

Initial work focussed on producing simple ways of predicting $\sigma(\Delta E)$, so that the spectra could be quantified using approximations of atomic structure like the Hydrogenic model [107] the Hartree-Slater model [101, 108]. This method was applied to the K, L and M shells of a range of elements in due course [109]. However, those models do not consider the effects of bonding and band structure, and so, near to the edge onset there were structures that might differ significantly from the simplistic models. This structure is known as the energy loss near-edge structure ((ELNES) see section 3.5.2) and can have a significant effect on any quantification attempts with realistic energy ranges (ΔE). Unfortunately, this means that simple quantification of ratios using only calculated cross-sections is subject to large errors. More advanced methods (now implemented in commercial software) for EELS modelling-based quantification using this earlier method (based on earlier work by Verbeeck and others [110, 111, 112]) still suffer from the fact that the calculated cross-sections may not well represent the ELNES.

For EELS, it was easiest initially to simply record a spectrum over a limited energy range, incorporating strong edges for any elements of interest, and

then perform a relative quantification. Absolute quantification was difficult to achieve, as it relies on the ZLP (its intensity is the I_0 in eq. (3.1)). The intensity of the ZLP is so much greater than the intensities of core-loss edges that recording them simultaneously was virtually impossible. The solution was to convert Equation (3.1) to instead account for the ratio of two elements in a spectrum

$$\frac{N_A}{N_B} = \frac{I_A \, \sigma_B}{I_B \, \sigma_A}. \tag{3.2}$$

Nevertheless, absolute quantification was occasionally performed by some groups [113, 114]. Standard samples could be used to determine cross-sections on an absolute scale, which could be compared to calculated cross-sections. The percentage discrepancies found in reference [113] are still reported, although what were reported as systematic errors in one direction are now often treated as random errors in either direction.

It must be stated that the assumption thus far is that there has only been single scattering in the spectrum and no plural scattering. If there is plural scattering then eq. (3.2) is no longer strictly true. However, while there is always plural scattering in any sample of reasonable thickness, it can be neglected if the sample thickness is less than 30% of the mean free path for inelastic scattering. Even in thicker specimens, all hope is not lost. The main effect of increased thickness is that the low-loss shape tends to get reproduced behind the edge onset. This is because the low-loss is the most likely event, and the edge shape becomes convoluted with the low-loss spectrum, resulting in a change in total edge intensity as well as differences in both its shape and position – the edge will move to a higher energy-loss. This means that in relative quantification it

will likely still work passably, provided the edges are far enough apart and a sufficiently large energy window post-edge is used in quantification.

As established, the difficulty of absolute quantification is that, for normalisation, it requires the ZLP intensity (I_0). As DualEELS takes the high- and low-loss simultaneously it overcomes this issue. DualEELS also allows for the absolute energy calibration by the ZLP position (although linearity of the energy dispersion may need checking and adjusting [115]); and the deconvolution of plural scattering from core-loss edges [96].

Work at Glasgow has definitively shown that absolute quantification is possible using EELS, with two caveats. The first is that DualEELS must be utilised, and the second is that there must be standard spectra for the elements of interest available at the relevant edges, in something close to the chemical form likely to be encountered in the experimental sample. This method has been used for the absolute quantification of nanoscale vanadium and titanium carbonitride precipitates in steels. The absolute cross-sections were determined experimentally from standard samples of TiC, TiN, VC, and VN [101]. These could then be used to quantify real nanoscale precipitates in steels [108] after some basic work on pre-processing the datasets [104].

3.4 Processing EELS Data

As the beam rasters across the sample the spectrometer will collect a full spectrum from every pixel, meaning that in STEM three-dimensional arrays are created for each dataset. These arrays therefore provide both spectral and spatial information about the sample and are known as a spectrum image (SI)[1]. Two SI arrays are collected initially; one high- and one low-loss, each covering different energy ranges within the spectrum and with different exposure times. After some processing it is possible combine the two into a single array. A visualisation of the arrays can be seen in fig. 3.2.

EELS spectra require processing before they can be analysed due to the amount of noise in the data. This processing is standard and is largely automated in DM. The processing steps are detailed below.

[1]It is possible to cycle through all the energies in DM by adjusting energy sliders. While much more information can be drawn from the SI, simply passing through an important energy (e.g. at the energy of a certain edge) can also provide insight. For example, there may be a sudden switch in the intensity of features in the SI at that energy, hinting that a particular element is present.

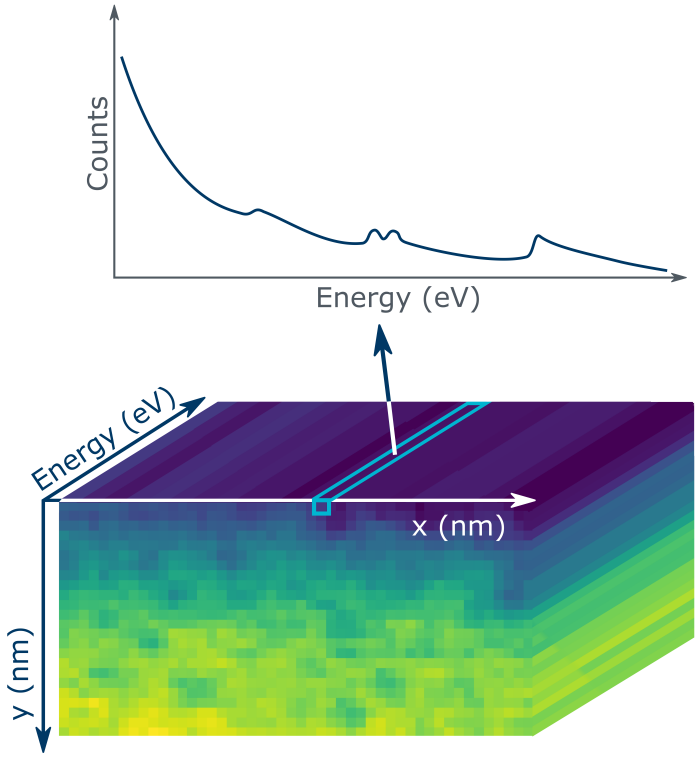

Figure 3.2: A visualisation of an EELS SI 3D array, showing the physical dimensions in the xy-plane and energy in the z plane. The high-loss EELS spectrum for a pixel is shown (note that this spectrum is only a sketch). This samples is aLIGO zirconia-doped tantala sample, deposited in argon.

ZLP Alignment

The first processing step is to align the datasets. As previously discussed, the ZLP is generally offset in the low-loss spectrum by a few electron Volts, meaning that all the following data is also offset by that amount. DM is able to identify the centre of the ZLP and adjust the spectrum accordingly, adjusting the high-loss data at the same time. This step also produces a correlation coefficient and a peak shift. The shifting of the centre of the ZLP is done by fitting a Gaussian and using the user-provided settings. In actuality, the ZLP is asymmetric and therefore not a Gaussian however, this is not noticeable at low dispersions (like

0.5 eV per channel, as is the most common dispersion for work in this thesis), where the point-spread function of the detector dominates the shape of the ZLP.

X-Ray Removal

There are frequently spikes in EELS spectra caused by external radiation such as cosmic ray muons and off-axis electrons. The latter can cause bright flashes if they strike an internal component of the spectrometer and excite a bright enough x-ray, which in turn could hit the phosphor and give a bright flash in one pixel. This step removes spikes that are greater than 10 standard deviations by setting the value of that pixel to the local median. This must be done for both the high- and the low-loss data.

Cropping

The edges of the low-loss spectra are then cropped out as the outermost channels may contain and cause errors. There may be shadowing or rescattering from the shutters that restrict which parts of the spectrum reach the CCD and so the very edges may contain odd spectrum shapes and artefacts. The far left of a low-loss spectrum contains no signal – only noise – so there is no point including it. The cropping is done by extracting the volume (i.e. number of channels). The exact energy range can vary by sample and collection conditions but between 150 and 200 eV is sufficient.

Principal Component Analysis

Principal component analysis (PCA) removes noise from the spectra and should[2] be performed for both high- and low-loss data [104]. It utilises a multivariate statistical analysis plug-in for DM [116] and works by assuming that the SI is a linear combination of information from individual components. It then separates the SI into eigenvalue plots of these different components, which can then be viewed in order of decreasing importance to determine which components contain signal and which ones are only noise. The SI is then rebuilt using only the 'useful' components to eliminate as much noise as possible. It is important to not also discard useful data during this step [117].

Splicing

Up until this point, the high- and low-loss SI have been kept separate, however to perform absolute quantification a unified spectrum is required. It is possible to splice two spectra together in DM with ease, as there is a plug-in for this purpose. The number of overlap channels to calculate the scaling factor can be adjusted as necessary, and if required, the dispersion can be automatically adjusted (although this was never necessary for this thesis' work). Once the splice has been tested on individual pairs of spectra the entire SI pair can be spliced together.

At this point it is important to extract the ZLP from the spliced spectrum, as this will be important in later when performing absolute quantification.

[2]While it is possible to analyse data without this step, it is heavily advised as the results with it are much cleaner.

Thickness

It is also possible to calculate the thickness of the sample from the spliced SI, as discussed in section 2.3.3 (see eq. (2.8)). DM has a plug-in that creates an image map of the relative thickness (i.e. t/λ) and it is possible to convert this to absolute thickness, if λ is known. At this point it is beneficial to also extract the ZLP intensity, for which DM has another plug-in.

Deconvolution

In the EELS spectrum a significant number of the counts are a result of electrons that have been multiply scattered, unless the sample has a thickness of less than a third of the mean free path. This can affect both the intensity and shape of the high-loss edeges. Therefore, the final step in the basic processing is performing Fourier logarithmic deconvolution, in order to remove plural scattering, ensuring that the high-loss edges' intensities and shape come from single core-shell scattering, and not any preceding or following interactions.

Deconvolution requires the whole spectrum to have been collected so that Poisson statistics hold. The recorded spectrum can then be described thus:

$$R(E) = B(E)[I_0 + S(E) + D(E) + T(E)...N(E)], \tag{3.3}$$

where $B(E)$ is broadening due to instrument resolution, I_0 is the ZLP intensity (i.e. the elastic contribution to the spectrum) and the following terms are the single, double, triple and so on, scattering terms. The plural scattering terms are made up of repeated convolutions of $S(E)$. For example,

$$N(E) = \frac{N(E)^N}{N!}. \tag{3.4}$$

Through a Fourier transform of eq. (3.3), it is possible to arrive at an equation for the single scattering distribution as a function of frequency. $S(E)$ can then be recovered by performing a reverse Fourier transform. A more thorough mathematical discussion of this process readers can be found by reading references [118, 119, 120].

Deconvolution is done through a DM plug-in, and it is possible to adjust the zero-loss model as well as the noise reduction parameters. The standard Gaussian model was generally used, as Gaussian fits are very robust. There is, however, also a theoretical justification for using this model. The energy distribution on the ZLP is not Gaussian or symmetric but this is not visible at the dispersions used in this thesis, where the dominant factor here is the point spread function of the detector, which gives a Gaussian blur. The model for the ZLP is of importance as both the height of the ZLP and its shape are components in the deconvolution calculations.

If there are any negative counts in the spliced SI this can either cause the deconvolution to fail or create artefacts in the deconvolved SI. Note that this step is optional but was performed for the work in this thesis. An example of how deconvolving can change the spectrum can be seen in fig. 3.3.

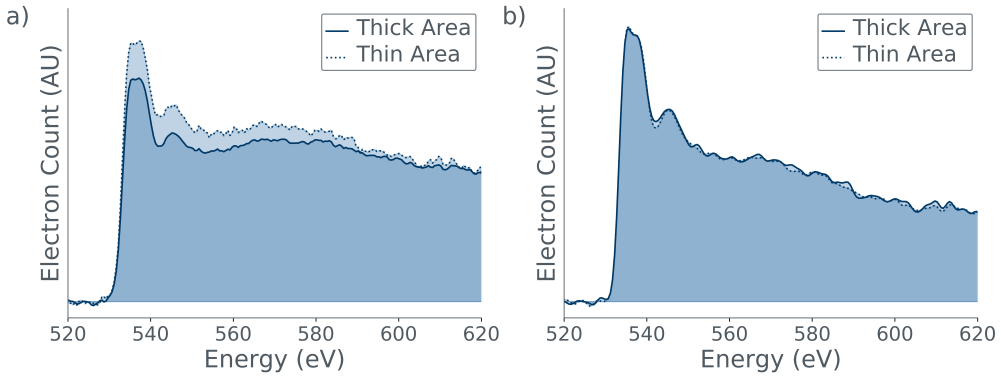

Figure 3.3: How thickness effects the spectrum; a) spectra extracted from thick area; spectra extracted from thin area; c) spectra extracted from deconvolved spectrum. Note that all of these spectra have been normalised by intensity, as otherwise the spectra taken from thicker areas would be weaker. This sample is an aLIGO zirconia-doped tantala sample, deposited in argon.

It is important to check around the splice point in the deconvolved SI as sometimes there is an issue with the splice resulting in Fourier noise about the splice overlap. If this occurs, then the splice parameters can be adjusted and then the splice and deconvolution performed again.

3.4.1 Multiple Linear Least Squares Fitting

Like PCA, multiple linear least squares fitting (MLLS) relies on multivariate statistical analysis. It works by linearly combining spectra provided by the user to recreate the SI

$$m(E) = c_1\sigma_1(E) + c_3\sigma_3(E) + c_3\sigma_3(E) + ..., \tag{3.5}$$

where $m(E)$ is the model created by the fit, c_n represents the amount of each component present and $\sigma_n(E)$ the components of the fit, which come from the user-provided spectra. Exactly what these components are depends on what the goal of the fit is and they can be qualitative or quantitative. A fit is often made up of a mixture of both qualitative and quantitative spectra. The reference

spectra can be taken from different points in the deconvolved SI or can be newly introduced, such as a elemental edge standard from the EELS Atlas [121].

The components can only be quantitative if their corresponding spectra are in absolute units of $Barns/eV$. If the components' spectra are taken from the SI they are qualitative, i.e. are used only for their shape [3]. This introduces an element of trial and error to the MLLS process; it can be unclear initially how many reference spectra are required, and from what region of the sample they should be taken. For much of the work in Chapter 4 two reference spectra were taken to represent the main body of the sample. The composition of the main body of the samples was of little interest as it was known but it was necessary to account for the matrix portion of the sample, and so spectra were taken from thicker and thinner areas of the sample. It may also be necessary to adjust the energy range of the fit. How 'good' the fit is can be assessed by the residual and the residual χ^2. The residual is an SI of the residuals in every channel of the spectrum. In general, both the residual and the residual χ^2 should be low when compared with the fit itself. The residual χ^2 can be used to see where the fit is failing (i.e. in which part of the SI) and the residual can be used to see how it is failing.

Even if there only qualitative spectra taken from the deconvolved SI, it is still possible to calculate elemental ratios. If using a quantatitve elemental standard it is then possible to perform absolute elemental quantification using the relevant fit component. The fit component will need to be adjusted, however before the intensity of the image is in number of atoms

$$C_{atoms} = \frac{C_{initial}\ a_p}{I_{zl}\ w_{ch}} \times 10^{28}, \tag{3.6}$$

[3] Arguably the spectra taken from the SI are quantitative, the calibration scale is just unknown.

where C_{atoms} is the fit component with its intensity in number of atoms; $C_{initial}$ is the original fit component; a_p is the area of the pixel; I_{zl} is the extracted ZLP intensity; w_{ch} is the channel width; and the 10^{28} is to convert from $barn^{-1}$ to m^2. A discussion of creating a semi-empirical standard for quantitative analysis can be found in section 4.1.

Figure 3.4 shows the different images and spectra that are part of producing a MLLS fit. This example is representative of the majority of the fits produced for chapter 4 and shows a titania-doped tantalum pentoxide GW mirror coating that contains small bubbles of argon. In DM there are also two SIs produced – the fit itself and the residuals of the fit – but these were omitted here as there is little use in showing a single slice of a three dimensional image. A HAADF image was included instead, to illustrate from where the reference were taken.

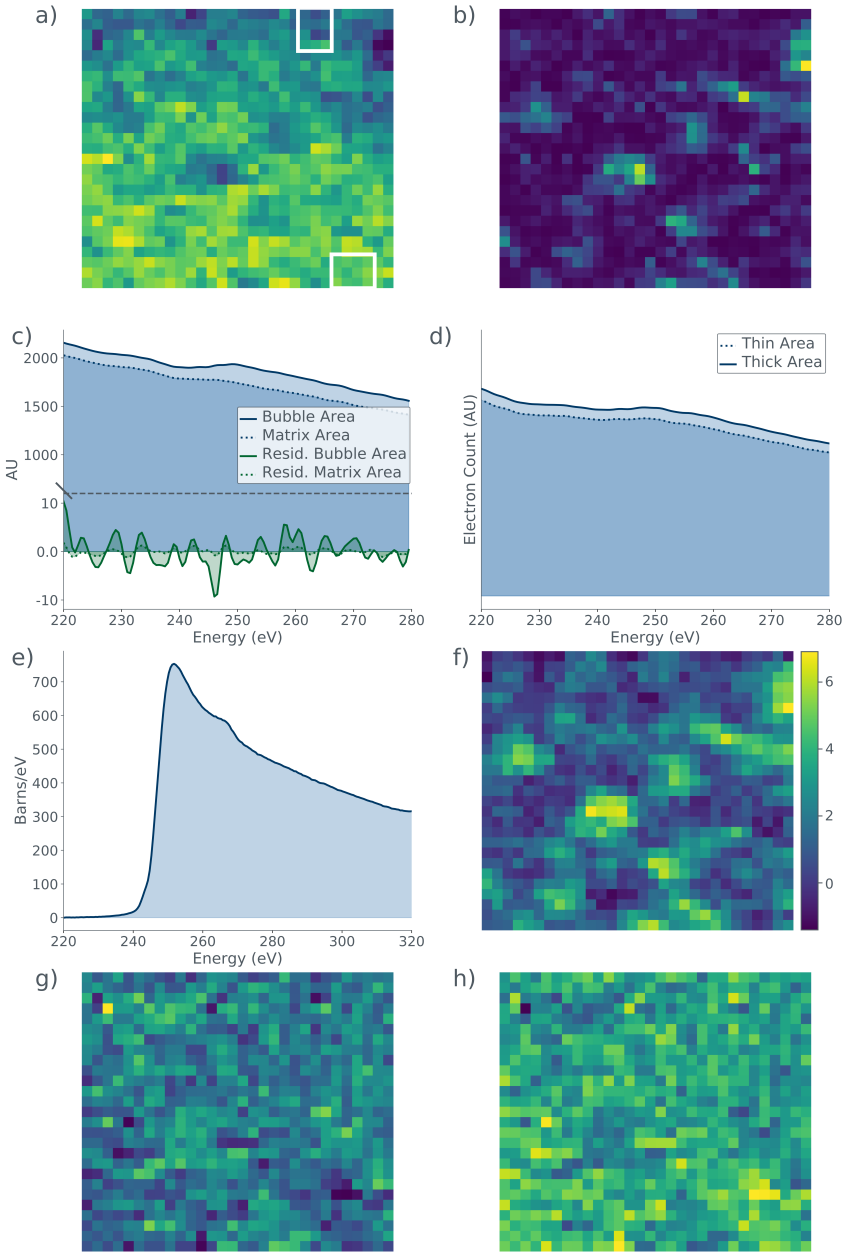

Figure 3.4: A breakdown of an MLLS fit, using two reference matrix spectra, showing what is used for and produced by the fit, done on a titania-doped tantala sample: a) a HAADF image of the dataset to show from where the matrix spectra were taken (top rectangle = thicker region; bottom rectangle = thinner region); b) the residual chi squared, showing areas of intensity where the fit performed worst, in this case it is likely due to the edge shape being different in the sample than in the qualitative reference spectra, and is not a cause for concern; c) the MLLS fit, taken from a matrix region and at one of the bubbles as well as the residuals to the fit taken from a matrix region and at one of the bubbles. How much larger the residuals are in the bubble region than the matrix confirms what can be seen in (b) – the residual chi squared. Note the break in the axis of this image, the residuals are <5% than the fit itself; d) extracted matrix spectra used to perform the fit; e) the qualitative reference spectra used to perform the fit, showing the argon $L_{2,3}$ edge; f) the fit coefficient for the qualitative reference spectra, showing areas of brightest intensity where the argon lies; g) the fit coefficient for the matrix spectrum taken from the thick region of the sample; h) the fit coefficient for the matrix spectrum taken from the thin region of the sample.

3.5 EXELFS

In the few hundred electron volts following an EELS edge there will be subtle oscillations. These are caused by the interference of the electron scattering from the atom and elastically backscattered electrons from the surrounding atoms. The electrons at this point can be thought of as spherical waves, as seen in fig. 3.5. These oscillations contain information about the short range structure of a material. Both the oscillations and the area in which they lie is generally known as the extended energy loss fine structure (EXELFS) and it was discovered in 1976 by Leapman and Cosslett [122].

EXELFS is analagous to a technique called extended X-ray absorption fine-structure spectroscopy (EXAFS), a subset of X-ray absorption spectroscopy (XAS). Most of the analysis techniques for EXELFS are taken from earlier work on EXAFS. Despite this, there are still several differences between the two techniques, perhaps most notably that EXAFS is generally used for higher losses (2keV to 35keV). There are advantages to using EXELFS; it can be done 'in-house' instead of at a beamline, it gives greater spatial resolution and full characterisation of the material can be done using the microscope's other capabilities [123]. Despite this, EXAFS is more commonly used; microscopy samples are more susceptible to radiation damage because they are smaller and so experience a higher beam energy dose per unit volume. The signal to noise ratio has also historically been a far greater problem in microscopes, although advances are being made in this area [1]. This difference in SNR can be seen in fig. 3.6, where the XAS spectrum is much smoother than the EELS spectrum. The difference between the subregions of EELS and XAS spectra can also be seen on this figure.

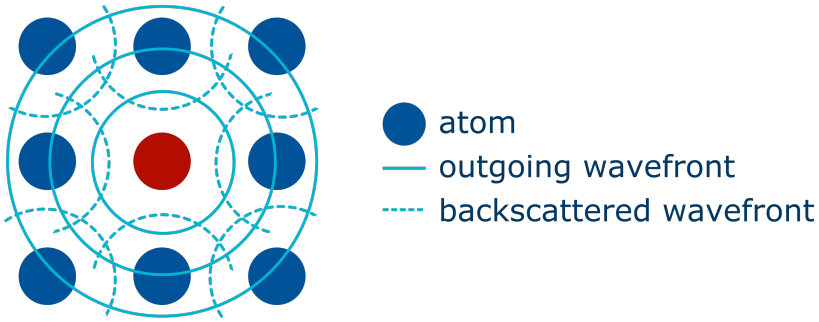

Figure 3.5: The electron wavefronts causing the oscillations following an EELS or EXAFS edge. The initial, outgoing wavefronts from the central/target (red) atom are shown with a solid line and the backscattered wavefronts from the neighbouring atoms are the dashed lines.

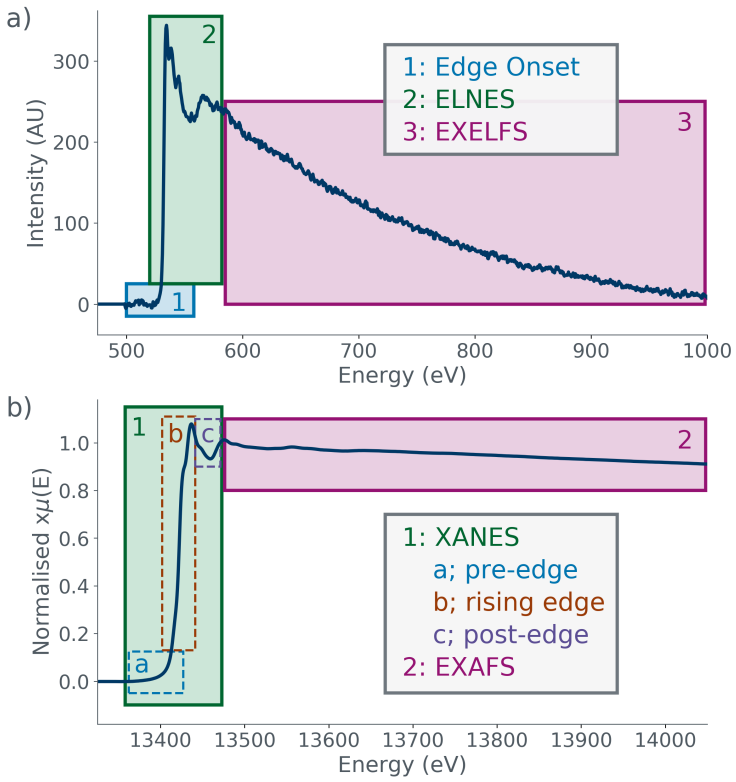

Figure 3.6: Breakdowns of the different regions around an edge for EELS and XAS: a) an EELS spectrum of the titanium k-edge showing the edge onset, the ELNES region and the EXELFS region; b) an XAS spectra of the bismuth k-edge showing the XANES region, further broken down into the pre-edge, rising edge and post-edge regions, as well as the EXAFS region. The bismuth data was collected by Ramos, Arnold and Rowe at the I20 Beamline at Diamond Light Source. Thank you to Dr. Ramos for her permission to use this spectrum.

EELS and XAS spectra appear very similar at first glance but there is one crucial difference. In EELS spectra the y-axis is initially photon counts, generally changing to arbitrary units of intensity post-processing, however in XAS the y-axis is the absorption coefficient.

3.5.1 Deriving the EXAFS Equation

EXAFS can be modelled using the EXAFS equation, and due to their similarities the EXAFS equation can also be used to model EXELFS. The EXAFS equation is derived below. This derivation has been included as the equation itself is used in analysing some of the data in chapter 5.

In XAS, following the interaction of the X-ray beam with the target atom, an electron will be emitted due to the photoelectric effect. As stated above, this photoelectron will be scattered and then backscattered. The backscattered photoelectron will return to the target atom before it has de-excited, therefore changing the probability of further X-ray absorption by changing the probability of creating a new photoelectron. As the scattered photoelectron is travelling at non-relativistic speeds its wavenumber (k) is given by

$$k = \frac{2\pi}{\lambda}, \tag{3.7}$$

where λ is the electron's wavelength. This equation can be approximated to

$$k \approx \frac{\sqrt{2m_0(E - E_k)}}{\hbar} \tag{3.8}$$

where m_0 is mass of the electron; $\hbar$ is the reduced Planck's constant; E is the energy transfer between the scattered electrons (or X-ray photon in the case of EXAFS); and E_k is the threshold energy of the edge. The threshold energy can be defined as the energy at which the electrons have just enough

energy to propagate through the material [124]. As E increases the interference oscillates between constructive and destructive, creating maxima and minima in the wavefront.

The EXAFS equation is used to describe the contribution a neighbouring atom has on the absorption coefficient of the central atom. The absorption coefficient describes how strongly the X-rays are absorbed; as their energy increases they become more penetrating and the absorption coefficient will gradually decrease. Similarly to EELS, at characteristic energies there will be a spike in absorption. These spikes occur when the energy of the X-ray is the same as the shell energy of an element within the material. The absorption coefficient (μ) is defined as

$$\mu = -\frac{dln(I)}{dx}, \tag{3.9}$$

where I is intensity of the incoming beam and x is the distance travelled within the sample. We then use μ to create a function that describes the EXAFS region

$$\chi(E) = \frac{\mu(E) - \mu_0(E)}{\delta\mu_0}, \tag{3.10}$$

where $\chi(E)$ is the change in the absorption coefficient caused by the neighboring atoms, $\mu_0(E)$ is the background and $\delta\mu_0$ is the normalisation factor, which in EXAFS is the large increase of the absorption coefficient at the edge. $\chi(E)$ can also be given as

$$\chi(E) = \frac{J_k^1(E) - A(E)}{A(E)}, \tag{3.11}$$

where $J_k^1(E)$ is the single scattering distribution and $A(E)$ is the theoretical energy-loss intensity that would occur without backscattering, using a single

atom model. Equation (3.11) assumes that plural inelastic scattering effects are negligible or have been removed by deconvolution, and that, similarly, instrumental resolution is not of concern.

Using this alongside Fermi's Golden Rule[4] it is possible derive an equation for the probability of absorption (χ, i.e. the EXAFS signal) as a function of the photoelectron wavenumber (k), which is known as the EXAFS equation

$$\chi(k) = {S_0}^2 \sum_j \frac{N_j}{k{R_j}^2} \left|f_j(k)\right| e^{\frac{-2R_j}{\lambda_e(k)}} e^{-2k^2{\sigma_j}^2} \sin(2kR_j + 2\delta_j(k)), \qquad (3.12)$$

where ${S_0}^2$ is the amplitude reduction factor, N_j is degeneracy, R_j is the distance between the absorbing and scattering atoms, $f_j(k)$ is the form factor, λ_e is the photoelectron wavelength, ${\sigma_j}^2$ is the Debye-Waller factor and $\delta_j(k)$ is the phase shift. The first exponential term in the equation accounts for suppression due to inelastic scattering of the photoelectron and the decay of the core hole. ${S_0}^2$ accounts for the difference between the initial and final states of the absorbing atom. This difference is due to the lost electron and the orbitals being more strongly affected by the positive charge of the nucleus. It is normally about 0.8 [125]. $f_j(k)$ accounts for the fact that atoms are not point-like; electrons are distributed in space around the atom. It is sometimes also referred to as the electron scattering amplitude. ${\sigma_j}^2$ factor accounts for the attenuation of X-rays as a result of thermal disorder – atoms vibrate at a frequency dependent on temperature (around 10^{13} Hz at room temperature) [126]. Through using the EXAFS equation to interpret the oscillations post EELS and XAS edges, it is possible to ascertain information about the fine structure of the sample.

[4]This states that the probability of absorption is proportional to the density of states (i.e. the number of possible routes for absorption) and the energy difference between the initial and final state.

3.5.2 ELNES

The region encompasssing the first 50 eV post edge onset is known as the ELNES. This region can give information on the density of states of the target atom, as it is caused by the excitation from a core orbital to unoccupied bands [127]. It can be interpreted in a variety of ways; through the band structure, the multiple scattering, molecular orbitals or multiplet processes [96]. There can be multiple scattering within the atom that contribute to the ELNES (see fig. 3.7).

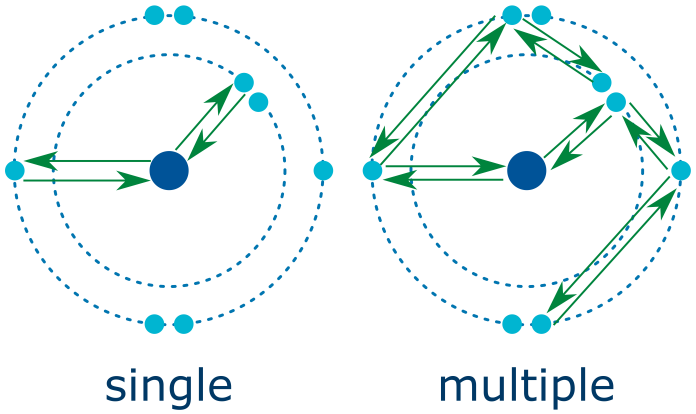

Figure 3.7: Single versus multiple scattering within an atom, both of which contribute to the ELNES.

Like EXELFS, ELNES is analagous to an XAS phenomenon, known as X-ray absorption near-edge structure (XANES, or NEXAS). XANES appears between 50 and 100 eV post edge onset. It is very sensitive to the geometric environment of an excited atom, as well as its spin state and oxidation. This is because these things affect the density of the unoccupied states available to the excited electron. The small changes visible in XANES therefore contain a lot of important information but typical XAS does not have sufficiently high resolution to track these changes. High resolution XAS overcomes this problem but it is rarely employed for solid state systems because it is much slower than conventional XAS. XANES and EXAFS are sometimes grouped

togther into X-ray absorption fine structure (XAFS). I have borrowed this convention and will be using the term energy-loss fine structure (ELFS) to when referring to both EXELFS and ELNES.

3.6 Adapting EXAFS Software for EELS Analysis

There are a lot of programs around for EXAFS data analysis and processing. Due to the similiarities between EXAFS and EXELFS it is possible to use these programs for examining EXELFS data, however there are naturally some limitations due to their differences.

Below are the programs used for the work in Chapter 5.

3.6.1 Demeter

Demeter was created by Ravel in 2005 and works as a graphical user interface for Newville's IFEFFIT library [128, 129]. It also makes use of the University of Washington's FEFF6 code [130] (see section 3.6.2). Demeter is split into three components; Athena for processing, Artemis for analysis and Hephaestus, which provides an enhanced periodic table. IFEFFIT is designed, amongst other things, to remove background, perform Fourier transforms, and fit data to FEFF results and other models.

For a comprehensive guide to using Demeter, readers are referred to the Athena and Artemis manuals.

Athena

Athena is used primarily for absorption normalisation (using eq. (3.10)) and background removal. Not only can it plot data in its original form of absorption coefficient (or intensity in the case of EELS data) versus energy, it can also transform it into k-space (wavespace, calculated using eq. (3.8)); R-space (real space, done using a Fourier transform of k-space); and q-space (a Fourier transform of the real space data back into k-space, retaining only the real components). Figure 3.8 shows a flowchart of the steps of processing data using Athena.

There are a variety of things that can be adjusted to improve the normalisation and background removal. The first is called the “rbkg" value, which is the value below which Athena removes Fourier components when plotting R-space. It defaults to one but may need to be adjusted for data with a lot of noise or two edges close together. Another is spline clamps, which tie the background more tightly to the data. They are generally only helpful at the high energy end of the data, where it tends to be steadier. Adjusting them can amplify data that is very small in the high-k range but can also affect the quality of the background.

The k-weight is a way to emphasise different sections of k-space spectra. A higher k-weight will amplify the high-k end of a spectrum. This value can be adjusted in two ways; first in background removal and second in plotting. The effect of changing the k-weight on both the k- and R-space spectra can be seen in fig. 3.9, where, as the k-weight increases, the oscillations in k-space are shifted to the right, and the peaks of the R-space plot become more discrete. It also affects the y-axes of all the plots, changing the units and values. The different k-weighting is used to help understand the data by magnifying differents portions of the spectrum and does not affect the data itself.

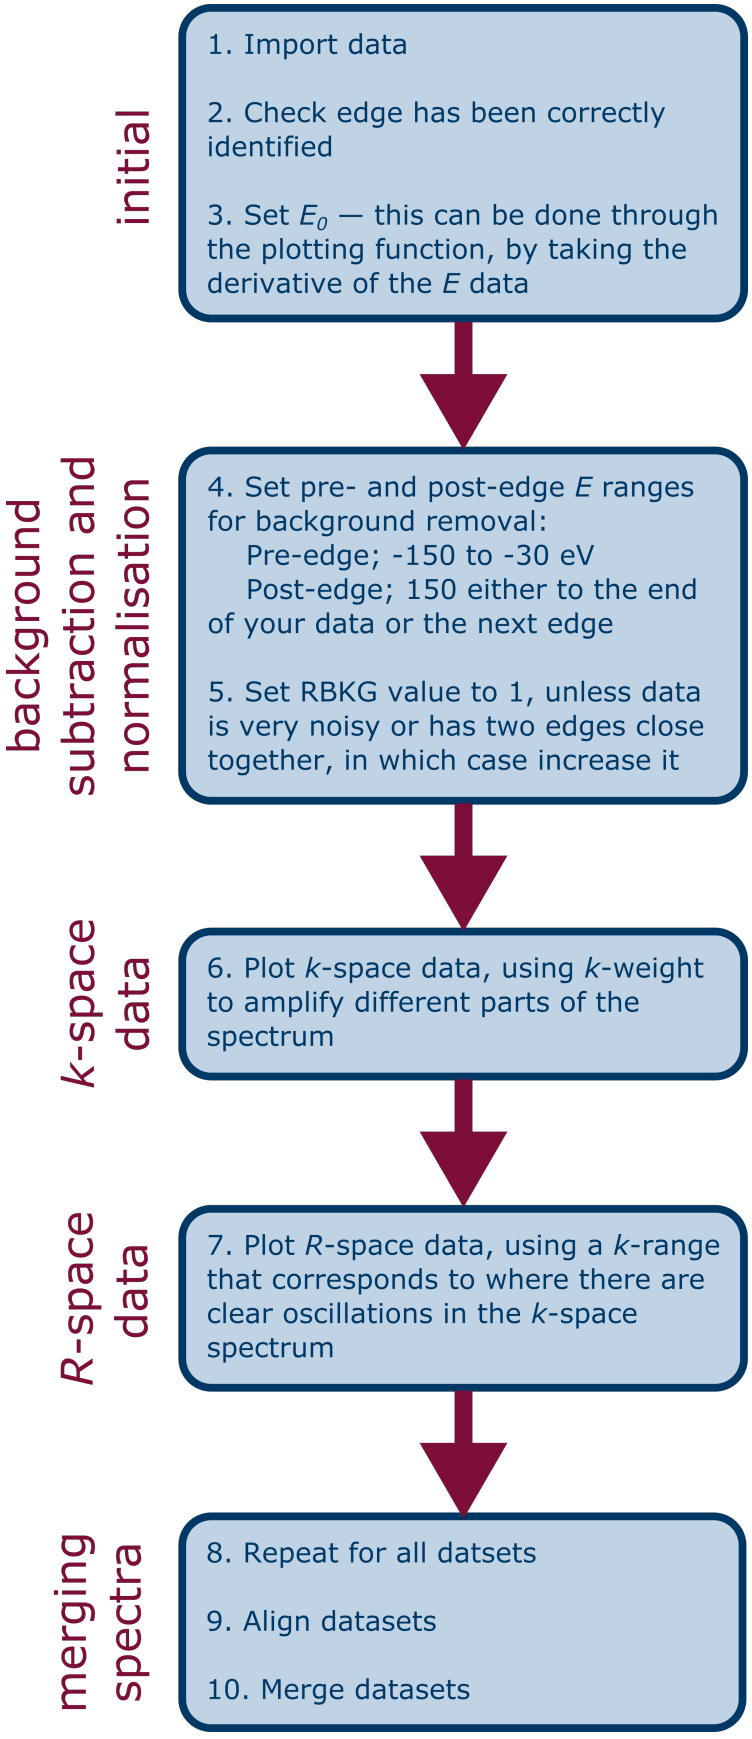

Figure 3.8: A flowchart detailing the steps of processing data in Athena.

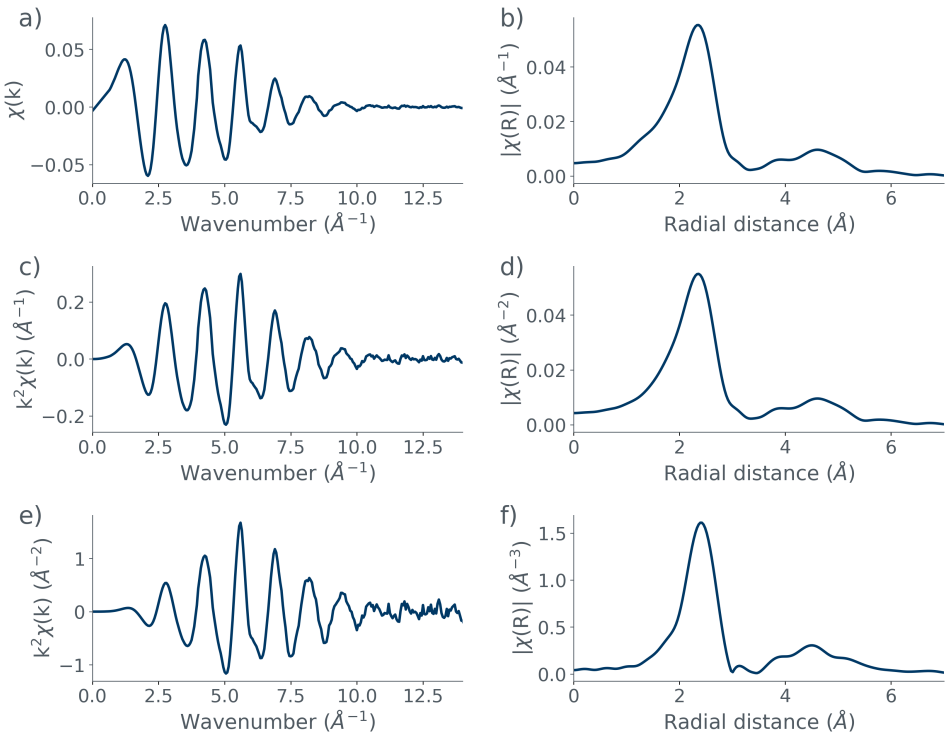

Figure 3.9: An example of data plotted in k- and R- space, showing how k-weighting can affect the data: a) the k-space plot, showing clear oscillations in the first half of the spectrum before it decays into noise on the right. This is plotted with a k-weight of 0; b) the R-space plot, also plotted with a k-weight of 0; c&d) k- and R-space plots with a k-weight of 1; e&f) k- and R-space plots with a k-weight of 2.

When plotting in R-space a k-range must be chosen. This is generally chosen to be between 3 Å^{-1} and wherever clear oscillations end in the k-space plot. Clear oscillations are defined as those not dominated by noise, which appears as irregular spikes. The lower value is the same regardless of the data due to the nature of the background removal, which is unreliable at low k-values. The higher the upper end the better and if there is not much space after the edge, it can be hard to produce meaningful k- and R-space plots.

Finally, there is an option for phase correction, which subtracts the central atom phase shift when calculating the Fourier transform. It does not include any phase shift due to the scattering atom.

While quantitative analysis cannot be done on the results of Athena processing it is still possible to draw some conclusions. The R-space plots show the average distribution of atoms around the central atom and the shape of the peaks can also be illuminating; broader peaks mean a disordered structure, while distinct peaks suggest a crystalline structure. It is important to note that the absolute values of the distances between the peaks cannot be compared between datasets due to phase differences. The focus should be on their relative separations and intensities.

Artemis

After the data has been processed by Athena, Artemis uses the EXAFS equation in order to calculate the respective positions of atoms within the sample, as well as their atomic species. It does this using crystallographic information files (CIF) as a model and the IFEFFIT library to run FEFF calculations, as will be discussed in section 3.6.2.

CIF are the standard form of storing crystallographic structure and contain unit cell values, atom names and their coordinates, as well as any modifying factors such as the temperature the data were collected at. They are used in Artemis to build an expected model of the material investigated. Artemis then uses the CIF file to build a model of the scattering paths based on FEFF calculations. This model is then fitted to the experimental data. Crucially, the parameters of the fit are *not* the parameters of the EXAFS equation but rather, the parameters of the EXAFS equation are written in terms of the parameters of the fit.

It is an iterative process, initially only the most important scattering path is used to create the fit and gradually more are added. The initial parameters are given expected values and new values are calculated during the fit. These

are the best indication of how reliable the fit is. Occasionally the values given are unphysical, such as a negative Debye-Waller parameter (impossible as it represents the fraction of elastically scattered X-rays), showing that there is a problem with the fit. Another indication of the quality of the fit is how strongly correlated the different parameters are. The goal is to keep them as low as possible, although sometimes correlations can be as high as 0.9 between similar parameters. Correlations can be dealt with by 'fixing' one parameter, which should only be done if the value is within the expected range.

Possible fitting parameters include degeneracy (N); change in absorption threshold (ΔE_0); amplitude (${S_0}^2$); Debye-Waller (σ^2); and thermal parameter (δR). The absorption threshold is used to set the zero value for k-space and is incorporated into the EXAFS equation. δR is the expansion due to temperature change

$$R_j = R + \delta R. \tag{3.13}$$

where R_j is taken from the EXAFS equation and is the distance between the absorbing and scattering atoms, and R is the bond length between the absorbing and scattering atoms, taken from the CIF file.

While Artemis was used for some of the analysis in Chapter 5 it is here that the differences between EELS and XAS, not their similarities, begin to dominate. Limited qualitative analysis was possible but EELS requires a different treatment and the FEFF6 upon which Artemis relies simply is not capable of accounting for it. Artemis is also of little use when examining XANES data, and so cannot be used for ELNES either.

3.6.2 FEFF

FEFF was created by John Rehr and his team at the University of Washington, Seattle. Its name comes from the effective curved wave scattering amplitude, which was denoted as f_{eff} in the original form of the EXAFS equation (represented by the form factor in eq. (3.12)). FEFF performs *ab initio* real space multiple scattering calculations using Green's functions and is based on spherical muffin tin scattering potentials. A muffin-tin potential is named as it can be visualised exactly like a muffin (or cupcake!) baking tray; it is made of spherical scattering potentials centred on each atom, with a constant value in the interstitial values between the atoms [131]. FEFF includes polarization dependence, core-hole effects and local field corrections, and was built using FORTRAN90 code. Its first iteration was in 2009 but has gone through many improvements and is now at FEFF10. Perhaps the most important improvement for this thesis was the inclusion of ELFS data as a possible input. This means it may be possible to better model the ELNES.

3.7 Summary

This chapter went into further detail about EELS, including the different areas of an EELS spectrum; the key advancement of DualEELS; absolute quantification; and the steps of processing EELS data. It also discussed ELFS and its similarities to XAFS; the underpinning theory of EXAFS that has been applied to EXELFS; and how programs originally designed for XAFS analysis can be used for ELFS.

The EELS spectrum is comprised of two areas; the low-loss region, which includes the ZLP and the plasmon peaks, and the high-loss region. The high-loss region includes characteristic edges that are caused by the electron beam interacting with the core-shell electrons of the atoms in the sample. DualEELS is the collection of high- and low-loss spectra simultaneously by splitting of the CCD and the use of an electrostatic shutter. This allows for absolute spatial quantification, calculation of the sample thickness and the removal of plural scattering. The underpinning theory and the history of absolute quantification were discussed. The steps of processing EELS data were laid out and discussed, including ZLP alignment, background subtraction, X-ray removal, cropping, PCA, splicing, thickness calculations, deconvolution and MLLS fitting. EXELFS are the subtle oscillations found in the several hundred electron volts following an EELS edge. These oscillations are caused by the interference of the electron wave scattered from the target atom and the backscattered wave from the surrounding atoms; and they can provide information on the short range structure of the material. EXELFS bears many similarities to EXAFS as the two are caused by the same effect and so EXELFS can be described by the EXAFS equation. ELNES is the 50 eV following the EELS edge onset and contains information on the density of states of the target

atom. Due to the similarities between XAFS and ELFS it is possible to use software designed for XAFS analysis for ELFS data.

An understanding of the details of EELS spectra is crucial to understanding the work in this thesis, as is understanding some of the underlying theory. Having an appreciation of the processing is also important as without it, the sensitive work in the following chapters would be impossible. MLLS fitting and absolute quantification is required for the work in chapter 4 and EXAFS techniques for chapter 5.

4

Quantifying Gas Bubbles

As discussed in section 1.5.2, thermal noise is one of the most significant limiting factors in the sensitivity of LIGO's detectors and the coatings of the mirrors within the detectors are one the greatest contributors to that thermal noise. Understanding the structure of the coatings allows their noise contribution to be characterised, which, in turn, allows this noise to be removed from the final signal or, better still, helps the search for new coatings without this noise contribution.

This chapter discusses work undertaken to characterise previously poorly understood characteristics of the high refractive index layer of the GW mirror coatings – tantalum pentoxide, commonly referred to as tantala.

It had been established for some time that there were voids in the coatings and that, based on simulations, they were likely spherical in shape [132, 133]. Rutherford scattering and ellipsometry data showed that the films contained

argon, although it was not then known where the argon was in the coatings [134]. The work in section 4.3.1 was begun knowing only these two facts; however, concurrent work by other members of the LIGO collaboration showed that in similar hafnia films the argon lay within the voids [135], and that in tantala films the argon-filled voids only formed on annealing [136]. Nothing was known about the form of argon within the voids or, beyond initial formation, to what extent they were affected by annealing.

From this point onwards, I will be using 'bubbles' rather than 'gas-filled voids' to refer to the noble gas areas.

Bubbles of noble gas are well documented, albeit not in GW mirror coatings. They mainly arise from ion implantation and have been investigated in a variety of materials [137], including nickel [138, 139, 140], aluminium [141, 142, 143], copper [144, 145], gold [145, 146] and silicon [147, 148]. They are of particular interest for materials used in nuclear reactors, where fission products implant into the protective cladding. These fission products are often noble gas nuclei; alpha particles that have captured electrons to become helium for example[149], although xenon [150] and krypton [151] are also common. Their implantation has been extensively studied in steels [152, 153], uranium dioxide [154, 155] and zirconium alloys [156, 157]. The main effect these atoms have on their surroundings is local strain, as noble gases do not interact strongly chemically with their surroundings, essentially making them insoluble. When there are mutiple gas atoms within a diffusion length of each other, the strain field interaction causes them to cluster together. If there are enough atoms, bubbles will start to form [158, 159, 160]. The larger the bubbles, the easier it is to collect other atoms. This can weaken the cladding [161], as well as the reactors' internal components [162] by causing swelling that can lead to cracking. Implantation can also affect the fuel, in both its performance [149]

and storage [163]. While the bubbles in GW mirror coatings are not formed in the exact same way, the process is analagous. Once present, the bubbles can be treated in the same manner when performing spatial quantification.

The effects of the bubbles on the mirror coatings is not yet fully known since their presence was only recently confirmed. Possible effects will be discussed later in this chapter.

The work in references [3] and [4] make up the majority of this chapter. Specifically, reference [3] covers the work on the argon bubbles in titania-doped tantala and reference [4] the work on xenon bubbles in Zircaloy-4. Both of these papers were written by me, alongside Ian MacLaren, with the other authors offering comments. The argon bubbles have been reanalysed to make the processing consistent with that of the other samples, as the technique naturally developed through repeating the process many times (specifically recording different bubble types and measuring neighbour distances). The discussion of these individual samples also comes from these papers. The graphs in this chapter have been redrawn but derive from the same code as was used for the papers. Section 4.1 is work originally written for reference [5].

4.1 Creating a Standard

Having an empirical standard for MLLS fitting (section 3.4.1) means that it is possible to produce fit coefficients that are quantitatively meaningful, thereby enabling absolute spatial quantification for our samples. However, when wanting to quantify noble gas atoms this is no simple request; as it is difficult to perform quantitative EELS on noble gases in the microscope, there are few standards available. Fortunately creating a semi-empirical standard is relatively straightforward. It is easiest to showcase this process using an example; here argon is used, but the same process was also used for the kyrpton and xenon standards. The individual spectra that contribute to the final semi-empirical argon standard can be seen in fig. 4.1.

The first step was finding data for the edge of interest. There was no need for them to be quantitative or collected under the same experimental conditions as the data of interest – it was merely a starting point. The EELS Atlas does contain $L_{2,3}$ edge data for argon gas [121], which will give the correct edge shape for the pressure used in that experiment. That spectrum was used for the shape of the edge but there were problems with the background. To deal with this, the signal was extracted by taking the background in front of the edge and saving both the signal and the background. Then the deconvolved spectrum was taken and using the elemental quantification plug-in in Digital Micrograph (DM) a model of the edge of interest was extracted. It was important to use a large range to get a complete tail after the edge, and to start as close as possible to the edge onset. In this example the range was 245 eV to 1000 eV. The next step was scaling the edge by a Hartree-Slater cross-section that had been calculated for our experimental conditions [108]. The EELS Atlas spectrum had a dispersion of 0.3 eV compared with 0.5 eV for our data. This means that the EELS Atlas

spectrum needed to be rebinned. While theoretically this should have been trivial, some non-linearties in the disperison meant it was necessary to used a DM script [115]. The scaled spectrum was then spliced with the EELS Atlas model – the first for the tail shape and the latter for the edge shape. Although not required here, it may be necessary to experiment with where to splice, the example case was spliced over 15 eV from 300 eV, to make sure the L_2 edge was included. If there is Fourier noise the splice point and range can be adjusted.

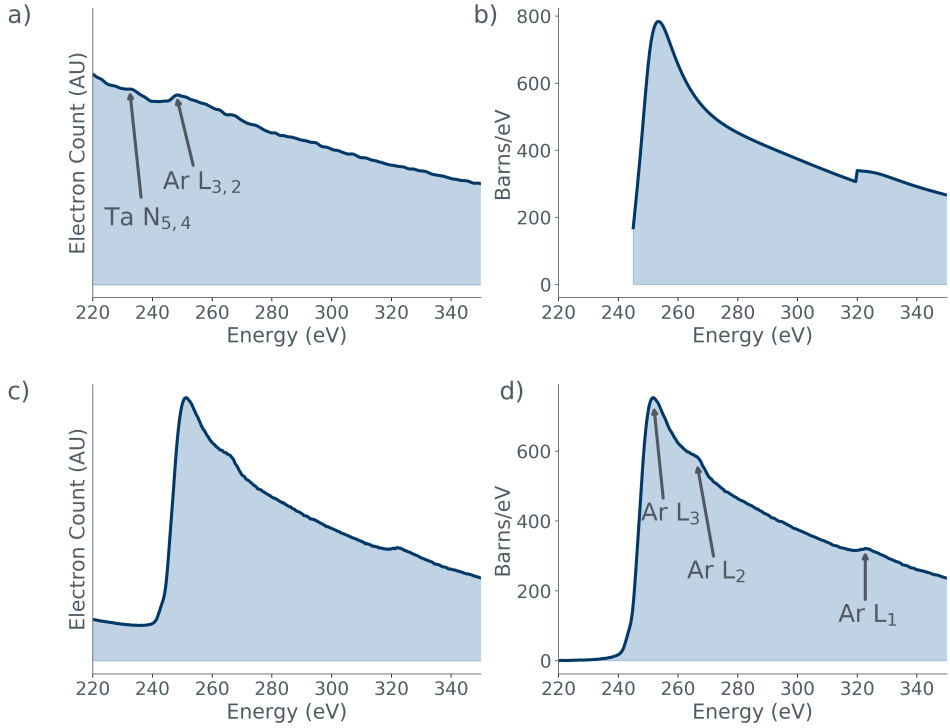

Figure 4.1: Creation of an argon semi-empirical standard using an argon edge collected from the tantala sample in section 4.3.1: a) the argon $L_{3,2}$-edge in an argon rich area in the deconvolved SI. The edge is visible but weak and its shape is not clear. It is also affected by the immediately preceding tantalum $N_{5,4}$ edge; b) the Hartree-Slater cross-section for the argon edge; c) the argon edge extracted from the deconvolved SI using elemental quantification in DM, with the background still present; d) the background-subtracted extracted argon edge, scaled by the Hartree-Slater cross-section, i.e. the final semi-empirical standard.

Figure 4.1 shows the stages of creating a semi-empirical standard for the argon $L_{3,2}$ edge. A Hartree-Slater cross-section (fig. 4.1.b) was used as it is well established that they are reasonably accurate for the $L_{2,3}$ edges of lighter

elements [105, 107], so long as the measurement is made at a high enough energy above the edge onset. Creating the argon standard (fig. 4.1.d) from the fig. 4.1.a presented a particular challenge due the presence of the tantalum $N_{5,4}$ edge at 229 eV, right above the argon $L_{3,2}$ edge. This can affect the pre-edge background. The difference between the initial edge in the experimental data (fig. 4.1.a) and the final semi-empirical standard (fig. 4.1.d) is stark, with the L_2 and L_1 edges now visible.

4.2 General Notes

4.2.1 Bubbles Nomenclature

Whole bubbles are those contained entirely within the STEM SI. Partial bubbles are those that sit at the edge of the image. In some cases, their diameters could be taken from the image and in others they were estimated based on what was in the image and other bubbles in the image. The number of atoms within the partial bubbles was calculated by taking the sum of the atoms within the visible bubble, and then estimating what fraction of the bubble was visible, and dividing the sum by this fraction. Multiple bubbles are those that overlapped and could not be easily separated. The total number of atoms was summed and divided by how many bubbles seemed to be joined together. Depending on their point of contact, the diameters could either be taken directly from the image or estimated based on their shape and other bubbles in the image. The whole bubbles in the image were used as a rough guide for the partial and multiple bubbles by giving an expected range of diameters and sizes. An example of a bubble fit coefficient map showing the different bubble types can be seen in fig. 4.2. All of the annotated bubble fit coefficient maps for every sample in this chapter can be found in the Appendix.

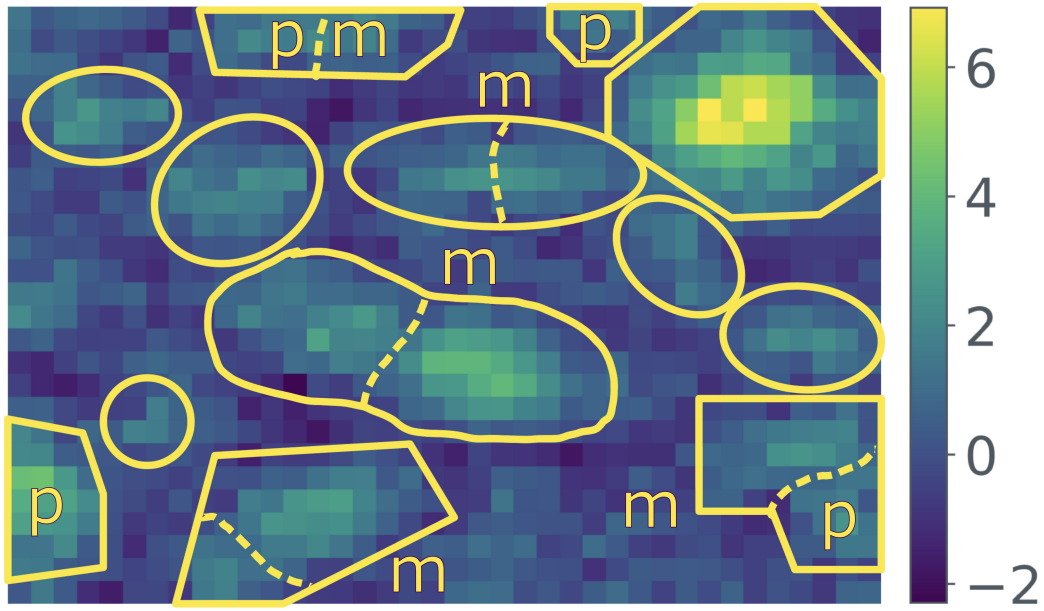

Figure 4.2: An example bubble fit coefficient showing the bubbles taken from the image. 'p' denotes a partial bubble, lying at the edge of the image and'm' denotes a multiple bubble, with a dotted line showing how they were divided.

4.2.2 Collection Conditions

The data were collected using a JEOL ARM200F at 200 kV and its Gatan GIF Quantum ER spectrometer. HAADF images were taken in order to show the areas of interest and see the physical structure of the films. Electron energy loss spectroscopy (EELS) was performed at both high- and low-loss in order that the composition of the sample could be determined with high spatial resolution. A convergence angle of 29 mrad was used for the probe and the collection angle into the spectrometer was 36 mrad. DualEELS was used to collect the low loss and the high-loss almost simultaneously. The high- and low-loss ranges were chosen such that the relevant edges could be investigated.

4.2.3 Accounting for Drift

Sample drift is a problem in the spectrum images (SIs) as they are taken by the beam rastering across the sample; along the fast scan direction drift is minimal but between rows of pixels it can be over 1 nm. Therefore, drift was compensated for by using only the horizontal diameter of the bubbles from the fit coefficient maps. The diameter was taken using the line-profile tool in DM, which produces a histogram of the pixels that the line passes through. The diameter was determined by taking the point at where the pixel values fall to background value, and adding an extra border of one pixel (pixels ranged from 2-6Å) each side of the bubbles, which had diameters ranging from 6 Å to 60Å. This was done after comparing the fit coefficient maps to the HAADF images. The outer edge of the bubbles appears to have very low noble gas content, but this is because the path through the outer edge is very short. Not including a border with low gas atom numbers gives unrealistically high densities.

DM's in-built drift correction was not used as for many of the datasets there was no fixed point to which to tie the drift correction.

4.3 In Gravitational Wave Interferometer Mirror Coatings

The work in this section starts with a titanium dioxide doped tantala pentoxide sample made via ion beam assisted deposition in an argon environment, which was the first sample investigated. The other two samples that make up the work in this section are zirconium dioxide doped tantalum pentoxide, one deposited in argon and one in oxygen. While these samples yielded lower quality EELS data, they still provide a valuable comparison to the inital work. Zirconium dioxide doped tantalum pentoxide was one of the potential mirror coatings that was suggested for aLIGO, before titanium dioxide doped tantala pentoxide was chosen.

4.3.1 Argon in Titanium Dioxide Doped Tantalum Pentoxide

Argon in titanium dioxide doped tantalum pentoxide was the first sample to be investigated. It was originally made when investigating new coatings for aLIGO and this composition (TiO_2-doped Ta_2O_5) was eventually chosen to be the high-refractive index layer for the aLIGO mirrors. It is nominally 25% titania and 75% tantala, however work by Isa et al. using similar methods to the work in this chapter showed it was actually closer to 10% and 90%, respectively [165].

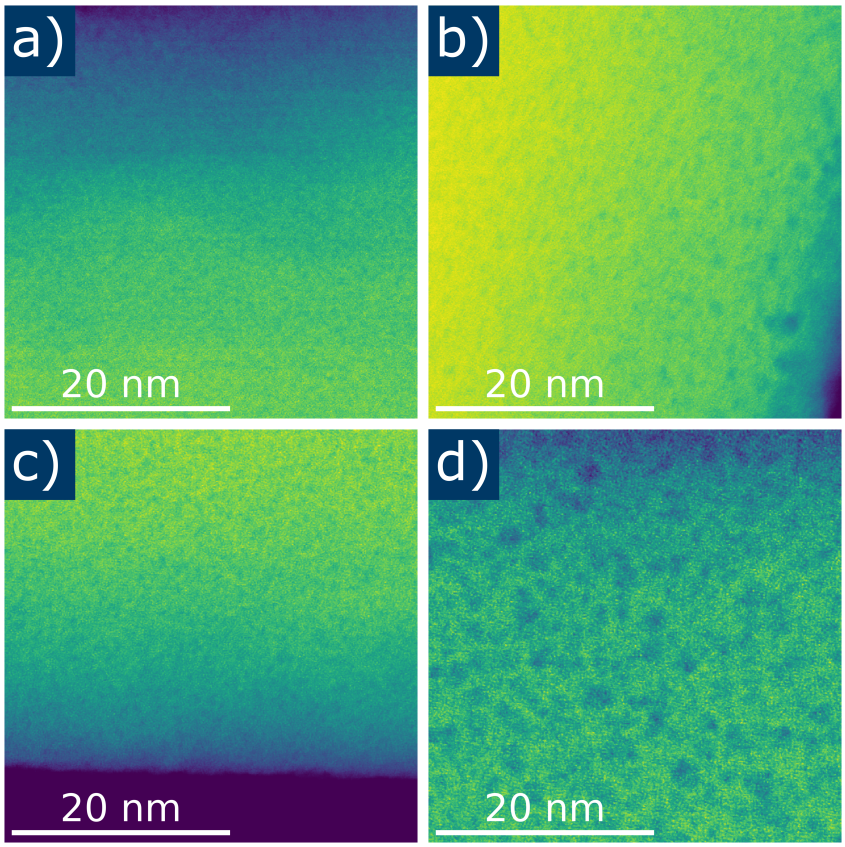

Figure 4.3: Example HAADF images taken from the four titania-doped tantala samples: a) the as-deposited sample, with no clear argon bubbles; b) the sample annealed at 300 °C, with some bubbles visible as darker regions in HAADF STEM, due to their reduced density with respect to the surrounding matrix; c) the sample annealed at 400 °C, with the bubbles becoming clearer and; d) the sample annealed at 600 °C, with clear bubbles visible.

4.3.1.1 Sample Preparation and Data Collection

There are four argon in titanium dioxide doped tantalum pentoxide samples, three annealed in air for 24 hours at 300, 400 and 600 °C respectively and one as-deposited sample. They were all created at CSIRO using ion beam assisted deposition onto fused silica substrates in an argon environment. The samples were prepared for transmission electron microscopy observation using a conventional cross-section preparation method [166].

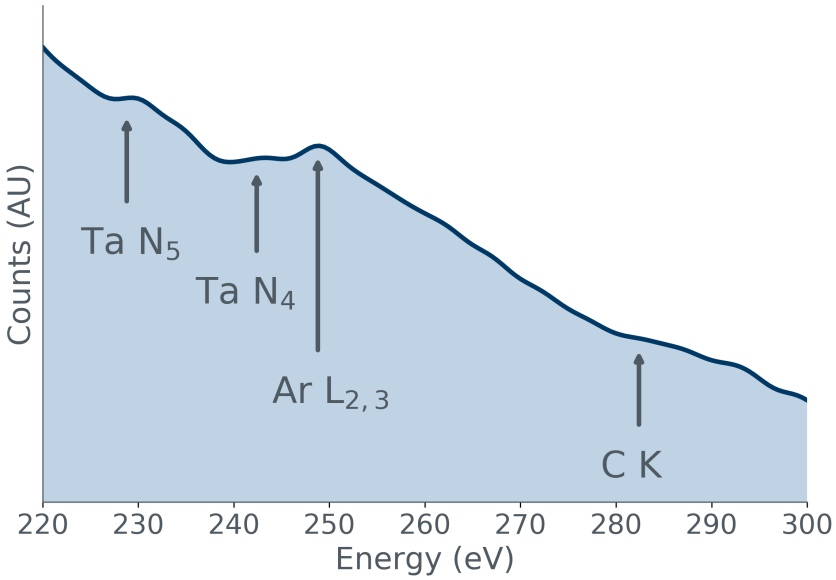

Figure 4.4: An example spectrum from the titanium dioxide doped tantalum pentoxide sample annealed at 600 °C, showing a small increase at the carbon K edge, at 284 eV, following the argon and tantalum edges.

The high- and low-loss EELS acquisition times were 89.5028 ms and 497.238 μs respectively, with energy offsets of 0 eV and 180 eV. The pixel size ranged from 0.2 nm to 0.6 nm.

The MLLS fit was taken between 220 eV and 280 eV in order to avoid the carbon K-edge at 284 eV, which can be seen in fig. 4.4. Carbon contamination can vary across the sample so eliminating it entirely from the fit prevents it skewing the results. The datasets also all had two 'matrix' spectra in the fit to account for the thickness gradient in the sample. Selecting the matrix spectra simply requires taking a qualitative spectrum from an area with no argon. This can be done by eye from the SI and then checked by viewing the spectrum; the Ar-$L_{2,3}$ edge has a strong peak at 250 eV, which is above the edge onset for the Ta-$N_{4,5}$ edge (226 eV) [167, 165], but below the Ti-$L_{2,3}$ (454 eV).

4.3.1.2 Results

Figure 4.3 shows HAADF images of the four titania doped tantala samples. One is the unannealed sample (fig. 4.3.a) and the other three were annealed at 300 °C (fig. 4.3.b), 400 °C (fig. 4.3.c) and 600 °C (fig. 4.3.d). In the 400 °C and 600 °C samples, lower-density (i.e. darker) rounded regions are seen, all a few nanometres in diameter. They are also visible in the sample annealed at 300 °C but unfortunately all the data from this sample is poor quality due to sample charging. While it would have been possible to recoat the sample to prevent this charging after attempting to analyse the 400 °C annealed sample it was deemed unnecessary; the 400 °C sample had very poor SNR and it was not possible to perform quantification on its datasets. No bubbles could be seen in the unannealed sample but there was still argon present, as can be seen in Figure 4.5.

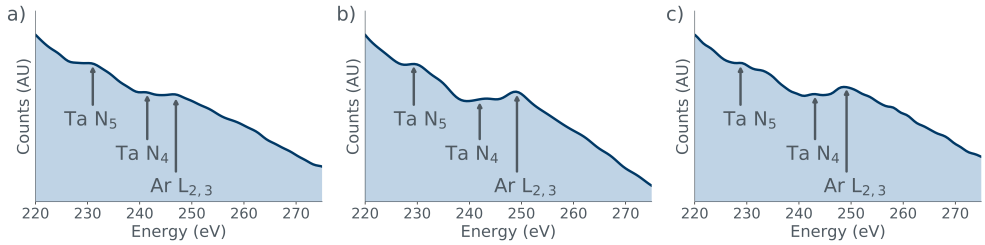

Figure 4.5: EELS spectra from three samples all annealed at different temperatures, showing clear evidence that argon is present in all three. The relative intensities of the edges are proportional to the amount of the element in the area from where the spectra were taken, which was from apparent bubble regions for all three. The tantalum N_5 (226 eV onset) and N_4 (238 eV onset) edges and the argon $L_{2,3}$ edge (248 eV onset) are visible in all of the samples; a) the as-deposited sample; b) 400 °C annealed sample; c) the 600 °C annealed sample.

While the bubble sizes can be hard to determine in the unprocessed datasets it does appear that the average size of the bubbles grows with annealing temperature. The bubble diameters are taken from the argon standard coefficient images not the initial HAADF image, even though they are visible in the latter because the edges are much easier to define when it is only the

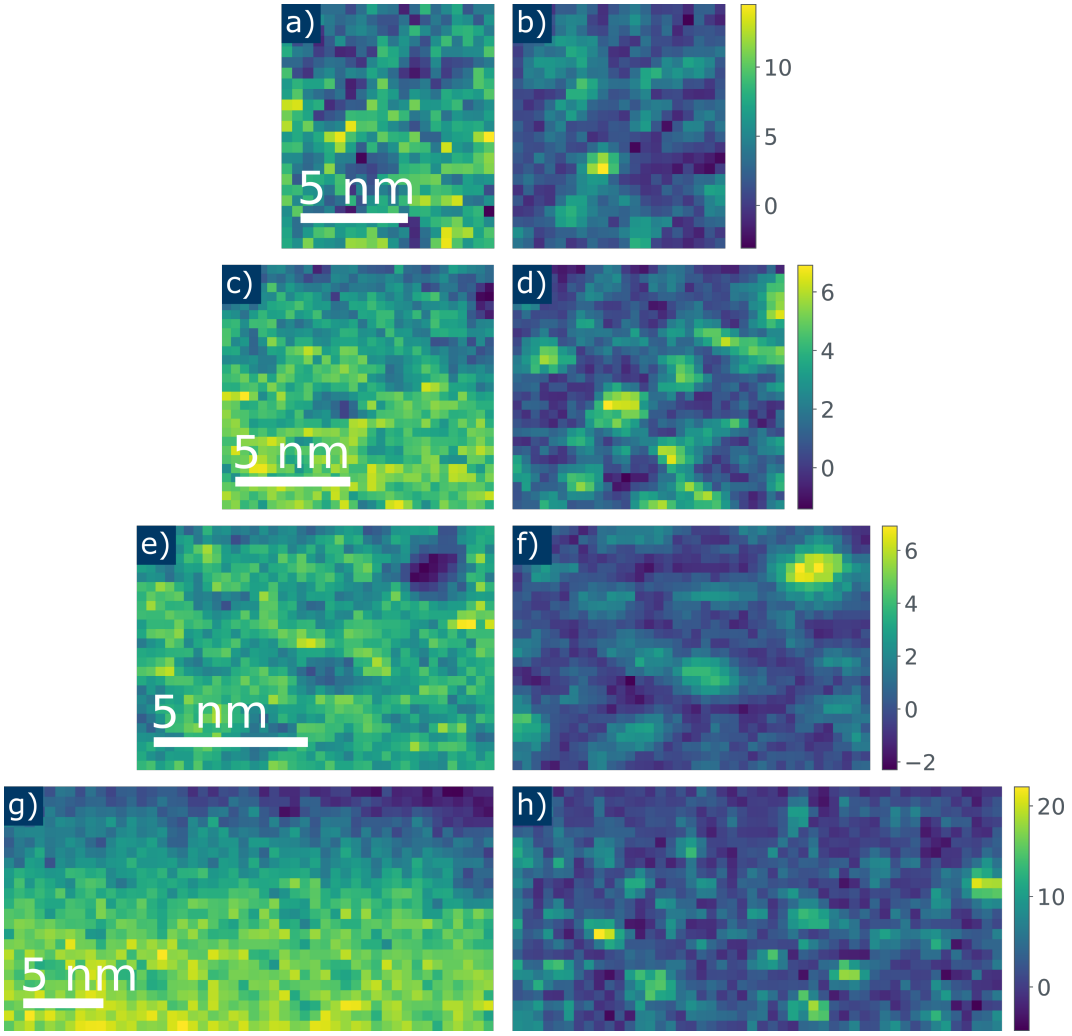

Figure 4.6: HAADF images (left hand side) and their corresponding argon fit coefficient images (right hand side) for the argon bubbles in the titania-doped sample annealed at 600 °C. The intensity of the bubble maps is in number of argon atoms.

argon contributing to the pixel intensity. It is likely that the lamella intersects some bubbles, which would release the argon but there do not seem to be any 'empty' bubbles (i.e. hemispherical voids) visible in the HAADF images and not in the bubble maps. This is not unexpected as although HAADF imaging is sensitive to density it is more sensitive to Z-contrast.

All data from this point onwards is from the 600 °C annealed sample, as there were not enough bubbles present in the 400 °C annealed sample to give

a reliable distribution of diameters. The results of analysing the bubbles in titania-doped tantala sample annealed at 600 °C can be seen in fig. 4.7.

The mean bubble diameter was 22.7 Å, very close to the mean diameter for just the whole bubbles, which was 23.5 Å, which suggests the multiple and partial bubbles' diameters were slightly underestimated due to user error. As would be expected, the number of atoms within the bubbles increases with bubble diameter.

The majority of the bubbles have densities and pressures suggestive of gas form, but some of the smallest may be a high density fluid. A few fall above the density of solid argon. The mean density for all of the bubbles was 959 kg/m^3 and for the whole bubbles it was 885 kg/m^3. This is much less dense than the Ta_2O_5 surrounding them, which has a density of ∼7750kg/m^3 [165]. This explains the HAADF contrast of fig. 4.6, where the bubbles are darker than the surrounding Ta_2O_5 matrix. Some of the bubbles, however, did start to approach this density, with the densest bubble being 6039 kg/m^3, although the densest whole bubble was less dense than this at 5059 kg/m^3. No background argon content subtraction was performed for this sample as the argon values in the matrix pixels of the bubble maps were low. How the densities and pressures were calculated will be discussed in section 4.5.1, as will possible errors.

The sample thickness, based on t/λ calculations, ranges from 250 Å to 400 Å. Any separation of bubbles in the z-plane (i.e. down into the sample) is more significant for the smaller separations, as bubbles separated by distances less than their diameter would likely have coalesced. Figure 4.8.a shows the measured separations in the $x - y$ and fig. 4.8.b shows how the distribution would change if all bubbles had a separation of 100 Å in the z plane. This example separation of 100 Å was chosen as it is less than half the thickness of the sample. It is very clear that the adjustment for possible z plane separation

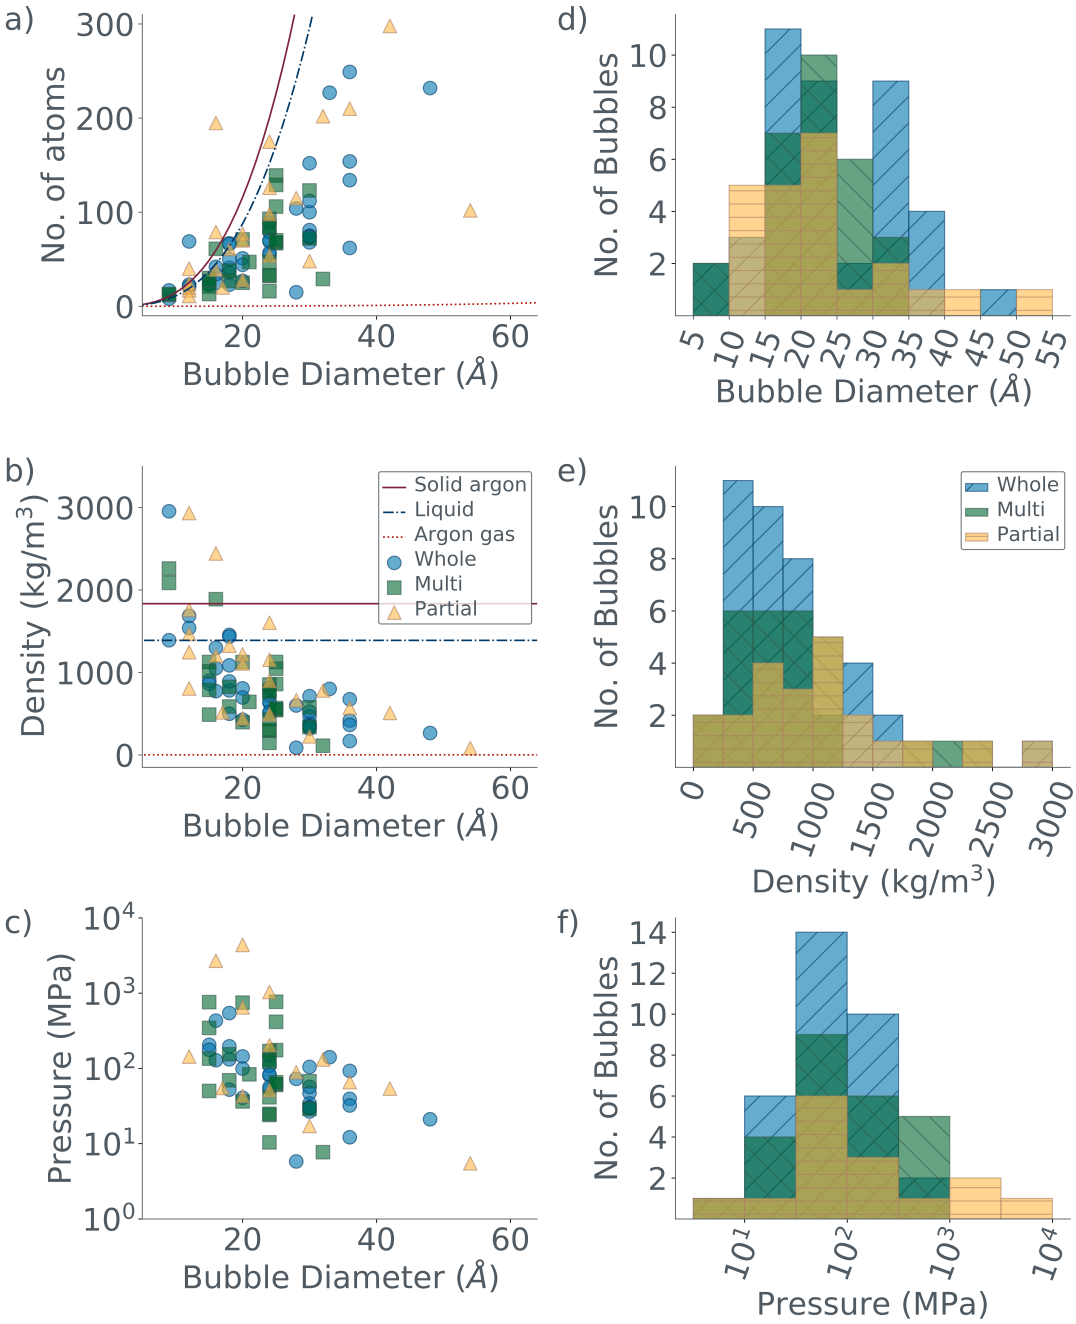

Figure 4.7: Analysis of the argon bubbles in the TiO_2-doped Ta_2O_5 sample annealed at 600 °C, showing the whole (blue circles), multiple (green squares) and partial (yellow triangles) bubbles, as well as the values for solid (arranged in the tightest possible face-centered cubic (FCC) packing) , liquid and gaseous (both at standard pressure and temperature): a) number of argon atoms in bubble against their diameter; b) distribution of bubble diameters; c) density of the argon bubbles; d) distribution of bubble densities; e) bubble pressure against their diameter; f) distribution of bubble pressures.

		Diameter (Å)	Density (kg/m^3)	Pressure (MPa)
	Mean	23.5	885	105
Whole (41)	Minimum	9	131	5.43
	Maximum	48	5059	5.81
	Mean	22.7	1282	641
Partial (18)	Minimum	12	81	5.48
	Maximum	54	6031	4380
	Mean	21.5	804	182
Multiple (8)	Minimum	9.0	112	7.69
	Maximum	32.0	2259	766
All (244)	Mean	22.7	959	241

Table 4.1: Statistical breakdown of the argon bubble properties in the TiO_2-doped Ta_2O_5 sample annealed at 600 °C, showing the mean, maximum and minimum values for the three bubble types (whole, partial and multiple). The number in brackets the number of each type of bubble. The final line gives the mean values for all of the bubbles, irrespective of type.

had the greatest effect on the distribution at smaller separations, and changed the shape of the distribution considerably. Where the distribution is truly centred about is unclear, but it is its shape will be closer to fig. 4.8.b, with a longer tail than either shown distributions to account for separations of up to 400 Å the z plane. The mean separation of the original distances in only the $x - y$ plane (i.e. plane of the lamella) was 36.9 Å

These results will be discussed fully in section 4.5.

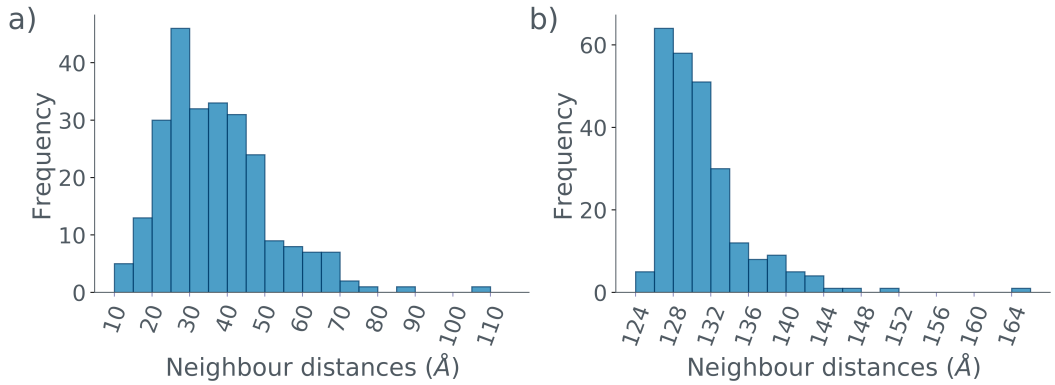

Figure 4.8: Distributions of the neighbour distances for the argon bubbles in the TiO_2-doped Ta_2O_5 sample annealed at 600 °C. Neighbouring bubbles are defined as any bubbles between which a straight line can be drawn without passing through another bubble: a) distances only in the $x-y$ plane, these distances were only taken in 2D as it was not possible to image the sample in 3D; b) an adjusted distribution assuming all bubbles also had a separation of 100 Å in the z plane, as an example of how including a separation in the z-plane affects the distribution of nearest neighbour distances.

4.3.2 Argon in Zirconium Doped Tantalum Pentoxide

The argon in zirconium doped tantalum pentoxide sample was made as a potential aLIGO coating, and was doped with zirconia.

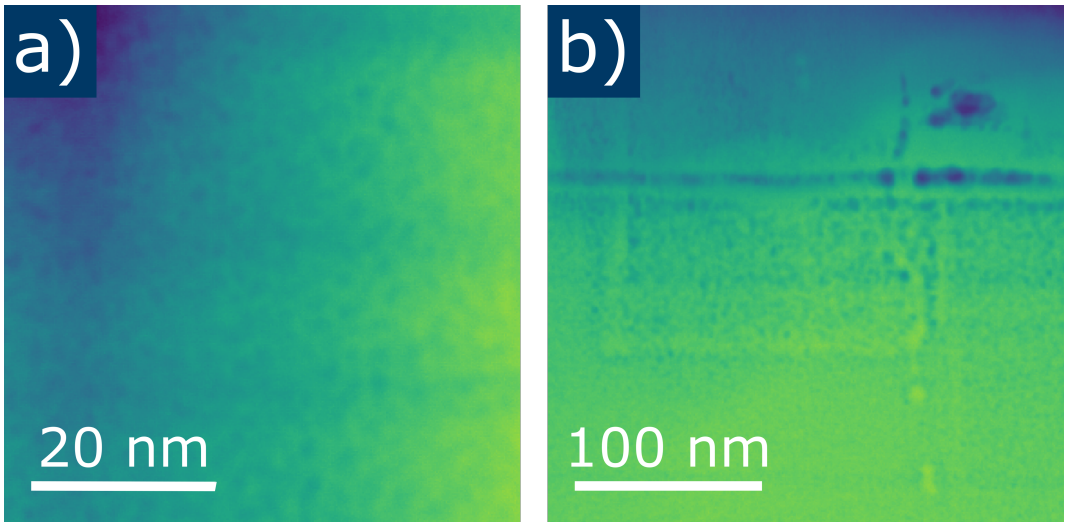

Figure 4.9: Example HAADF images taken from the two zirconia-doped tantala samples: a) the sample deposited in argon, with visible bubbles and; b) the sample deposited in oxygen, showing very clear damage from the beam – the darker areas streaking horizontally across the sample.

4.3.2.1 Sample Preparation and Collection Conditions

This sample was made using ion beam assisted depositon in argon, onto fused silica substrates. It was made at Strathclyde by Stuart Reid's group. The exact stochiometry of this sample is unknown, and although it could be determined using similar techniques it was not deemed necessary as the interest was in the bubbles.

The SI pixel size ranged from 0.25 nm to 1 nm. The high- and low-loss collection times were either 0.000495 s and 0.0495 s respectively, or 0.0004878 s and 0.01951 s. The dispersion was 0.5 eV/channel.

4.3.2.2 Results

Bubbles can be seen in the HAADF images for the sample deposited in argon (fig. 4.9) but it can be seen in the bubble maps (fig. 4.10) that there is a large amount of argon not in the bubbles. It is not entirely clear why the argon has not coalesced (likely due to different annealing conditions) but it presents a challenge for calculating the bubble properties. This background increases the number of atoms in each pixel and creates a very high number of atoms per bubble. The argon is likely spread through the sample and, as a consequence, many bubble values have been lost when calculating their pressures due to the nature of the van der Waals equation eq. (4.1). At extreme values the van der Waals equation 'flips' and will give negative or unphysically high pressures. This is clearly visible in table 4.2, where the majority of the pressures for all of the bubbles types are negative. In order to account for this, a background argon value was taken from each image. A rectangular region of interest (ROI) was drawn and summed on the bubble map in DM. This was then divided by the area of the rectangle, to get the mean argon value per pixel. This value was

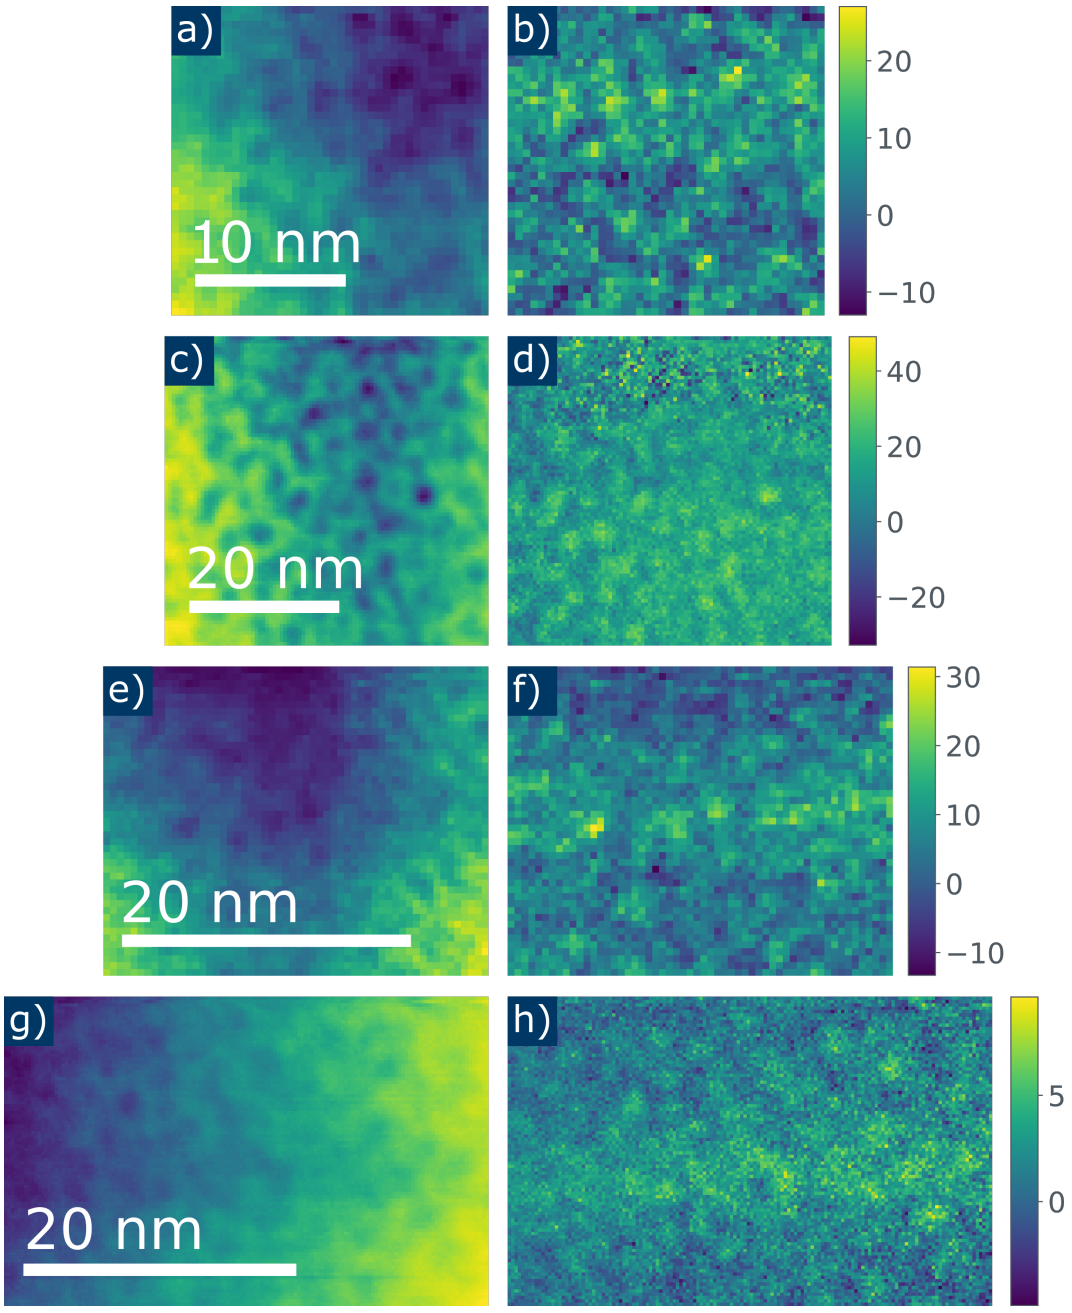

Figure 4.10: HAADF images (right hand side) and their corresponding argon fit coefficient images (left hand side) for the argon bubbles in the zirconia-doped sample. The intensity of the bubble maps is in number of argon atoms.

then used to subtract the background argon from all the bubbles before any further calculations were done. A mean value for the background argon content, rather than a minimum, was chosen as there was a lot of variation between pixels in the argon content, and the minimum value was often very small and not representative of the majority of the pixels; and therefore, subtracting by the minium value had little effect on the data.

Performing background subtraction meant that far fewer bubbles gave unphysical pressure calculations, and placed their densities mostly in the gaseous region. How much the data were affected by background subtraction can be seen in fig. 4.11 and table 4.3. Some also have a density above solid argon, but it is possible they are high-density liquids rather than true solids. Some of the bubbles have been 'lost' in background subtraction as they ended up with negative bubble counts and therefore were discarded before the density and pressure calculations.

On average it appears that the diameters for the partial and multiple bubbles were under- and overestimated, respectively. This lead to the multiple bubbles driving up the mean value for pressure when compared with the whole bubbles. The mean across all three types for the background subtracted bubbles was 29.3 Å, with the whole bubbles having a mean of 27.5 Å. As with the previous sample the number of atoms in the bubble goes up with diameter, although the exact relationship is hard to define. The pixel size, which ranged from 2.5 to 10 Å was the biggest source of uncertainty in the measurement. The density decreases with increasing diameter and the mean density across all all three types for the background subtracted bubbles was 960 kg/m^3 – a little lower than the mean for the whole bubbles which was 1009 kg/m^3. The bubbles are all well below the density of the surrounding tantala ($\sim$7750

kg/m^3 [165]). A breakdown of the results can be found in table 4.3 and the results themselves are visible in fig. 4.11.

The bubble separation was calculated using the same method as the previous sample and has the same limitations. This sample's thickness ranged from 500 Å to 700 Å, and an additional separation of 100 Å was chosen again. The mean separation of the original 2D distances was 67.3 Å.

These results will be discussed fully in section 4.5. The original, pre-background subtraction data have been included in this section to fully illustrate how much background subtraction can improve the data.

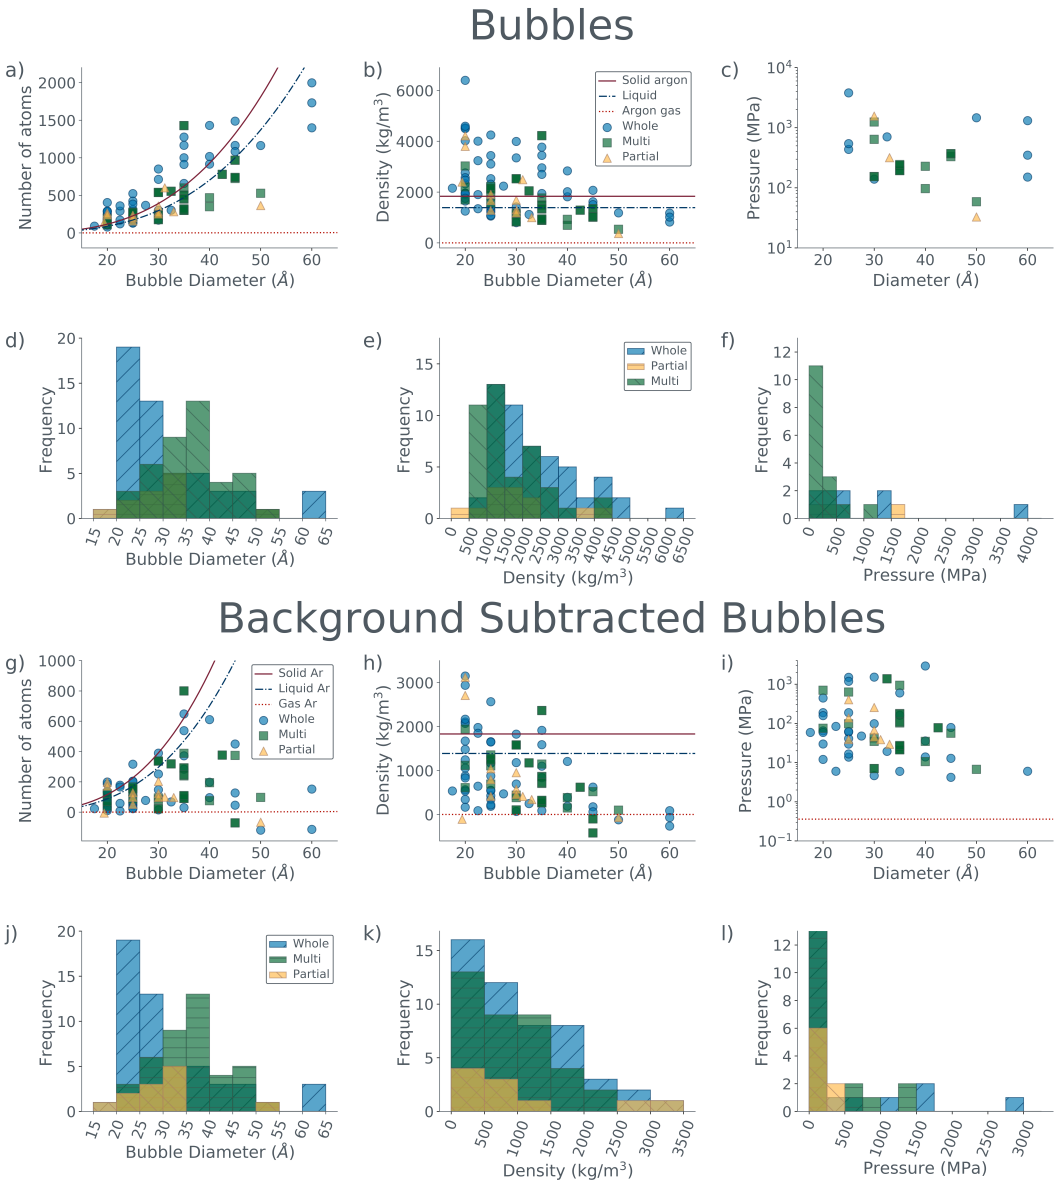

Figure 4.11: Analysis of the argon bubbles in the ZiO_2-doped Ta_2O_5, showing the whole (blue circles), multiple (green squares) and partial (yellow triangles) bubbles, as well as the values for solid (arranged in the tightest possible face-centered cubic (FCC) packing), liquid and gaseous (both at standard pressure and temperature), comparing initial data with background subtracted data: a) number of argon atoms in bubbles versus their diameter; b) density of bubbles versus their diameter; c) pressures of bubbles versus their diameter; d) histogram of bubble diameters; e) histogram of bubble densities; f) histogram of bubble pressures; g-l) as previous but for the background-subtracted bubbles. An explanation of the three bubble types can be found in the main text.

		Diameter (Å)	Density (kg/m^3)	Pressure (MPa)
	Mean	29.2	2376	-226
Whole (53)	Minimum	17.5	798	-9600
	Maximum	60.0	6398	3810
	Mean	28.2	1941	-2932
Partial (12)	Min	19.4	369	-2470
	Maximum	50.0	4212	1560
	Mean	33.7	1621	-1922
Multiple (41)	Minimum	20.0	535	-2230
	Maximum	50.0	4220	1250
All (106)	Mean	30.8	2034	-220

Table 4.2: Statistical breakdown of the argon bubble properties, showing the mean, maximum and minimum values for the three bubble types (whole, partial and multiple) in the zirconia-doped sample. The number in brackets lists the number of each type of bubble. The final line is the mean values for all of the bubbles, irrespective of type. This table shows how badly the final pressures are affected by the presence of the background argon content, with the majority of the bubbles giving negative pressures.

		Diameter (Å)	Density (kg/m^3)	Pressure (MPa)
	Mean	27.6	1009	286
Whole (50)	Minimum	17.5	64	4.20
	Maximum	60.0	3150	2930
	Mean	26.9	1082	127
Partial (10)	Minimum	19.4	341	29.3
	Maximum	33.0	3121	403
	Mean	32.4	862	831
Multiple (37)	Minimum	20.0	99	6.70
	Maximum	50.0	2364	17000
All (97)	Mean	29.3	960	486

Table 4.3: As in table 4.2, but after background subtraction, of particular note are the pressures, which are no longer overwhelmingly negative, and now give realistic values.

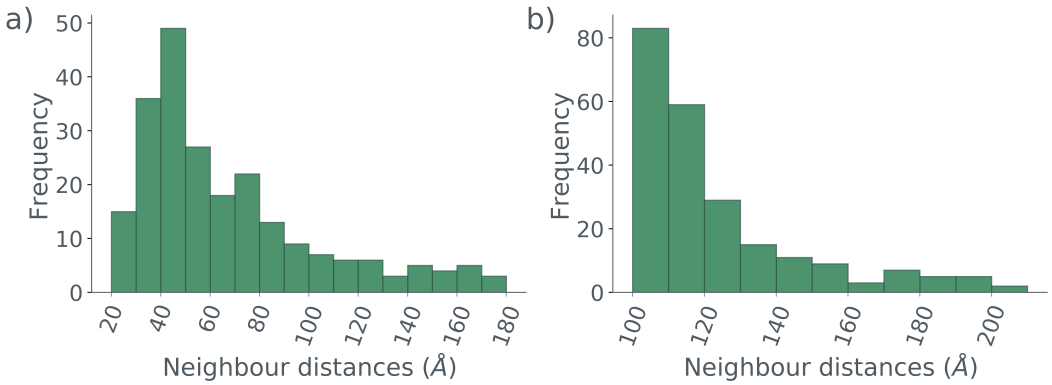

Figure 4.12: Distributions of the neighbour distances for the argon bubbles in the zirconia-doped sample. Neighbouring bubbles are defined as any bubbles between which a straight line can be drawn without passing through another bubble: a) distances only in the $x-y$ plane, these distances were only taken in 2D as it was not possible to image the sample in 3D; b) an adjusted distribution assuming all bubbles also had a separation of 100 Å in the z plane

4.3.3 Oxygen in Zirconium Doped Tantalum Pentoxide

The oxygen in zirconium doped tantalum pentoxide sample was made in the exact same way as the previous sample, except it was deposited in oxygen not argon. It was also made by Stuart Reid's group and its exact stochiometry has not been determined.

There was hope this sample could be used as a control for the argon samples, however imaging it immediately damaged the sample. The damage from the beam rastering across the sample can be seen very clearly in fig. 4.9. There are not any bubbles visible in the two HAADFs taken for this sample but it is not possible to make definitive statements. The oxygen bubbles may have existed and been burnt away by the beam or oxygen may not form bubbles. It is possible that the sample contained more oxygen than was stoichiometric, which was metastable under ambient conditions but as soon as the beam provided energy, the excess oxygen escaped.

Despite the lack of information available from this sample the data have still been included to show the limitations of this technique.

4.4 In Nuclear Reactor Cladding

The following samples differ from the samples in the previous section (section 4.3) in that the xenon and krypton were deliberately implanted, rather than being unintentionally trapped during a deposition process. This is because these samples are made of Zircaloy-4, a common cladding material of nuclear reactors. These samples were made with the intent of mimicking fission products implanting in the cladding in order to investigate the effects on the cladding. Here, they are analysed with a view to comparing them with the GW mirror samples. HAADF images of xenon and krypton can be seen in fig. 4.13 and the semi-empirical samples created for the follwoing analysis can be seen in fig. 4.14.

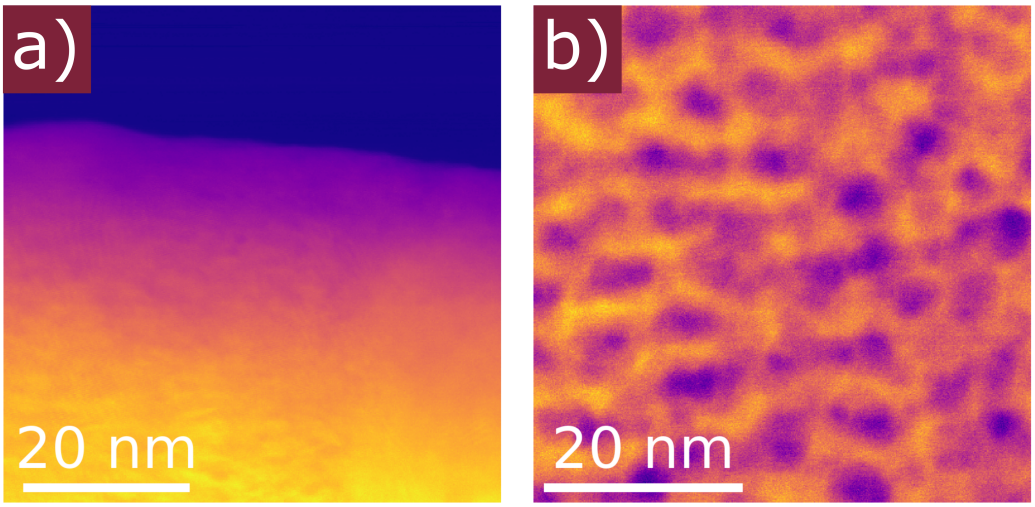

Figure 4.13: Example HAADF images of Zircalloy-4: a) with no clear xenon bubbles visible and; b) with large krypton bubbles clearly visible.

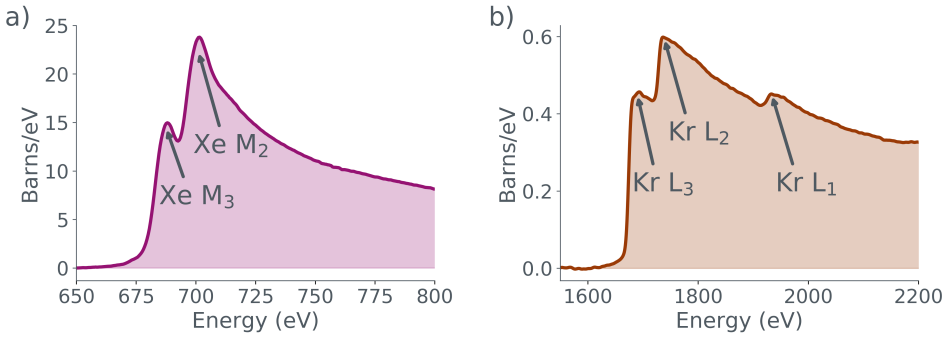

Figure 4.14: The semi-empirical standards for: a) xenon, made using the sample from section 4.4.1 and; b) krypton, made using the sample from section 4.4.1.

4.4.1 Xenon in Zircaloy-4

4.4.1.1 Sample Preparation and Collection Conditions

The Zircaloy was cut into 10×10×2 mm tiles suitable for ion implantation by wire electro-discharge machining. Prior to cutting the larger coupons of material were pickled in hydrofluoric acid. The ion implantation was performed using a 2 MV tandem accelerator. Samples were mounted on a 3 inch diameter silicon wafer using double-sided carbon tape, with the xenon beam rastered over this area during implantation. A dose of 1 MeV $_{129}Xe^{+}$ was implanted at room temperature, ∼25 °C, which gives a peak xenon concentration at ∼200 to 300 nm below the sample surface, with all the xenon contained in the top 500 nm of the sample, as seen in fig. 4.15. Implantation was performed to give xenon fluences of 5×10^{15} ions/cm^2. These calculations were not done at Glasgow but at the National Nuclear Laboratory. Following xenon implantation, samples were exposed to 5 bar deuterium gas for 30 minutes at room temperature. The reason for this exposure was for other research, examining the interactions of hydrogen and xenon (following earlier work on H-He interactions in bubbles in Zircaloy [168]), which will not be discussed here. The preparation of these samples, and the calculations determining

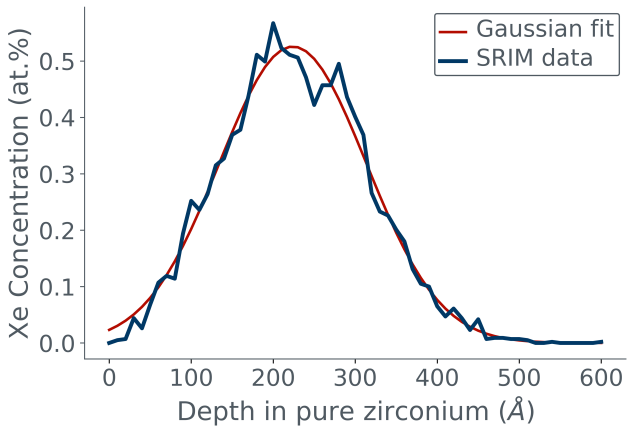

Figure 4.15: The calculated stopping range of the xenon ions in pure zirconium. The amplitude of the Gaussian fit is 0.53, it is centred around 224 Å and its standard deviation is 89.6 Å. The data within this figure were provided by the National Nuclear Laboratory.

Samples were prepared for STEM analysis using a standard focused ion beam liftout procedure using a FEI Helios PFIB instrument [1]. Low energy thinning of the final lamella was performed at 2 kV or 3 kV to improve the surface quality prior to analysis using STEM. In thinner parts of the sample, there were some darker regions that appeared to be bubbles in the HAADF images but contained no gas. These are presumably bubbles that were opened to the surface while the sample was being made, and had their gas escape. The thinnest parts of the sample could not be analysed due to the frequency of these 'empty' bubbles. Thicker areas containing many bubbles were chosen for analysis, focussing on areas that were not so thick that the signal to background ratio in the EELS was too poor to perform signal fitting.

The datasets had a low-loss range of -300 to 724 eV and a high-loss range of 100 to 1124 eV. The SI pixel size was 3 Å for all of the datasets but the high- and low-loss times varied. The MLLS fits were performed between 640-740 eV, over the xenon $M_{4,5}$-edge., and used the xenon semi-emprical standard visible in fig. 4.14.a. Two matrix spectra were taken to account for any thickness effects, as described in section 3.4.1.

4.4.1.2 Results

In the xenon implanted Zircalloy-4 sample, the bubbles are not clearly visible in the HAADF images of fig. 4.16, however using the SI the xenon peak ($M_{2,3}$ at 672 eV) can still be observed in certain areas. This shows that xenon is present in the sample and after MLLS fitting using the xenon standard visible in fig. 4.14, it is clear the xenon has also coalesced into bubbles. These bubbles were then characterised. The background level of xenon in this sample was not high enough to warrant background subtraction, as was done in section 4.3.2.

The mean bubble size across all three types is 21.2 Å, slightly above the mean for the whole bubbles, which is 20.1 Å. The multiple bubbles have a higher mean, suggesting that their diameters were overestimated when splitting two or more bubbles apart. Conversely, the partial bubbles had a lower mean (18.3 Å), suggesting their diameters were underestimated. As would be expected, the number of atoms within the bubbles rose with the diameter of the bubbles. The precision of the diameter measurement is limited by pixel size, which was 3 Å. This is the largest source of uncertainty in the data.

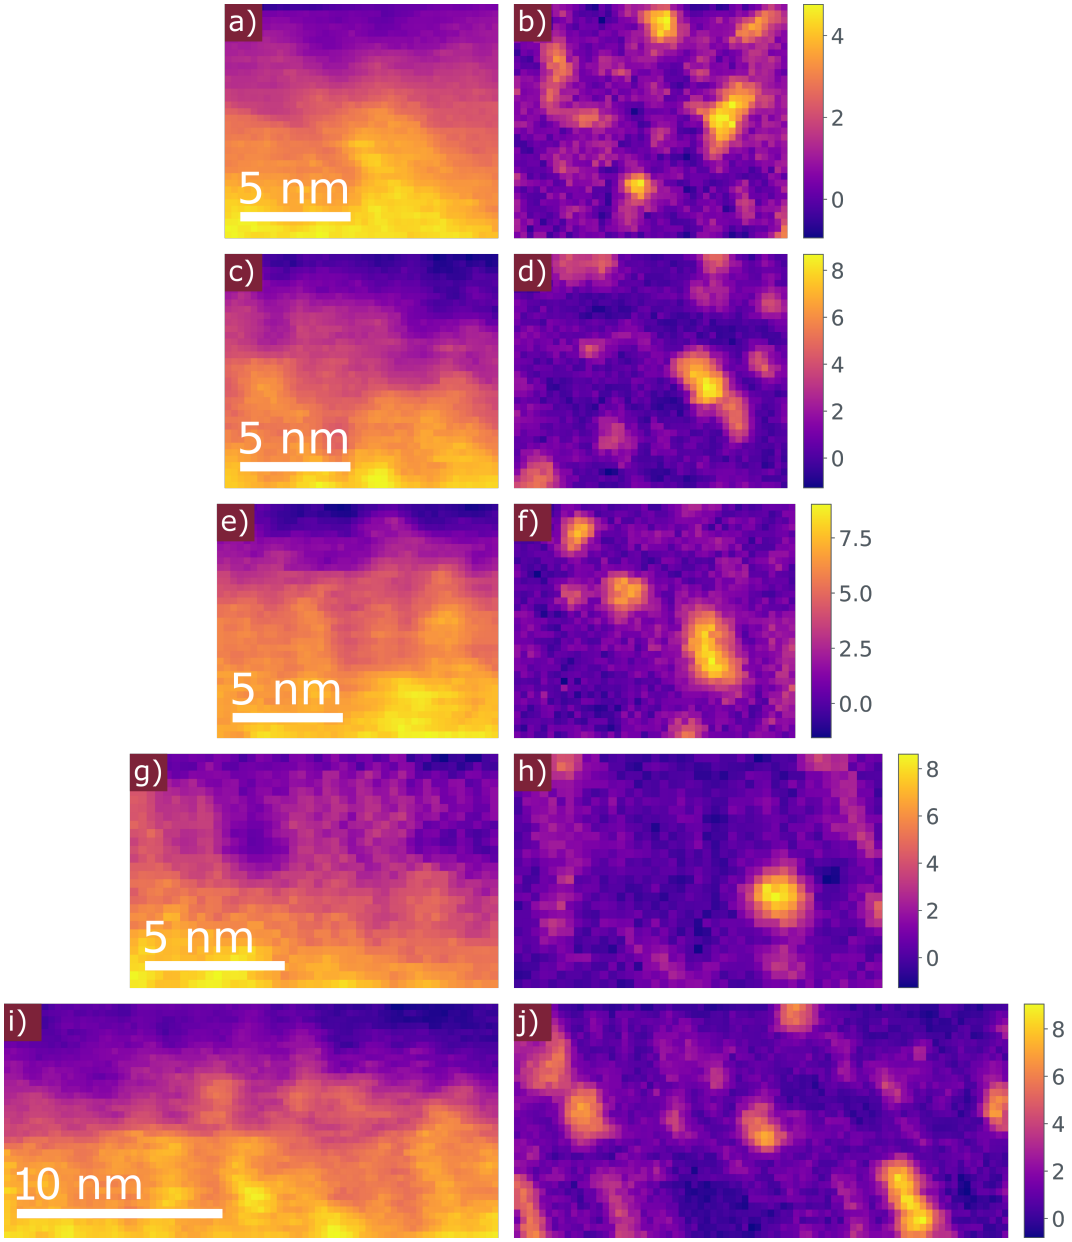

Figure 4.16: HAADF images (images with scale bar) and their corresponding xenon fit coefficient maps (images with intensity scale) for the xenon bubbles in Zircaloy-4. The intensity of the bubble maps is in number of xenon atoms. Continued overleaf.

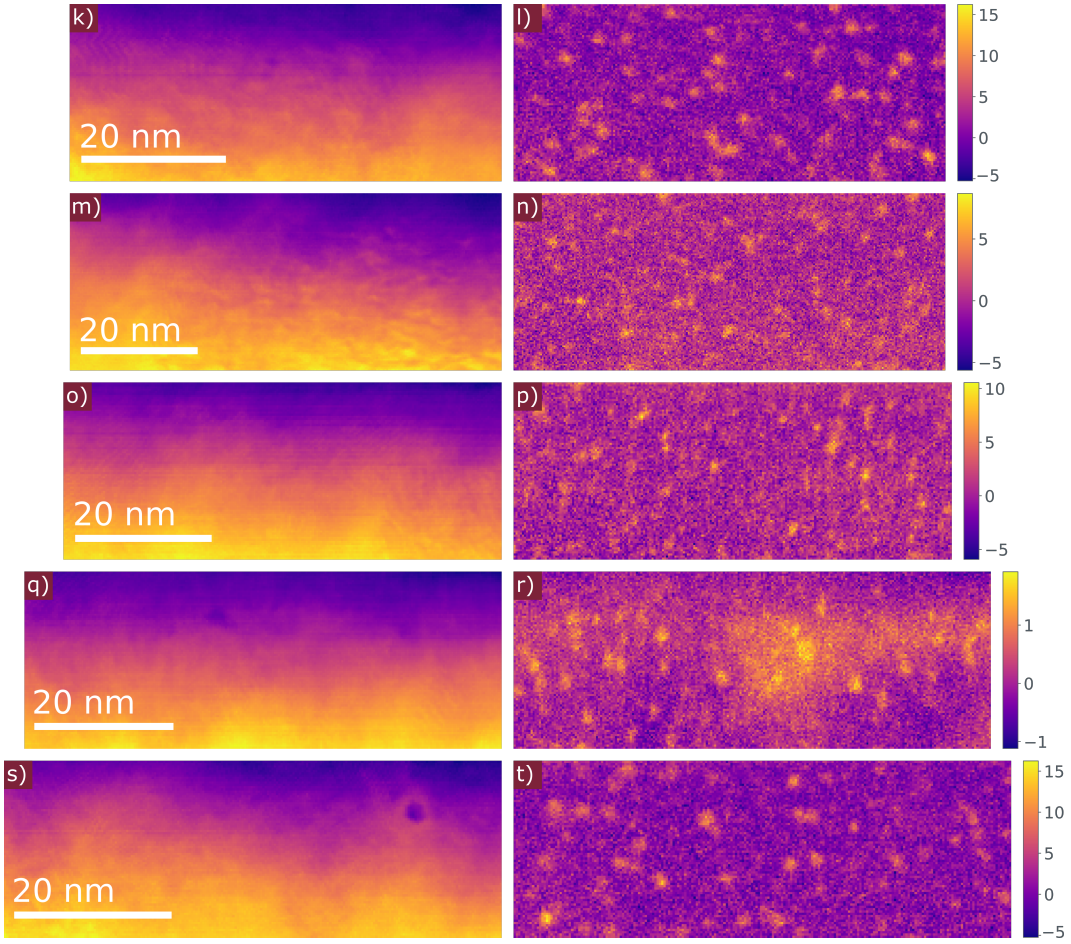

Figure 4.16 cont.: HAADF images (images with scale bar) and their corresponding xenon fit coefficient maps (images with intensity scale) for the xenon bubbles in Zircaloy-4. The intensity of the bubble maps is in number of xenon atoms.

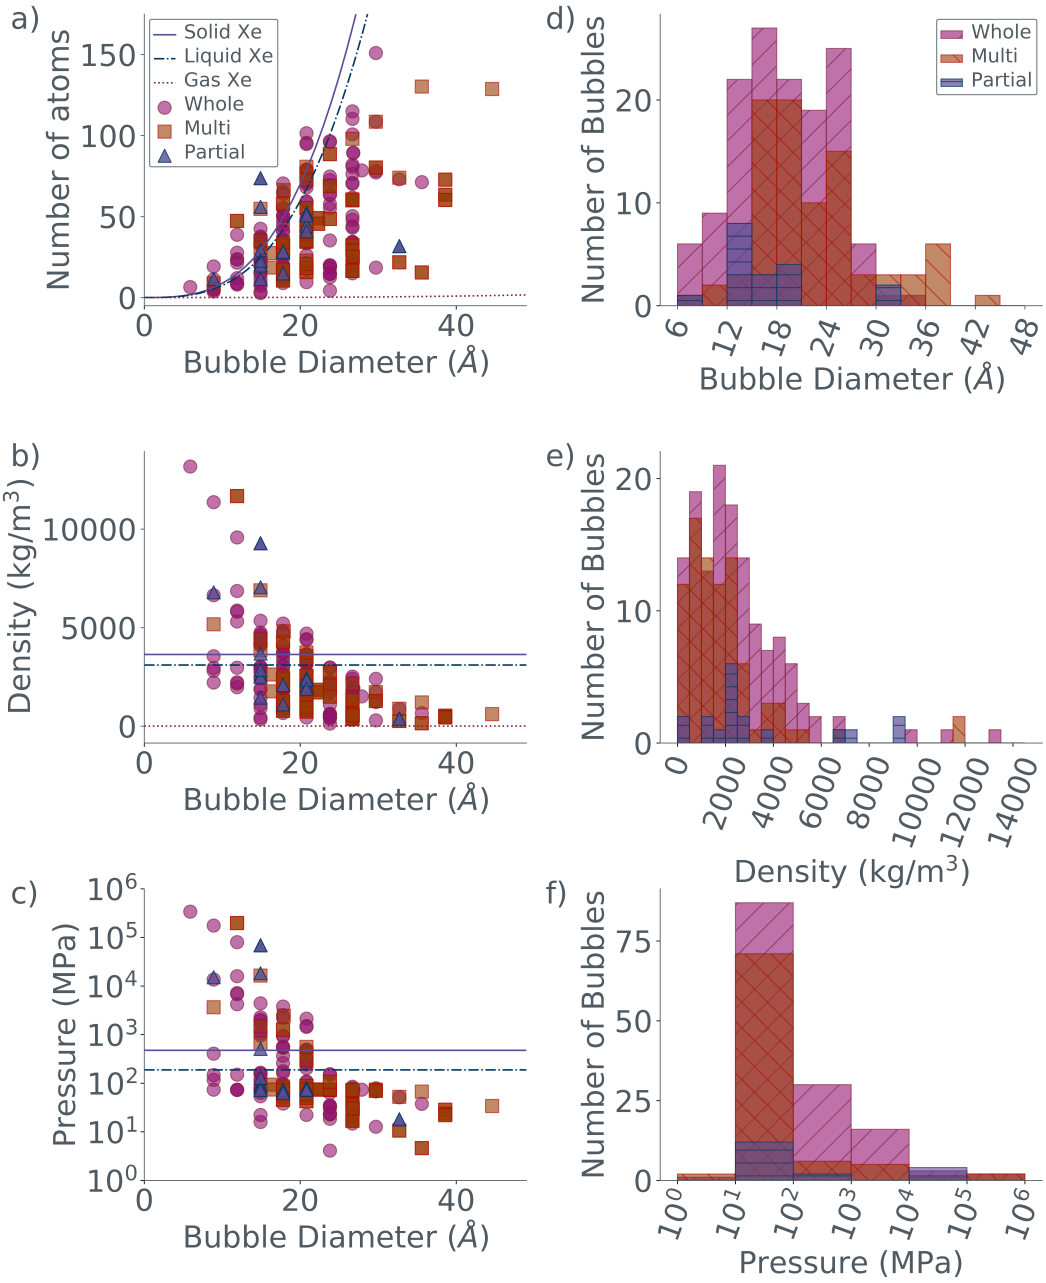

Figure 4.17: Analysis of the xenon bubbles in the xenon-implanted Zircaloy-4 sample, showing the whole (orange circles), partial (pink squares) and partial (purple triangles) bubbles, as well as the values for solid, liquid and gaseous xenon: a) number of xenon atoms in bubbles versus their diameter; b) density of bubbles versus their diameter; c) pressures of bubbles versus their diameter; d) histogram of bubble diameters; e) histogram of bubble densities; f) histogram of bubble pressures.

		Diameter (Å)	Density (kg/m^3)	Pressure (MPa)
	Mean	20.1	2493	4931
Whole (139)	Minimum	5.9	131	4.11
	Maximum	35.6	13178	339544
	Mean	18.3	3346	9609
Partial (18)	Minimum	8.9	378	18
	Maximum	32.7	9283	668986
	Mean	23.6	1931	4916
Multiple (87)	Minimum	8.9	143	4.62
	Maximum	44.6	11675	197507
All (244)	Mean	21.2	2355	5271

Table 4.4: Statistical breakdown of the xenon bubble properties, showing the mean, maximum and minimum values for the three bubble types (whole, partial and multiple). The number in brackets the number of each type of bubble. The final line is the mean values for all of the bubbles, irrespective of type.

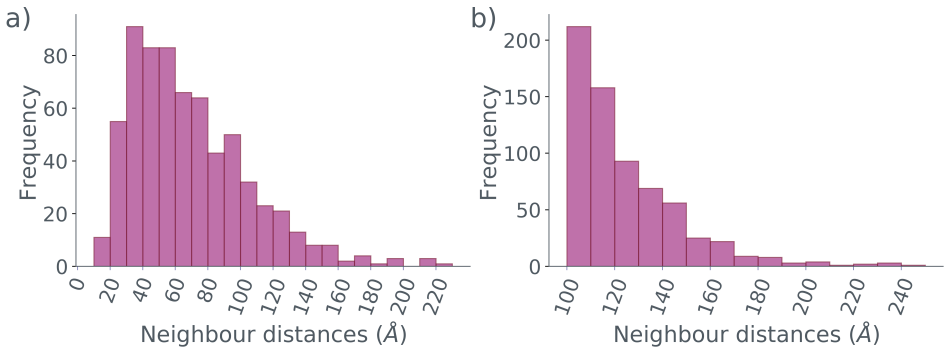

Figure 4.18: Distributions of the neighbour distances for xenon bubbles. Neighbouring bubbles are defined as any bubbles between which a straight line can be drawn without passing through another bubble: a) distances only in the $x - y$ plane, these distances were only taken in 2D as it was not possible to image the sample in 3D; b) an adjusted distribution assuming all bubbles also had a separation of 100 Å in the z plane.

Due to the large variation in number of atoms in any given bubble size, it is hard to determine the exact nature of the relationship between the two (linear, polynomial etc) but it is clear that the density drops as the diameter increases. The bubbles had a mean density of 2355 kg/m^3. It seems the bubbles are mostly in the density/pressure range corresponding to the gas phase, with some of the smaller ones falling into densities above that of solid

xenon. Implanted groups of xenon forming solid regions at room temperature have been previously observed [169].

On average, the bubbles had a mean pressure of 5.27 GPa. This average seems to have been dragged up by the partial bubbles, as the whole and multiple bubbles had a mean of 4.9 GPa. The bubble separation was calculated using the same method as the previous samples and has the same limitations. The sample thickness, based on t/λ calculations ranges from 250 Å to 400 Å, so an additional separation of 100 nm was chosen again. The mean separation of the original 2D distances was 69.1 Å.

4.4.2 Krypton in Zircaloy-4

4.4.2.1 Sample Preparation and Collection Conditions

The krypton in Zircalloy-4 sample was made in the same way as the xenon-implanted sample except in the case of the ion beam liftout procedure; for the krypton samples it was done using Ga ions in a FEI Nova Nanolab 200 instrument.

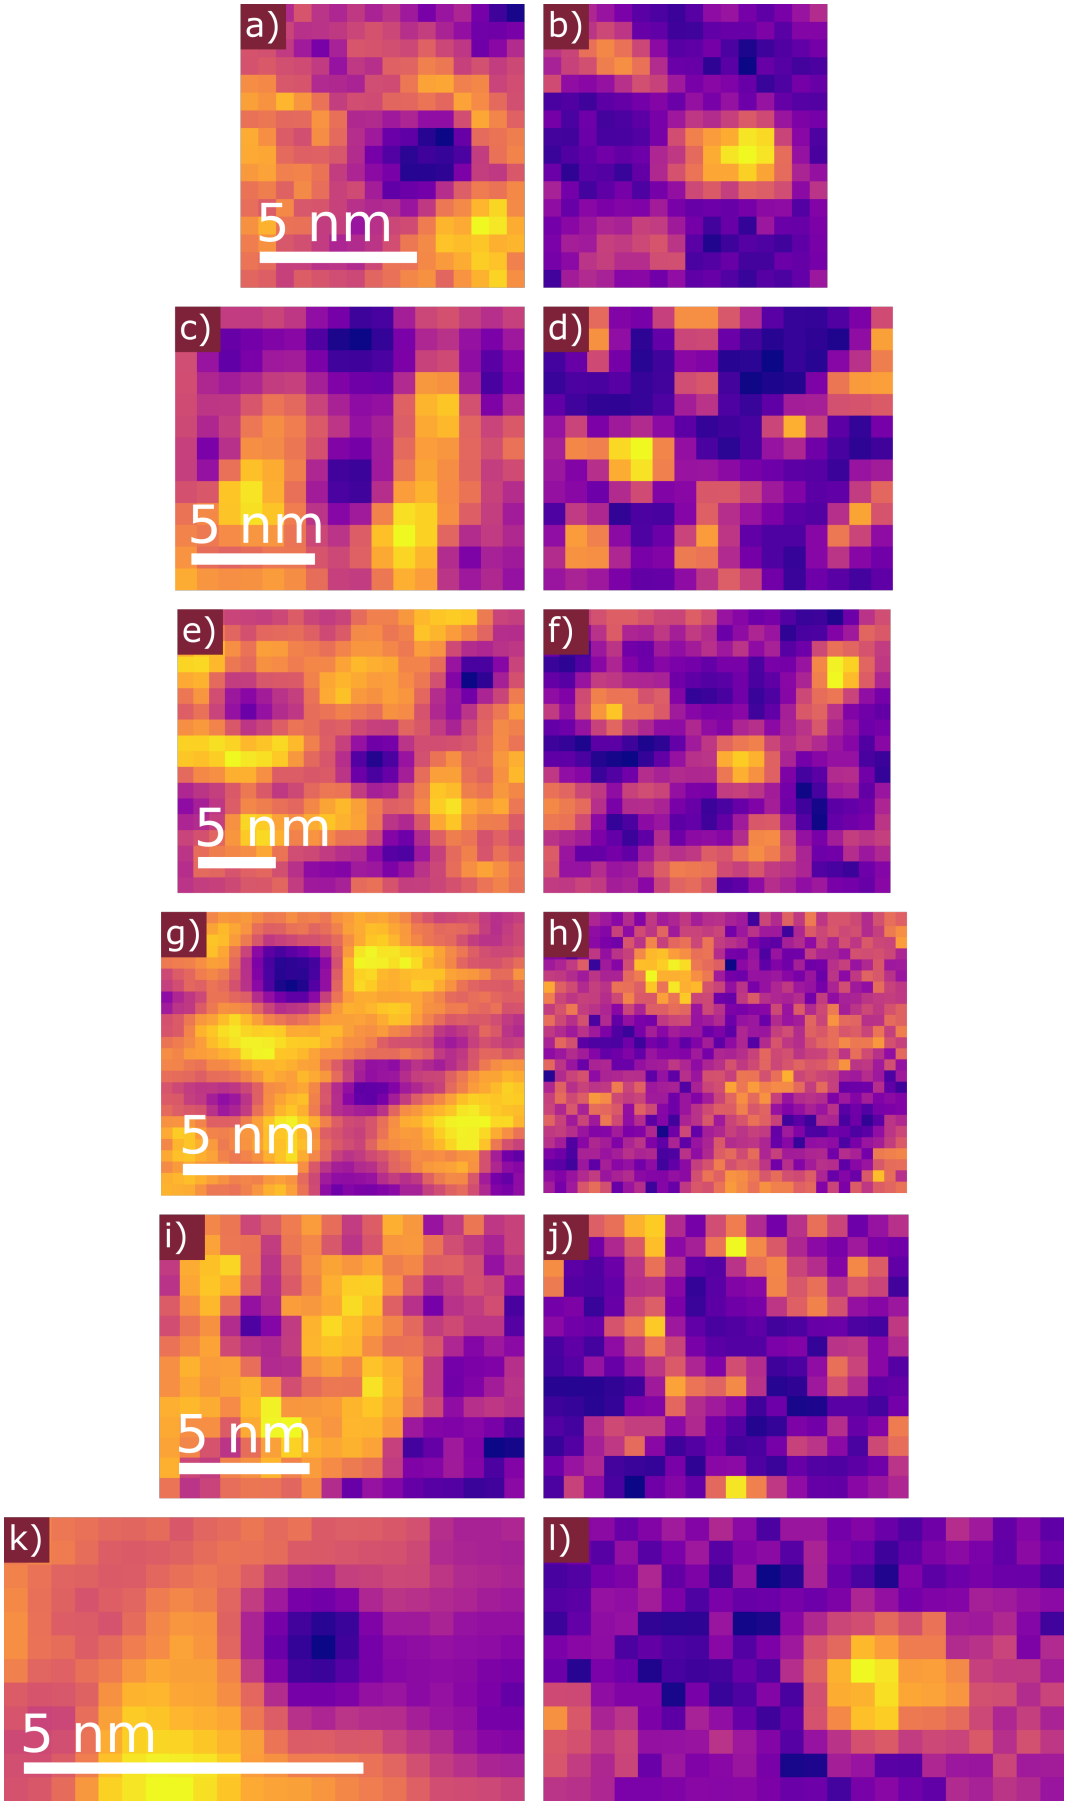

Figure 4.19: HAADF images (left hand side) and their corresponding krypton fit coefficient maps (right hand side) for the krypton bubbles in Zircaloy-4.

The pixel size, and high- and low-loss times varied between the datasets. The MLLS fits were performed between 1550-1900 eV to account for the krypton $L_{3,2}$ edges.

4.4.2.2 Results

The krypton regions are clearly visible in the HAADF images of fig. 4.19 as less dense regions, and analysis yields strong EELS krypton edges in the SI spectra. Despite this, analysis on the krypton samples was much less successful. There were fewer krypton bubbles when compared with the xenon.

The mean diameter of the bubbles was 37.0 Å, below the mean diameter for the whole bubbles, which was 39.3 Å. The krypton bubbles presented far higher densities than any of the previous samples.

The bubble separation was calculated using the same method as the previous samples and has the same limitations. The sample thickness, based on t/λ calculations ranges from 350-500 nm but an additional separation of 100 nm was chosen again. The mean separation of the original 2D distances was 50.7 Å.

Calculation and analysis of the densities and pressures of the bubbles were not possible, for reasons that will be discussed in section 4.5.2.

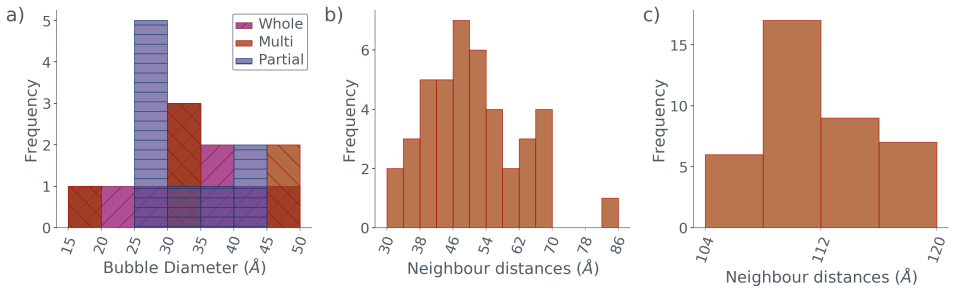

Figure 4.20: Properties of the krypton bubbles: a) histogram of bubble diameters; b) separations of the bubbles only in the $x-y$ plane, these distances were only taken in 2D as it was not possible to image the sample in 3D; b) an adjusted distribution assuming all bubbles also had a separation of 100 Å in the z plane.

4.5 Discussion

The discussion section will first discuss the tantala and the Zircaloy-4 films separately before comparing them and then wrapping up with a final discussion. Before that, however, there are some general notes on the analysis.

The bubbles were assumed to be spherical. This is the standard shape throughout the literature and is supported by their shape in the HAADF images. Additionally it matches the initial modelling by Prato [132]; while this work assumed the bubbles were empty voids, the best match to the experimentally measured optical properties came from a simulation that assumed the bubbles were spherical. Their volumes were calculated using the equation for the volume of a sphere.

The pixel size is the largest source of uncertainty for this data. The EELS quantification is very accurate and its associated uncertainty is on the order of single atoms. This is because these values are dependent on the intensity of the ZLP and the relevant edges, both of which contain a huge number of counts; therefore, their associated random error is very low. Any errors introduced in this step will be systematic, such as miscalculating the semi-empirical standard. The diameter uncertainty, however, may account for much of the spread of densities and pressures at the same diameter, and vice versa. This uncertainty is not included in the graphs as it makes them too visually cluttered but it is important to consider. To give a numerical example of the effect the pixel size has on the calculations; increasing the diameter of a bubble of 15 Å by 3 Å would almost halve its pressure, whereas decreasing it by the same amount would double it. For a bubble of diameter 36 Å, applying the same calculations would decrease the pressure by a fifth or increase it by a third, respectively.

4.5.1 Gravitational Wave Coatings

When calculating the pressures inside of the bubbles, an equation of state is required. Inititally the ideal gas law was considered, but this was quickly discarded as its assumptions do not hold in this case; at such small volumes and high densities the argon atoms cannot be assumed to be non-interacting point particles. Other equations of state (EoS) were investigated, including the Wohl model, which was discarded as it is not fit for high densities. Also investigated were the Beattie-Bridgeman model [170] and the Redlich-Kwong EoS [171], which superceded the former. The Redlich-Kwong EoS is essentially an empirical adjustment to the van der Waals EoS, and many later adjustments are based on it. It gives a more realistic description of the pressures of gases above their critical temperature. Despite the Redlich-Kwong's EoS improved accuracy the van der Waals was chosen to calculate the pressures inside the argon bubbles. This decision was made for two reasons; the first is that tracking down the empirical constants for argon was proving less than simple and the second is that the pressure calculations were only ever intended to be first approximations. Accurately calculating the pressures in such extreme conditions requires molecular dynamics simulations, which are outwith the scope of this thesis.

The van der Waals EoS

$$(P + \frac{an^2}{V^2})(V - nb) = nRT, \tag{4.1}$$

where P is pressure, V is volume of gas, n is the number of moles, T is temperature, R is the gas constant and a and b are experimentally derived constants unique to each material (1.345 L^2 bar mol $^{-2}$ and 0.03219 L mol $^{-1}$, respectively, for argon [172]).

The van der Waals EoS did fail for some of the smallest and most dense bubbles – giving either negative or unphysically high pressures. The van der Waals equation is known to break down in extreme cases and if the smallest gas bubbles are high-density fluids, then a gas equation of state would not be an appropriate way to describe them. Recent work by Jelea [173] has shown that for xenon and krypton bubbles the van der Waals EoS tended to overestimate the bubbles' interior pressures. As the van der Waals EoS is for gases, it is naturally inappropriate for all the bubbles more dense than liquid argon. It is also possible that the more dense bubbles were actually two bubbles overlapping in the beam path in the z-plane – essentially falsely doubling the number of atoms within the bubbles. As the bubbles do not affect the thickness map of the samples and the images taken are two-dimensional there is no way to ascertain for certain whether the densest bubbles are two bubbles that are completely overlapped.

The titania-doped sample was the first to definitively prove that argon implants in this and other similar materials and forms bubbles on annealing. The question is then; why do the bubbles occur? The presence of argon atoms themselves can easily be explained – they are trapped during sample deposition and are therefore dispersed throughout the sample, which is hundreds of nanometers thick for both single- and multi-layered coating samples. In materials with noble gas bubbles formed by ion implantation the gas atoms displace the surrounding atoms, causing highly strained positions in the lattice. This strain is then reduced by the aggregation of the noble gas atoms [160, 158, 159]. In the coatings sample however, although the argon atom does not bond with the surrounding atoms it is assumed that it does disrupt the bonds of the surrounding Ta_2O_5 glass. This means the glass is still in a high energy state that could be reduced by the formation of noble gas clusters but this will

only occur when there is enough thermal energy for the argon atoms to diffuse randomly. This means that bubble size will vary with annealing temperature.

The total argon content of the sample is not trivial to calculate. A similar sample investigated using RBS contained 3% argon by atomic ratio [134]. Comparing the number of argon atoms in the bubbles in our sample, with the total number of atoms in the matrix (done using number densities of atoms and the MFP of Ta_2O_5 and Ti_2 [165, 100]), gives a value of <1% argon. It is likely that the rest of the argon has not coalesced into the bubbles and is still within the matrix as individual interstitial atoms and sub-nanometre clusters. This is supported by the bubble maps in fig. 4.6, where there are still argon counts in pixels outwith the bubbles. It is possible that these argon atoms would precipitate into bubbles upon longer or hotter annealing. The argon that was close to the surface may have also escaped during the annealing or sample preparation. It was not possible to calculate the bubble volume fraction as there was no argon-free Ta_2O_5 film with which to compare it. Preliminary calculations suggest, however, that it may be around 4-5%, on account of the relatively low density of the bubbles, when compared with the Ta_2O_5. This calculations were done using the dimensions of one of the images and the thickness calculated using t/λ calculations. From this the volume contained within the image can be calculated, and as the bubble volumes are known it is possible to calculate what percentage of the volume is bubble.

The zirconia-doped tantala bubbles presented different difficulties. As previously discussed, a large portion of the argon in this sample had not coalesced into bubbles and was instead dispersed throughout the sample. This made calculating the bubble properties particularly challenging. The edges of the bubbles were not clearly delineated; the smaller bubbles were not always distinguishable from the background; and the background artificially increased

the atom count, thereby inflating the densities and pressures of the bubbles. Subtracting the background was done in very broad strokes – summing the argon count in a 'typical' area of the argon coefficient map. While this seemed to work well for the majority of the bubbles it is not a perfect solution and gave some of the bubbles a negative atom count. In one of the datasets two backgrounds were taken as there were clearly two distinct regions when it came to argon presence. Perhaps a better solution would be summing an area around the bubble and using that as a localised background for each individual bubble. One way to do this may be by doubling the bubble radius and using the gas atom intensity of the 'outer ring' as an approximation of the local background. This could also cause problems however, as many of the bubbles are close enough that this 'background bubble' would overlap with neighbouring bubbles.

The reason for the very different argon presence in the zirconia-doped sample than in the titania-doped sample is unknown. It is unlikely that the discrepancy is due the zirconium. Although zirconium is larger than titanium, it should behave the same chemically and not present any additionally challenges to the argon atoms difusing and coalescing into bubbles. The annealing conditions for this sample are unknown but, compared with the three titania-doped samples, it seems probable that it is at least partially due to insufficent annealing. There may also be more argon in this sample as it was made with a different instrument, or that less was lost during sample preparation.

There is little to say about the oxygen sample other than as a cautionary tale – it is not even possible to tell if bubbles are present in this sample as it was so damaged by the beam. It was initially hoped this would be a good control with which to compare the other samples but the data was too poor in quality. Spatially resolving these bubbles would have presented a similar challenge to the argon zirconia-doped sample, due to the presence of oxygen

in the matrix. If these samples had not damaged so easily, however, the first step would have been quantifying the matrix, and then seeing if there was more oxygen in this sample when compared with the argon sample. Sadly due to the quality of the data this was not possible.

The titania-doped tantalum pentoxide matrix was not quantified as this work had been done previously by Isa [165], using similar elemental quantification methods as described in this chapter. The zirconia-doped tantalum pentoxide sample was not quantified as it was deemed unnecessary; the focus was on developing the bubble analysis technique and applying it to more samples. Additionally, the coating was not chosen for aLIGO so there is no wider benefit to the GW community by characterising it.

The bubbles in the zirconia-doped sample are, on average, 50% smaller than their counterparts in the titania-doped sample. When considering the large amounts of argon not in the bubbles, it does seem likely that the bubble diameters were being overestimated due to the large argon background values. This also probably there are more bubbles from the zirconia-doped sample at higher densities with larger diameters. Creating a better background subtraction, or annealing these samples further, may lower the number of higher-density bubbles.

Despite the problems with the zirconia-doped sample, broadly speaking the results from the two samples are consistent. There is a large spread of number, density and pressure values at the same diameter but generally number increases with diameter, whereas density and pressure decrease with diameter. Most of the bubbles still exhibit densities consistent with a gas, although some may be liquids, solids or high-density liquids. The average densities were similar, albeit slightly more dense in the zirconia-doped sample. The pressures were the same order of magnitude, although about double for the zirconia-doped sample.

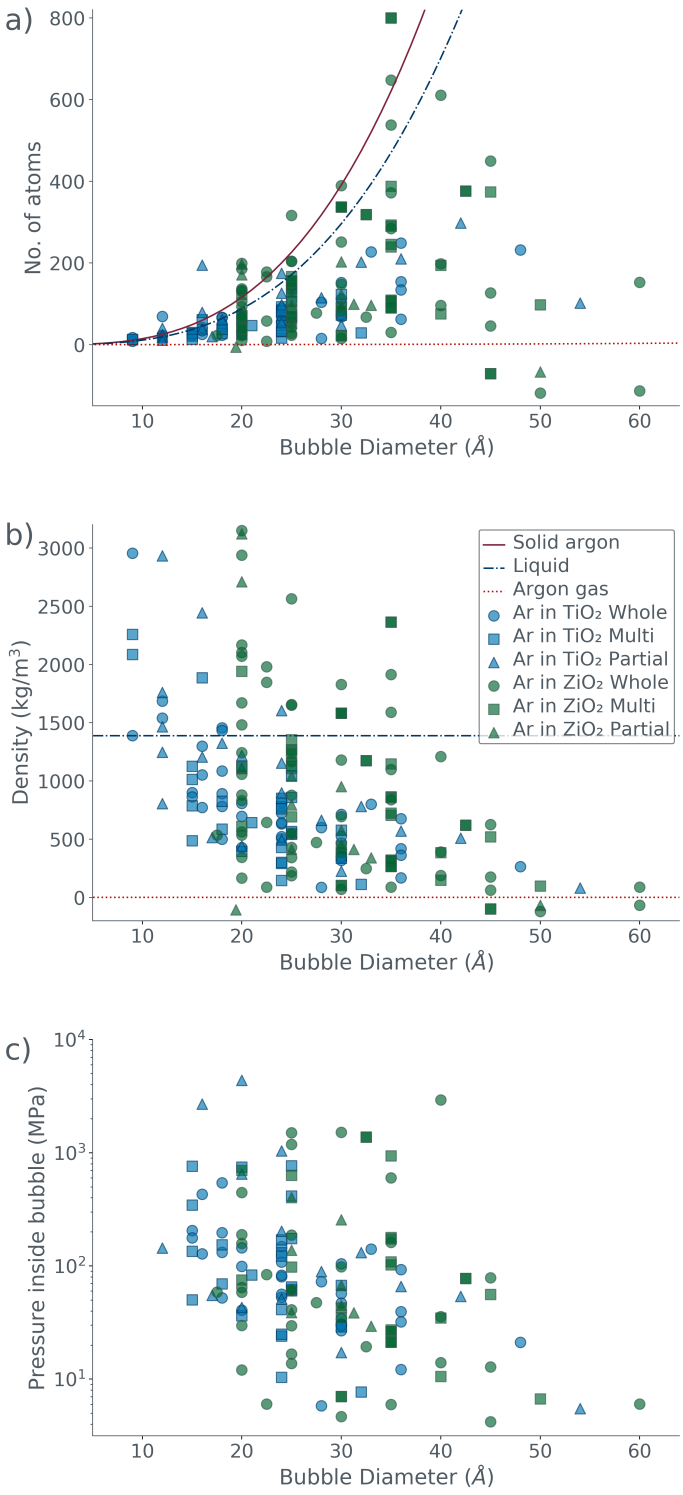

Figure 4.21: Comparison of the argon bubbles from both tantala samples. The bubbles from the titania-doped sample are blue (the whole bubbles are circles, multiple bubbles squares, and partial bubbles triangles) and the background-subtracted bubbles from the zirconia-doped sample are green: a) number of argon atoms in bubbles versus their diameter; b) density of the bubbles versus their diameter and; c) pressures of the bubbles versus their diameter

4.5.2 Nuclear Reaction Cladding

As established, the van der Waals equation works less well for xenon and krypton bubbles [153]. As xenon and krypton are fission products they have been much more extensively studied than argon. Consequently, it was possible to find an equation tailored to xenon and krypton bubbles. Jelea put forth an EoS for xenon/krypton mixtures confined in nuclear fuels

$$P_b = \rho_b kT + \sum_{i=1}^{3} F_i(x,T)\rho_b^{i+1}, \tag{4.2}$$

where P_b is the pressure of the gas within the bubble and ρ_b is its density, k is Boltzmann's constant, T is temperature and $F_i(x,T)$ is a function of the molar fraction of krypton (x) of the gas and temperature (T). The equation was tested at the extremes of 100% krypton and 100% xenon. While the equation is designed for bubbles in uranium dioxide it was felt this equation was a close enough match to provide a similar insight to the pressures as the van der Waals equation could provide for the argon bubbles.

The krypton bubbles are much more clearly visible than the xenon bubbles in the HAADF images, despite the difficulty in analysing the krypton bubbles. This is somewhat surprising as xenon has a larger atomic number than krypton, which is what gives the HAADF images contrast. However, the krypton bubbles are both denser and bigger than the xenon bubbles. The latter is unsurprising, as a similar effect has been observed with other pairs of elements in the same group of the periodic table. If the same high manganese steel is made with either vanadium and niobium, there will be precipitates of the carbides VC or NbC present. The NbC precipitates will be smaller because Nb is larger and therefore the diffusion coefficient will be smaller. This means it will be more difficult for the atoms to move around and form bubbles, meaning there will be

more but smaller precipitates [104]. It is possible the same is happening here, as this is a phenomena based on the respective sizes of the precipitate/bubble atoms and the matrix; larger atoms have smaller diffusion coefficients. As krypton is smaller than xenon it can move about more easily and therefore form fewer but larger bubbles. It would be expected however, that the krypton bubbles' greater diameters would, on average, make them less dense not more, as Ostwald ripening is expected [174, 175, 185].

It is not possible to compare the two Zircalloy-4 samples, due to the problems with the krypton datasets. In many cases, the krypton bubbles gave unphysical densities and the average density of the krypton bubbles was above the density of the surrounding Zircaloy-4 (6550 kg/m^3 [176]) by several orders of magnitude. Initially, it was thought that some of the extremely dense bubbles were more than one stacked on top of each other, as the krypton bubbles' larger average diameters would make this more likely. This could not, however, explain away the densest of the bubbles. Additionally, the bubbles are visible in the HAADF images (fig. 4.19). This suggests that they are not denser than the surrounding Zircaloy-4; although HAADF imaging is primarily sensitive to Z-contrast, it is also sensitive to density. Further investigations revealed that the Zircaloy-4 appeared denser than it is. This suggests that there is a problem with the data taken from the krypton sample. The exact problem is unknown but there are a few possible causes.

The first is that the Hartree-Slater cross-section may be underestimated at these energies. Hartree-Slater cross-sections work well between 300 eV and 800 eV [101], and, even if they do fail at higher energies, it is unlikely the error would be so large [107]. This means that this alone could not account for the discrepancies of the krypton data. The second reason is that optical effects may cause a different collection angle than the standard calibration [1,

177]. The collection angle is a function of energy loss and its correlation is stronger at higher energies. By optimising the microscope setup this correlation can be weakened but the krypton data was collected before the optimisation process was developed and so it is possible this is affecting the data. As with the standards however, these optical effects also could not account alone for the krypton values. The final possible cause is the most likely and it relates to the time ratio of the high- and low-loss spectra. The krypton data is the only data that did not need to be spliced together, as the high- and low-loss did not overlap. This means the high-loss data had to be manually adjusted for the time ratio of it and the low-loss data. The true time ratio in the spectrometer may deviate from the stated ratio, resulting in a systematic effect of the high-loss being more or less intense than expected. It is not simple to check this at high-losses, where splicing is not possible. It seems likely that this is a significant effect. It would require a systematic investigation to ensure that absolute quantification from cross sections works for edges above 800 eV and to estimate the likely scale of this problem. It is not possible with only the one sample to make such an estimation now, especially as the krypton values may be affected by more than one issue.

Unfortunately, the krypton sample is not longer at the University of Glasgow, so it was not possible to collect more data from the sample.

4.5.3 Comparing the xenon and argon bubbles

Overall the argon and xenon datasets seem to be in good agreement – despite using different pressure calculations for argon and xenon the final pressures are in the same region. The biggest outliers come from the zirconia-doped sample's argon bubbles, whose problems have already been discussed. While there are smaller xenon bubbles this comes from limitations of the technique – the pixel size was smaller on average for the xenon bubbles (3 Å versus between 3 Å and 6Å for the argon bubbles) and failures of the van der Waals equation. While there are smaller argon bubbles initially they are lost in the pressure graph as they give either negative or unphysically high pressures. The xenon bubbles are more dense on average than the argon bubbles, which is not surprising as xenon is heavier than argon but it is also because of the smaller xenon bubbles. Smaller bubbles are more dense, so the greater number of smaller xenon bubbles will drive up the average density. The higher average density of the xenon bubbles, and the lower density of the Zircaloy-4 than the tantala may be why the xenon bubbles are less visible in their HAADF images (fig. 4.16) than the argon bubbles in theirs (fig. 4.6 and **??**).

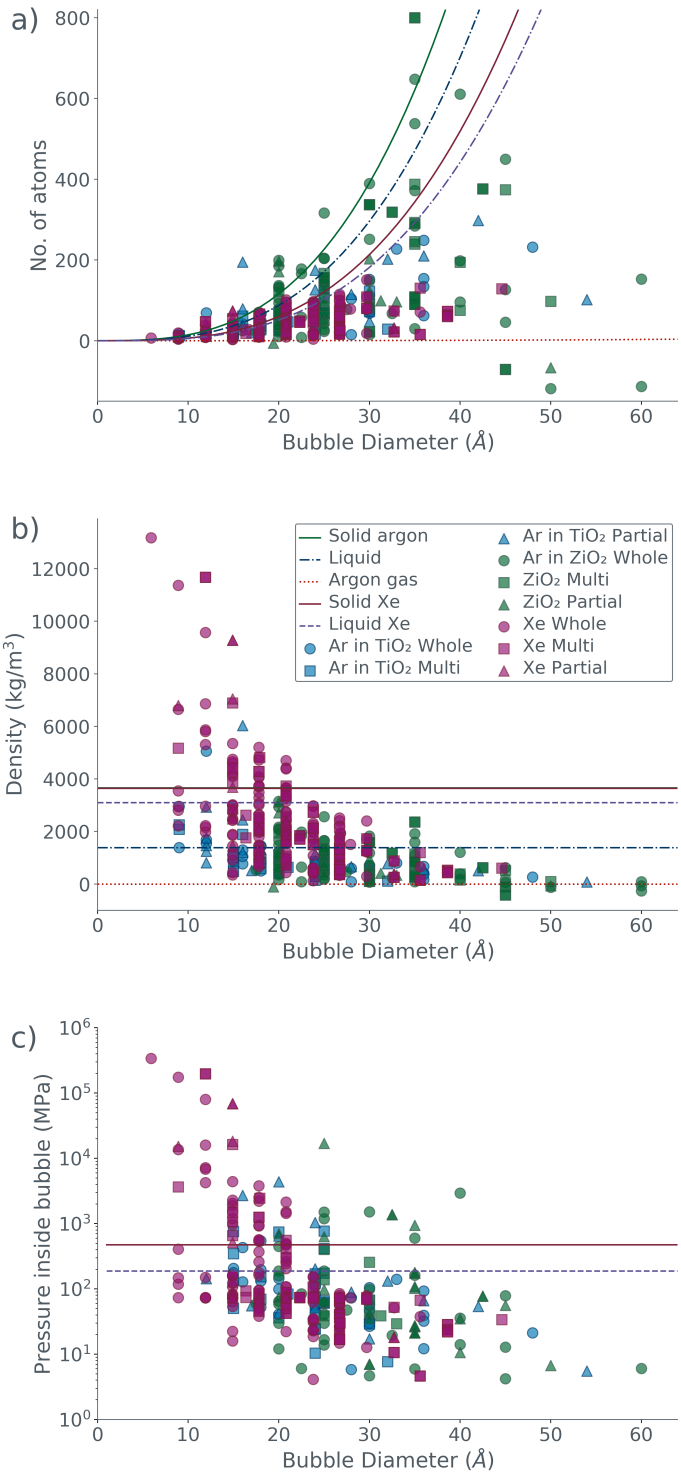

Figure 4.22: Comparison of the argon and xenon bubbles. The xenon bubbles are pink (whole = circles, multiple = squares, partial = triangles) and the argon bubbles from the titania-doped sample are blue and the argon bubbles from the zirconia-doped sample are green: a) number of atoms in bubbles versus their diameter; b) density of the bubbles versus their diameter and; c) pressures of the bubbles versus their diameter.

	Diameter (Å)	Density (kg/m^3)	Pressure (MPa)
Argon (TiO_2-Ta_2O_5)	22.7	959	241
Argon (ZiO_2-Ta_2O_5)	29.3	960	486
Xenon	21.2	2355	5271

Table 4.5: Mean values for all of the bubbles, taking all bubble types into account.

As there was only one krypton sample and one xenon sample it is not possible to compare how annealing affected them in comparison to how it affected the argon in the titania-doped tantala sample.

It was considered whether the Jelea equation 4.2 could be adjusted to account for argon but as it relies on several constants derived through molecular dynamics equations this was not possible. Despite this, the pressures being in similar ranges despite the different equations suggest that both are working, at least as a first approximation.

It is not possible to tell from the EELS spectra whether the bubbles are solid or still fluid. The signal to background ratio for the relevant gas-edge is so small for any individual bubble that, when compared with the overall EELS spectrum, any fine details of the ELNES are not distinguishable, which is where such information would be found.

4.5.4 Final Discussion

While the bubble properties have been discussed in great detail; the importance of characterising the bubbles will now be put into context. Potential further work will also be discussed. First however, some context will be provided for the high densities of some of the bubbles.

While densities suggesting that solid argon and xenon bubbles are present in the films may initially seem surprising, solid noble gas bubbles are a common phenomenon [178, 137, 142, 141, 179, 140, 142, 145, 180, 181, 100, 169]. Although 'bubbles' is commonly used, the term often refers to solid precipitates or particles. Some studies have determined the bubbles have a crystalline structure [182, 183], with lattice images obtained [184]. It has been observed that bubbles coarsen with annealing [152, 185], just as was observed with the argon bubbles in the titania-doped tantala sample. As discussed at the beginning of this chapter, bubbles of noble gases are of particular concern to the nuclear industry, where fission products implant in reactor components, potentially causing swelling and cracking [161, 162]. The mechanics of how the bubbles form and affect the surrounding matrix is outwith the scope of this thesis but both are active areas of research [135, 186, 187, 188, 189].

The final goal of this analysis is reducing the noise in LIGO's detectors. Creating a technique to analyse the samples is simply the first step and analysing the current mirror coatings specifically is the second. After this; having a better model for the bubble pressures, understanding how they chemically and physically affect the surrounding material, as well as loss measurements in tandem with well described bubbles will be required. Some of this work has already begun [190, 191].

The optical properties of the films will almost certainly be affected by the bubbles, most notably the mechanical loss. It has been found that mechanical loss initially decreases with increased annealing temperature [192], however at annealing higher temperatures the benefits begin to plateau [134]. The initial decrease of mechanical loss is probably due to structural relaxation from an initial 'frozen' state that is far from equilibrium. Above a certain annealing temperature, however, it may be that the bubbles become a critical size and where previously they were small enough to have little or no effect on the loss, they might now be large enough to appreciably affect it [135]. Alternatively, it may be that their formation is positively affecting the mechanical loss, but beyond a certain size their contribution decreases. They may also be affecting the scattering loss [193], as the larger the bubbles, the more likely the incoming photons are to be scattered or absorbed by them rather than the tantala glass. This also cause laser damage to the coatings, as the argon bubbles do not have the same desirable low-absorption properties that the glass does. It has long been suspected that the presence of the bubbles (or what we have now experimentally confirmed to be bubbles) may be contributing to laser damage in the coatings.

While the xenon bubbles in Zircaloy-4 present less immediately obvious benefit to reducing mirror coating noise it does not mean there is none. Recent work done on dielectric multilayer films for ultra-high power and energy laser applications investigated using xenon rather than argon during the sputtering process [194]. There, it was found that using xenon reduced laser damage for UV-lasers, as xenon bubbles were less likely to form and when they did, were closer to the density of the hafnia films. While their sample preparation meant that the xenon behaves differently in their work than here, it shows that understanding xenon bubble properties is still beneficial for GW coatings.

It also shows that bubble formation happens in many different thin films, so figuring out either how to eliminate the bubbles or understanding their noise contribution is crucial to improving LIGO's mirror coatings.

The next steps for this project come in three parts. The first is having more samples[1], the second is molecular dynamics work and the third is an investigation into cross-sections for losses above 800 eV. Being able to collect statistically significant measures of mean bubble diameter after annealing at multiple temperatures would make it possble to estimate the average energy barrier to diffusion for an argon atom in the Ta_2O_5 glass. From this the average diffusion distance for a given annealing time and temperature could be estimated, as well as what proportion of the total argon will form visible nanobubbles. It would also be beneficial to compare samples made in the same argon environment with the same annealing conditions but of different composition to see how that affects bubble formation. These experiments would likely be done in tandem with X-ray scattering [134] and ellipsometry [133]. The use of a direct detection detector could improve this technique by lowering the noise. Atomic modelling could potentially identify the modes by which the bubbles affect the loss. Molecular dynamics simulations could also shed light on the bubble formation if a suitably large cell could be constructed, with a reasonable number of argon atoms interspersed with the tantalum, oxygen and titanium atoms. Investigating the cross-sections comes in two parts. The first is establishing, experimentally, if Hartree-Slater cross-sections are underestimated at high losses. This would require collaboration with a group that has a gas cell for TEM and STEM/EELS. The second part is a systematic study of the actual time ratios of high- and low-loss spectra from high energy edges.

[1] And a reliably functioning microscope...

When considering applications outside of GW detection; this method for studying small noble gas bubbles is directly applicable to other noble gases in a range of materials, not just coatings. Additionally, most other studies on the properties of gas bubbles in similar materials are based on indirect measurement methods and modelling [173, 195, 182, 196]. Now that a direct measurement method has been proposed, work can be done on comparing the results. More generally still, this method can also be applied to any fluid bubbles or even solid precipitates in materials where there are no good standards, with the caveats of a suitable EELS edge being present and the material not being susceptible to damage by the beam.

4.6 Conclusion

The history and importance of bubble characterisation was discussed. The method for creating a semi-empirical standard and using it perform spatially-resolved absolute-quantification was then outlined. This technique was then used to analyse bubbles in four different materials materials using STEM; argon in titania-doped tantala annealed at a range of temperatures; argon and oxygen in zirconia-doped tantala; and xenon and krypton in Zircaloy-4. The bubble properties, namely diameter, density and pressure, were determined. The results were then discussed and the samples were compared. The sample deposited in oxygen could not be examined as it suffered severe beam damage, and the titania-doped tantala samples annealed at the lower temperatures could not be quantified due to the low SNR. It was not possible to fully analyse the krypton bubbles but possible reasons for this were discussed. The argon bubbles in the highest temperature annealed sample, the argon bubbles in the zirconia-doped tantala sample, and the xenon bubbles did have consistent

results. Despite using different EoS to calculate the pressures of the argon and xenon bubbles they gave similar pressure ranges, as would be expected. The xenon bubbles had a higher average density and pressure, which is consistent with other research [194].

How the bubbles may be affecting the coatings was discussed. To summarise: while it is currently unknown exactly what effect the bubbles have on the coatings, it is believed they affect the mechanical loss (although in what way is unclear) of the coatings, as well as increasing laser damage due to photon scattering and absorption by the bubbles rather than the tantala glass. There was also a discussion of potential further work in this project, in particular how to better determine what effect the bubbles are having on the surrounding materials and on the optical properties of the coatings. This would likely require simulation and modelling work, in combination with further experimental work using both microscopy and other techniques.

5

What We Can Learn from EXELFS

This chapter investigates whether recent developments in high energy-loss EELS mean it would be possible to use EXELFS to investigate the short range order (SRO) of GW mirror coatings. The goal is twofold: firstly, to advance EXELFS as a technique, as with quantification in Chapter 4; and secondly, to provide more insight into the structure of the GW mirror coatings. The importance of understanding the structure of the coatings has already been discussed at length but there are some nuances to understanding the SRO specifically. Heat treating of the tantala layer may lead to some local ordering, which may influence the loss. For the silica layer, the intention was to investigate whether the fabrication process had any effect on the SRO.

We start with a brief recap of extended energy-loss fine structure (EXELFS): following the edges in an EELS spectrum, there are oscillations caused by the interference of the outgoing electron wave with the backscattered waves from

the neighbouring atoms. These oscillations are known as EXELFS and give information on the atomic environment in which the target atom sits. EXELFS was first proposed by Leapman and Cosslet in 1976 [122]. More information can be found in section 3.5.

EXELFS is analagous to extended X-ray absorption fine structure (EXAFS) but is less commonly used, despite the advantages it offers. This is because attaining high-quality EXELFS data is much more challenging: EXELFS samples are typically more damaged by the beam as they require a higher beam dose per unit volume; the lower energy edges that can be studied with EXELFS are closer together, meaning the oscillations and background fitting overlap; the background itself can be difficult to calculate for EELS data; and finally, SNR is generally lower, in part due to the comparitively high readout noise of the spectrometer (see section 3.2).

Much of the early work on EXELFS was primarily to establish it as an equivalent technique to EXAFS. It was found that it could compete, but only below 1.5 keV. Martin and Mansot used EXELFS to examine the differences between amorphous and crystalline silicon carbide at K and L edges [197]. Serin et al. [198] looked at EXELFS modulations above the K edge in low Z elements. They reaffirmed the technique's usefulness in investigating inhomogeneous materials due to its high spatial resolution but also discussed the limitations, which included the fact that overlapping edges can cause a loss in resolution and that the effects of multiple scattering reduce the useable thickness of the sample. They did note, however, that recent improvements in deconvolution techniques may reduce the effects of this problem. In 1996, Haskel et al. also demonstrated some of the possibilities of the technique, introducing a variety of ways to improve the data acquisition such as realtime alignment and accumulation of a large number of spectra all while monitoring sample drift, radiation damage and

changes in energy resolution; therefore improving the EELS data quality. They then used EXELFS to analyse the K, L_1, and L_2 edges in a range of elements [199]. The same group later investigated the effects of radiation damage on EXELFS, finding that in pure metallic materials the effects were minimal but became more serious in ionic solids [200]. In more recent years, EXELFS has been used to analyse the short-range order of multi-component systems [201]; investigate magnesium vandate catalysts [202]; study spatial resolution and crystal anisotropy [203]; probe the structure of iron nanoparticles in carbon nanotubes [204]; and quantify volume expansion at grain boundaries [205]. Improvements in aberration correction (section 2.2.1) in the past few decades have meant that EXELFS is not widely used to analyse crystalline materials because it is now possible to directly visualise atomic positions with relative ease. However, this is clearly not useful for studying amorphous materials like the GW coatings studied in this chapter. EXELFS, however, gives information on the SRO of even amorphous materials. The SRO of GW mirror coatings is of interest as it may shed light on why coatings of the same material but made using different processes can have different properties.

At the University of Glasgow, in 2017, it was established that high energy-loss EELS (in excess of 5 kV) could, theoretically, be used to study local atomic ordering with EXELFS [206]. It would, however, require sufficiently long data acquisition times and optimising the microscope and spectrometer optics to reach a suitable SNR. It was then found that optimising the camera length resulted in a monotonically decreasing background, which allows for good-quality background subtraction, allowing better extraction of the oscillations from the raw microscope data [1]. Previously, due to imperfectly designed EELS lens programs the background did not have a perfectly monotonical decay. This made extracting the signal from the data challenging. Coupled

with other improvements to the microscope design, this advancement is a step towards overcoming one of EXELFS main limitations.

Some of the work in references [1] and [2] has been included in the initial work portion of this chapter. The transition metal edges and silicon data were collected for [1], although no EXELFS processing on the transition metal edges was included in that work. The paper was written by Ian MacLaren, with the silicon analysis done by me. I also wrote a short literature review on EXELFS for the paper that has been expanded and rewritten into this introduction. The titanium data were collected by Jamie Hart at Drexel, and used in reference [2], a conference paper for Microscopy and Microanalysis 2019. That work was done by Jamie Hart, following discussions between Jamie Hart, Ian MacLaren and myself. It focusses more on comparison between traditional EELS done with a spectrometer and EELS done using direct detection. The titanium data in this thesis is the same raw data but was processed by me, with a view to improving EXELFS done with the current EELS set-up at the University of Glasgow.

The results will be presented first, before being discussed in detail in the second half of the chapter.

5.1 Initial Work

This initial work was not done on GW coatings. This was partly just to make use of available data and this early work focussed mostly on checking that EXAFS software and analysis techniques worked with EELS data. We start by investigating the L-edges of five transition metals, then moving on to a silicon sample before finally looking at titanium data collected by collaborators at Drexel University.

The data were all processed using the typical processing methods discussed in section 3.4.

5.1.1 L-Edges of Transition Metals

Four different transition metals were looked at; zirconium, molybdenum, ruthenium, rhodium and palladium. Performing EXELFS analysis on the $L_{2,3}$ edges was never possible – the separation of the L-edges of these samples were all below the 300eV threshold for extracting reliable oscillations. Between the L_3 and L_2 edges there is, at most, 150 eV; between the L_2 and L_1 edges there is a larger gap but still always less than 300 eV. An attempt was made to see if there were any meaningful oscillations after the L_1 edge. The edges and their k-space oscillations can be seen in fig. 5.1 and all five datasets have a k-weighting of one. Initially, there do seem to be some oscillations in the zirconium and molybdenum data, but the data are too noisy for the oscillations to be of use. At first glance the rhodium data (section 5.1.1.g&h) has smooth and clear oscillations but their magnitude is approximately a million times larger than the other datasets, which casts doubt on their veracity. There are some small peaks after the edge in the rhodium data, it seems likely that Athena is misinterpreting these peaks as oscillations. The palladium data is

cut too short after the edge to be of use. There does not seem to be any relation between increasing atomic number and the oscillations. There are some similarities between zirconium and molybdenum (fig. 5.1.b&d), despite the poor SNR in the ruthenium oscillations. Both are hexagonally close-packed, so the similarities make sense.

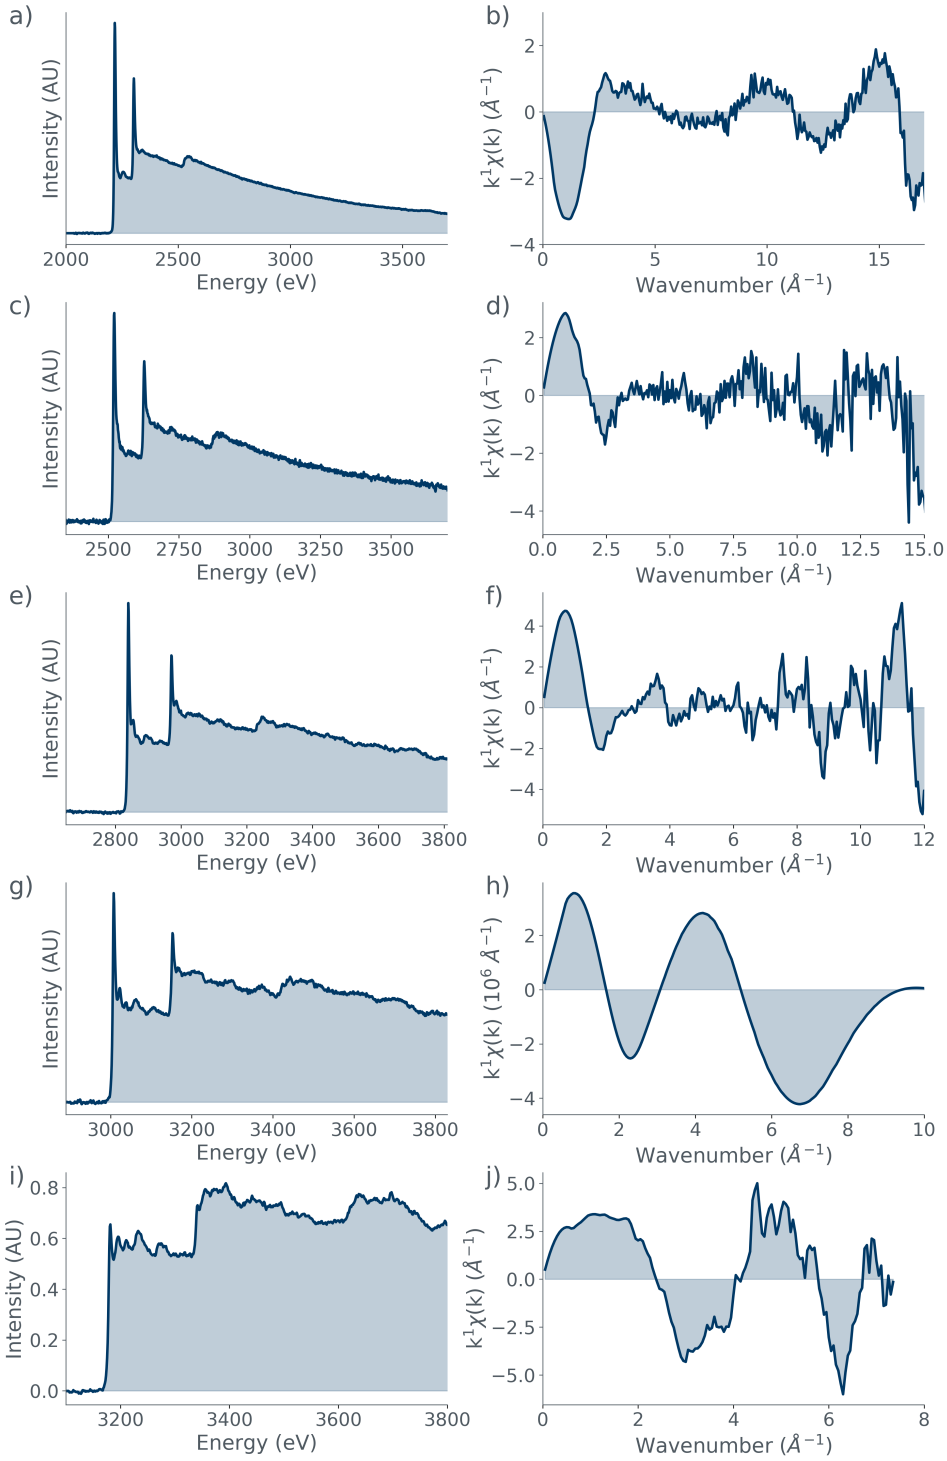

Figure 5.1: The data from the transition metals, showing the edges and then the *k*-space data, all of which show very poor and noisy oscillations: a,b) zirconium; c,d) molydenum; e,f) ruthenium; g,h) rhodium; i,j) palladium. The energy ranges used to extract the k data were 100-200, 150-1190, 170-900, 100-700 and 150-550 eV post-edge, respectively.

5.1.2 Silicon

The data in this section was collected by Ian MacLaren and comes from a silicon wafer tilted off the 001 plane so as to reduce EELS oscillations. EELS oscillations arise due to Bloch waves in a crystal, which are excited in different ways depending on the crystal orientation close to a strongly diffracting crystal plane. This can have serious implications for interpreting the structure since it can completely change the near-edge structure or the apparent ratio of different atoms in a material [207, 208]. Therefore, close to a pole, results change significantly with orientation and moving the sample stage or scanning the beam can give you position dependent changes in EELS. The solution is simply to significantly tilt the sample away from any pole or low-index plane [209].

After initial processing, the silicon data were summed and then passed through a low-pass numerical filter, with the hope of reducing the noise in the spectrum without removing the EXELFS oscillations. The numerical filter replaces the value in each spectrum channel by the average number of counts per channel, in an interval w units wide, and centered on the channel in question. Four filters were attempted, with the $w = 5$ spectrum chosen for further analysis. The $w = 10$ lost too much data and also had its edge energy shifted due to excessive smoothing. The $w = 1, 2$ were insufficiently different from the original, unfiltered spectrum. The final spectrum can be seen in fig. 5.3.a. The transformation into R-space (fig. 5.3.c) was performed over the k-range $4 < k < 12$ Å^{-1}. This is a relatively small k-range, but increasing the values only gave psuedo-structure in the lower part of the radial distribution function (RDF) thanks to the noise at the higher end of the k-range (fig. 5.3.b). Adjusting the k-weighting resulted in no improvement.

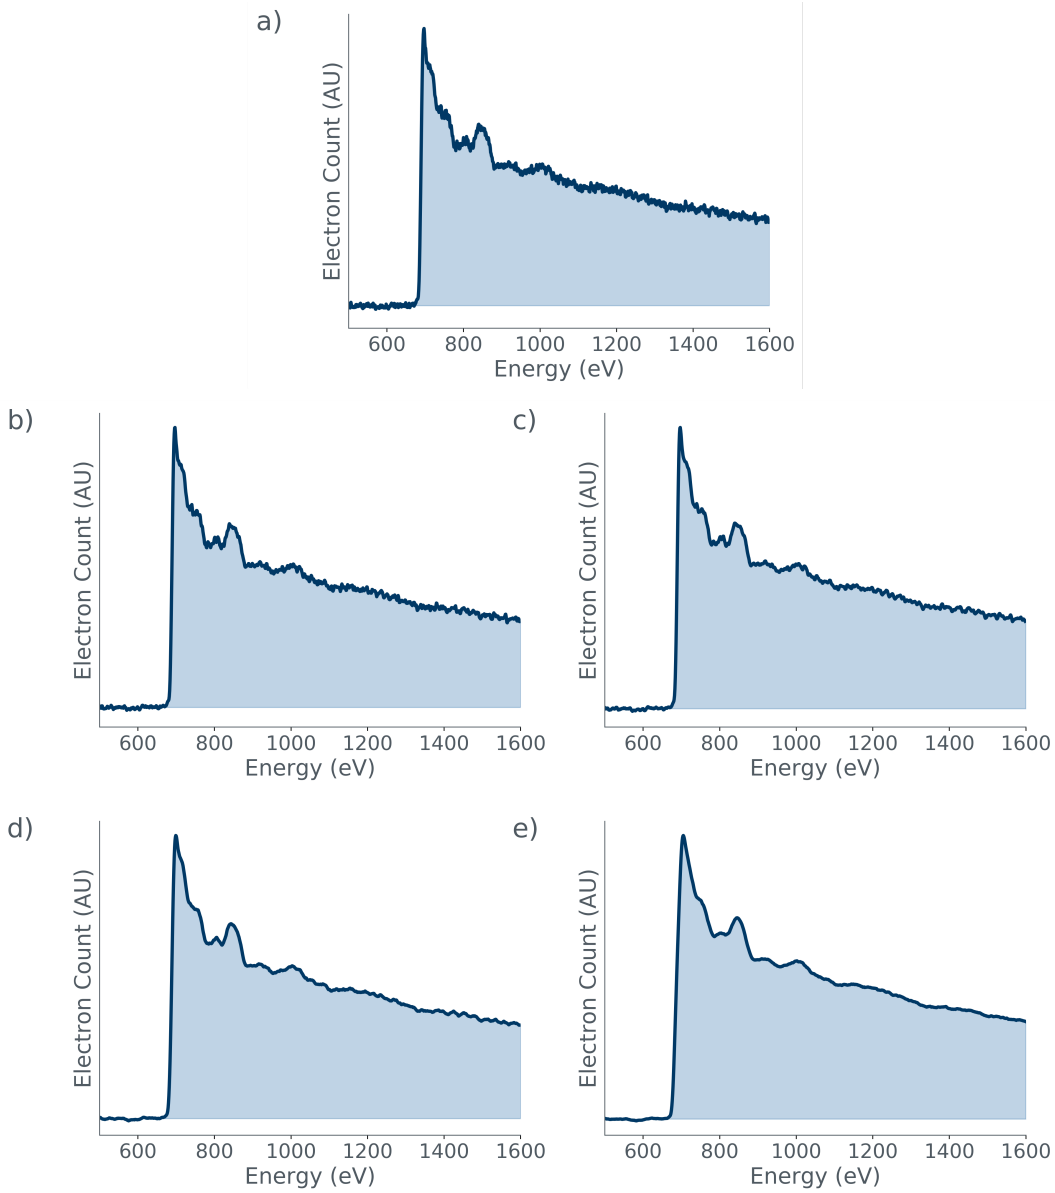

Figure 5.2: The effects of a numerical filter on EELS data, which replaces the value of each spectrum channel by the average counts per channel in an interval w units wide, centered on the channel in question: a) the original deconvolved abd background subtracted spectrum; b) $w = 1$; c) $w = 2$; d) $w = 5$; e) $w = 10$.

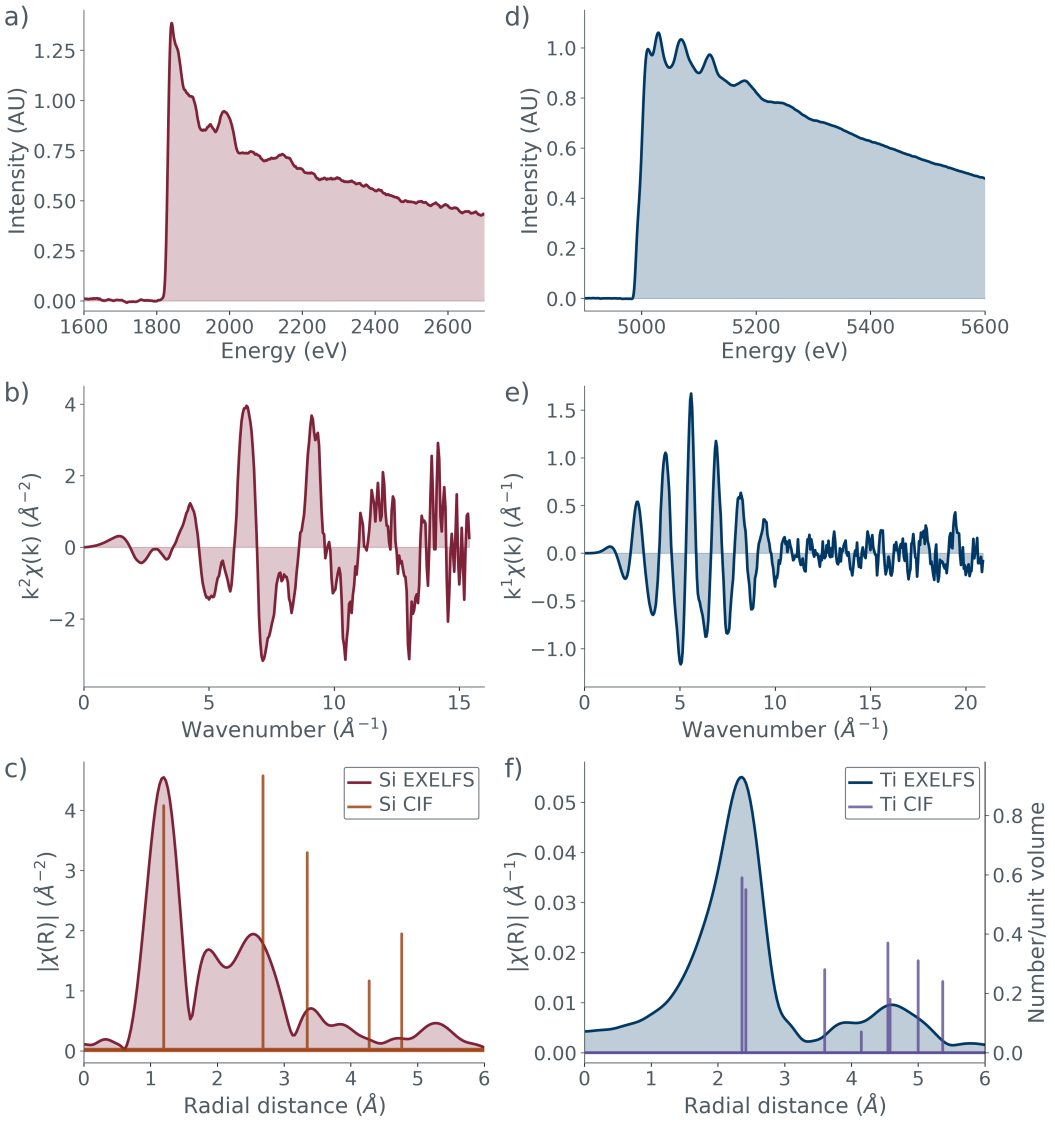

Figure 5.3: Results from the silicon and titanium samples examined: a) the silicon K-edge, showing a significant amount of noise in the post-edge tail; b) the k-space plot, showing some noisy oscillations, note that this has been weighted by k^2; c) the R-space plot (burgundy line), with the accepted positions of silicon atoms overlaid (orange spikes). These spikes have the correct relative distances but have been shifted by 1.155 Å to align with the inital peak of the EXELFS data. A k-range of 4-12 $Å^{-1}$ was used. The y-axis is in number of atoms per unit volume, and can be seen on the far right of plot f.; c) the titanium K-edge; b) the k-space plot, showing clear oscillations in the first half of the plot before descending into noise. Note that this has been weighted by k; c) the R-space plot (blue line), with the accepted positions of titanium atoms overlaid (purple spikes). A k-range of 3-11 $Å^{-1}$ was used.

5.1.3 Titanium

Two sets of titanium data were provided by Jamie Hart at Drexel University, PA [210]. The first was made up of ten spectra. These spectra were then run through median and Gaussian filters, and then combined. The second set of data were made up of 21 spectra, taken at a higher current to the previous data set. These spectra were then deconvolved, aligned and summed.

The combined spectrum from the second data-set was the one used for analysis as the EXELFS oscillations were stronger. It was processed using Athena and the transformation into R-space was performed using $3 < k < 11$ Å^{-1}. The RDF was compared to a histogram produced by CrystalMaker [211]. These data can be seen in fig. 5.3. The results were much more promising than the ones produced for the silicon data, and so it was decided to try and produce a fit using Artemis.

Initial results were promising. As mentioned in section 3.6.1, there are a variety of ways to establish how successful a fit is and they were consistently positive, even as multiple scattering paths were added. None of the values were unphysical; the fit appeared good in the graph produced and none of the correlations were too high.

Figure 5.4 shows the scattering paths (fig. 5.4.a) that contribute to the experimental RDF (fig. 5.4.b) – i.e. the expected form of the RDF were only that neighbouring atom present. Combining paths produces the actual RDF. As paths are added to the fit, the expectation is that the fit (also infig. 5.4.b) grows closer to the experimental data. This is indeed the case for this data.

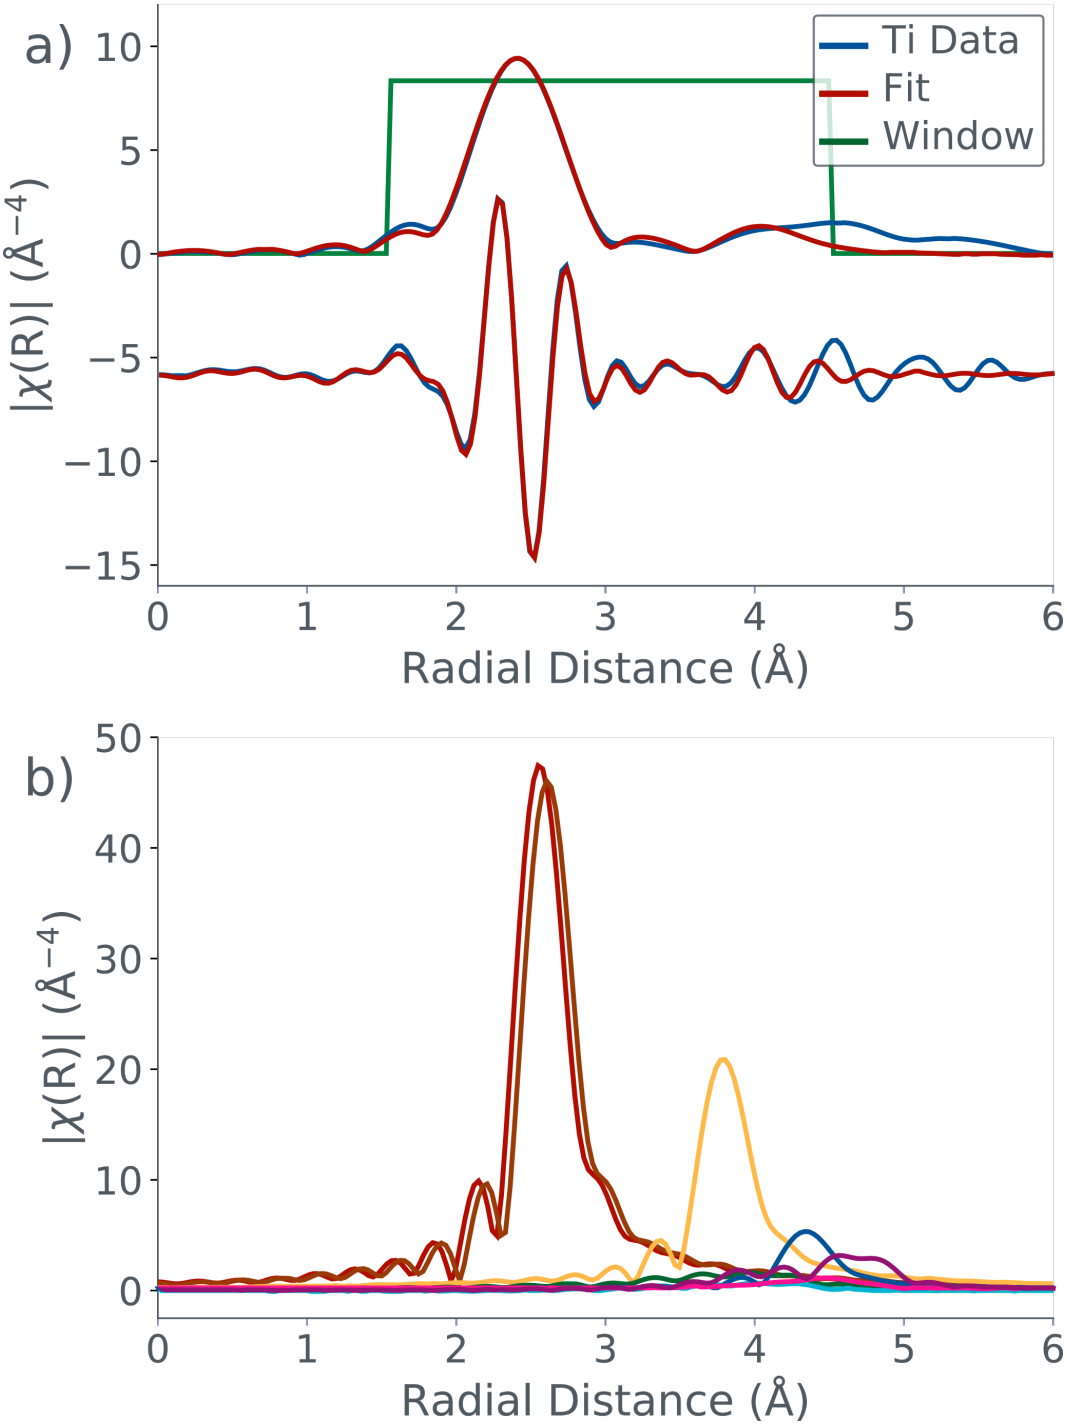

Figure 5.4: Graphs from artemis analysis of titanium data: a) the fit made with Artemis, using the scattering paths below to match the EXELFS data; b) the scattering paths as calculated by FEFF. Each line represents a path from a different neighbouring atom. The importance of the paths to the fit decreases with their peak intensities.

5.2 Silicon

This sample is a commercial silicon wafer topped with hafnia and was investigated for several reasons. The first was to see if collecting mutliple spectra to be combined would have a positive effect, as was suggested by the results of the initial work, before investigating the GW coating samples. The second was so that there was a Si K-edge taken from silicon that could be compared with the one taken from the GW coating silic. Finally, it was felt there was some benefit in collecting the silicon data purely for its own sake – although the current aLIGO coating uses silica, silicon is being suggested as a future GW coating.

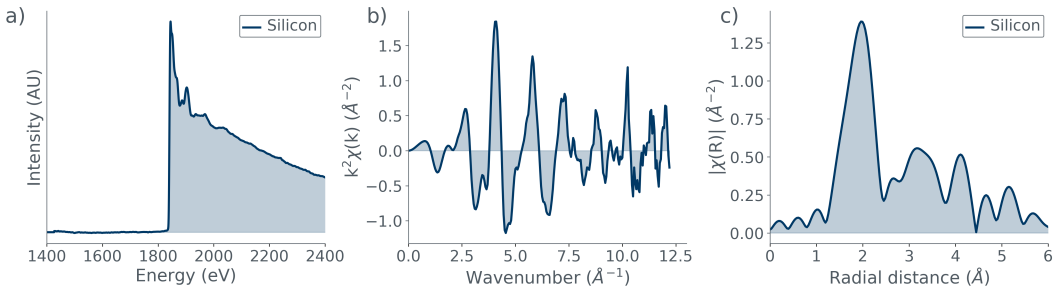

Figure 5.5: Results from the silicon wafer: a) the silicon K-edge; b) the k-space plot, extracted from 150-474.5 post edge; c) the R-space plot made using a k-range of 3-10.28 Å^{-1}.

The data are shown in fig. 5.5. The transformation into R-space was performed using $3 < k < 9.5$ Å^{-1}. The EELS data is much cleaner (fig. 5.5.a), than the first sample, and the k oscillations are clearer and less noisy. The k-range, however, is still limited and the R data (fig. 5.5.c) shows signs of noise at the lower end because of the noise at the upper end of the k data (fig. 5.5.b). A third pass at this data would benefit from adjusting the energy window over which the EELS data were collected, by extending further post-edge.

5.3 Silica

Amorphous silica makes up a huge fraction of LIGO's mirrors; it is the low refractive index layer of the coatings; the fibres the mirrors are suspended with are made from silica [164] and the test masses themselves are made from silica. As stated in section 1.5, the tantala layer dominates the loss contribution of the coatings, however the silica still contributes to the noise. As the required sensitivity of the detector increases, even smaller noise contributions must be considered.

Silica films can be made by a variety of methods; ion-beam sputtering; magnetron sputtering [212]; reactive electron beam evaporation; plasma ion assisted deposition [213]; sol-gel synthesis; chemical vapour deposition; or physical vapour deposition [214]. It is well known that the method used affects its properties, but the reason for this is much less well established [215]. Silica has applications far outside GW mirrors: it is used in hemispherical resonator gyroscopes [216], atomic clocks [217], semiconductors [218] and photovoltaics [219].

All of the data in this section were processed in the same manner, using the basic EELS processing as described in section 3.4. Once the individual spectra had been processed, they were summed and then passed through a low-pass numerical filter ($w = 5$), in order to reduce high-frequency noise. The resulting spectra were then prcoessed in Athena. All of the k and R-space data were appropriately weighted by k^2, to allow for better comparison.

5.3.1 aLIGO Test Sample

In the case of the aLIGO test sample, there is not a lot of information other than its composition. It dates from approximately 2011 and is thus likely to be a test composition for the aLIGO coatings. Two[1] sets of data were taken from this sample; one from the silica substrate and the other from one of the deposited silica layers. The intention was to compare the two, to see if differences were visible in their EXELFS.

5.3.1.1 Deposited Silica Layer

This data can be seen in fig. 5.6. There are clear, strong oscillations for both the oxygen and silicon K-edges (visible in fig. 5.6.b&e, respectively). The oxygen R-space plot (fig. 5.6.f) has pseudo-structure below the first major peak. The transformation into R-space was performed using $3 < k < 10.85$ Å^{-1} for the silicon edge and $3 < k < 10.2$ Å^{-1} for the oxygen edge.

[1]In fact, a third set of data were also taken from this sample, from the tantala layer. Argon bubbles are visible in these data, as in the sample from section 4.3.1. These data were not included in Chapter 4 as it was felt this data would add little to that chapter when so little is known about how this sample was made.

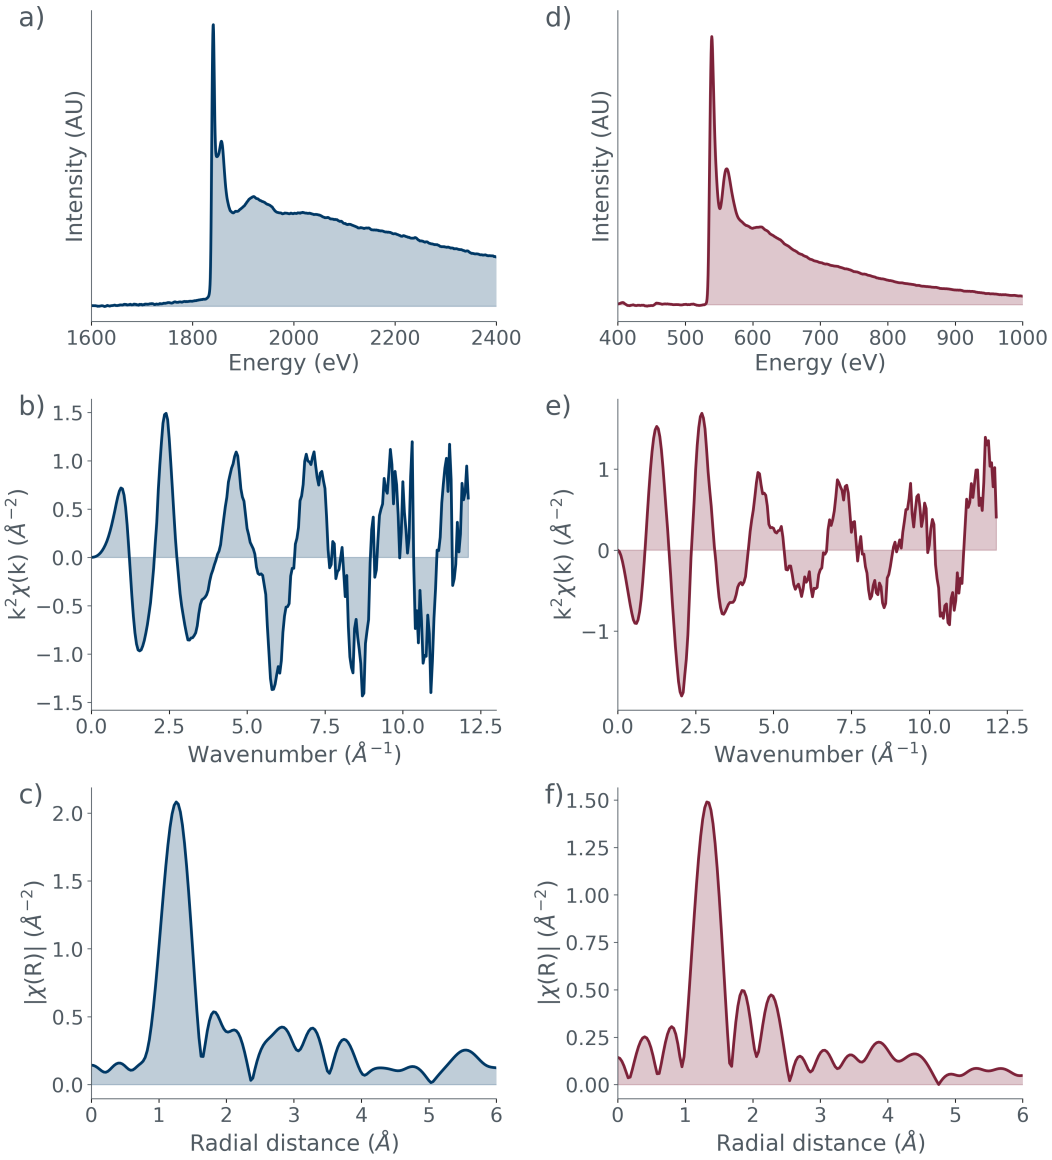

Figure 5.6: Results from the deposited silica layer from the aLIGO test sample: a) the silicon K-edge; b) the silicon k-space plot, extracted from 150-461.5 post edge; c) the silicon R-space plotmade using a k-range of 3-11.5 Å^{-1}; d) the oxygen K-edge; e) the oxygen k-space plot, extracted from 150-1005.273 post edg; f) the oxygen R-space plotmade using a k-range of 3-11.08 Å^{-1}.

5.3.1.2 Substrate Silica

This data can be seen in fig. 5.7. There are clear oscillations for both the oxygen and silicon K-edges (visible in fig. 5.7.b&e, respectively) but both become noisy at the higher end of the k-range. The silicon R-space plot (fig. 5.7.c) has pseudo-structure below the first major peak.

The transformation into R-space was performed using $3 < k < 10.3$ Å^{-1} for the silicon edge and $3 < k < 10$ Å^{-1} for the oxygen edge. The results from both the deposited and substrate silica will be discussed in depth in section 5.4.3.

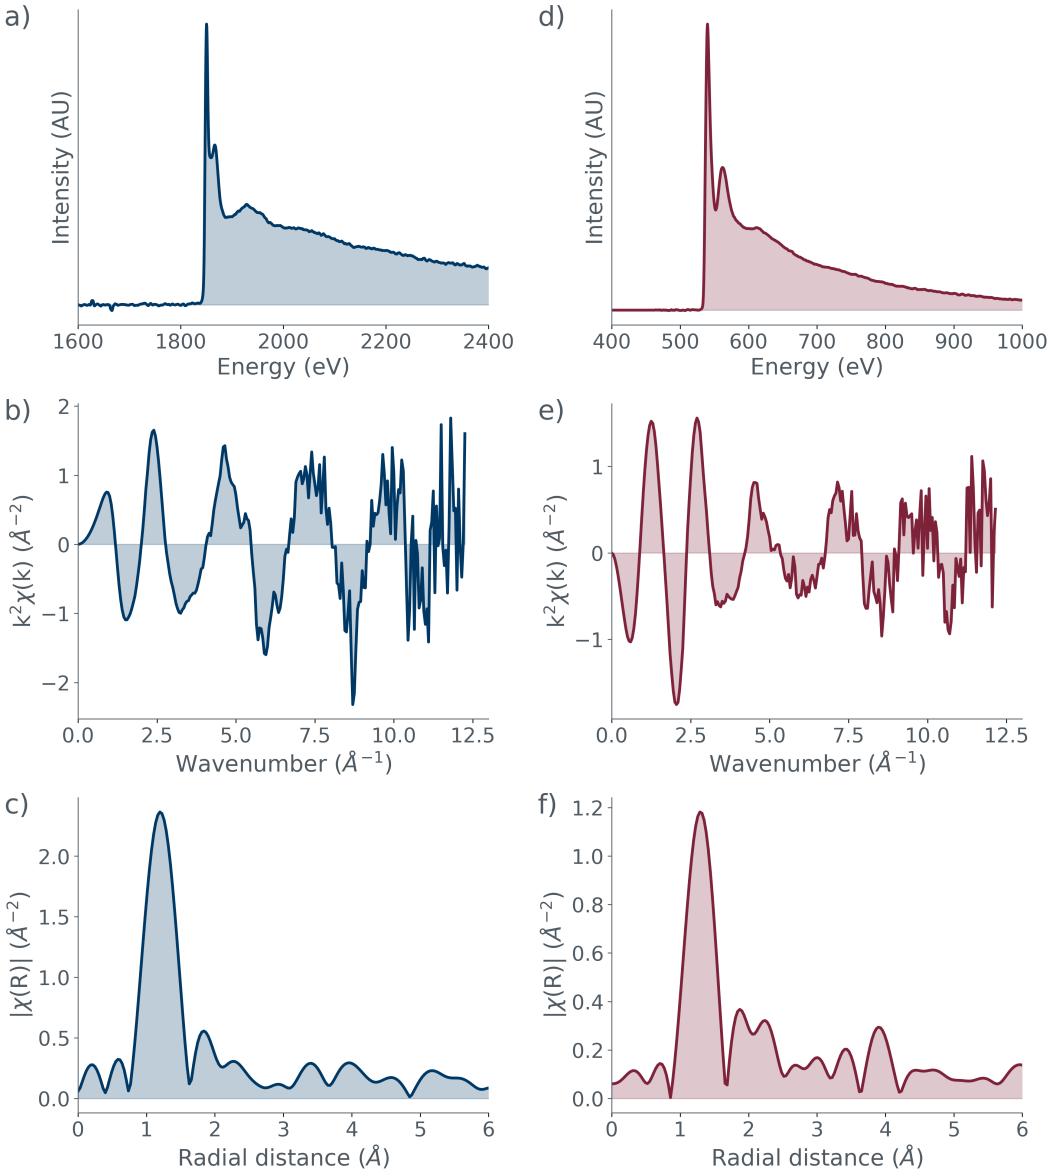

Figure 5.7: Results from the substrate silica layer from the aLIGO test sample: a) the silicon K-edge; b) the silicon k-space plot, extracted from 150-479 post edge; c) the silicon R-space plot made using a k-range of 3-10.5 Å^{-1}; d) the oxygen K-edge; e) the oxygen k-space plot, extracted from 150-1005.489 post edge; f) the oxygen R-space plot made using a k-range of 3-10.7 Å^{-1}.

5.3.2 Amorphous Silica and Zirconium Doped Tantalum Pentoxide Multilayer

This sample was created by Martin Hart and was made in the same way as the current mirror coatings on the aLIGO mirrors. It is a cross-section of the multilayer coatings and therefore contains both tantala and silica. A HAADF image of this sample can be seen in fig. 1.5. For the oxygen K-edge, spectra were taken from each silica layer and then summed together for each dataset, and then all the resulting spectra were also summed. The silicon K-edge data were processed slightly differently. These datasets were made up of only high- and low-loss spectra, with no HAADFs or SIs. Consequently, the spectra were deconvolved and then summed with no additional processing.

The data are visible in fig. 5.8. The transformation into R-space was performed using $3 < k < 10.85$ Å^{-1} for the silicon edge and $3 < k < 10$ Å^{-1} for the oxygen edge. This sample would benefit from more data on the silicon K-edge. The results can be seen in fig. 5.8.

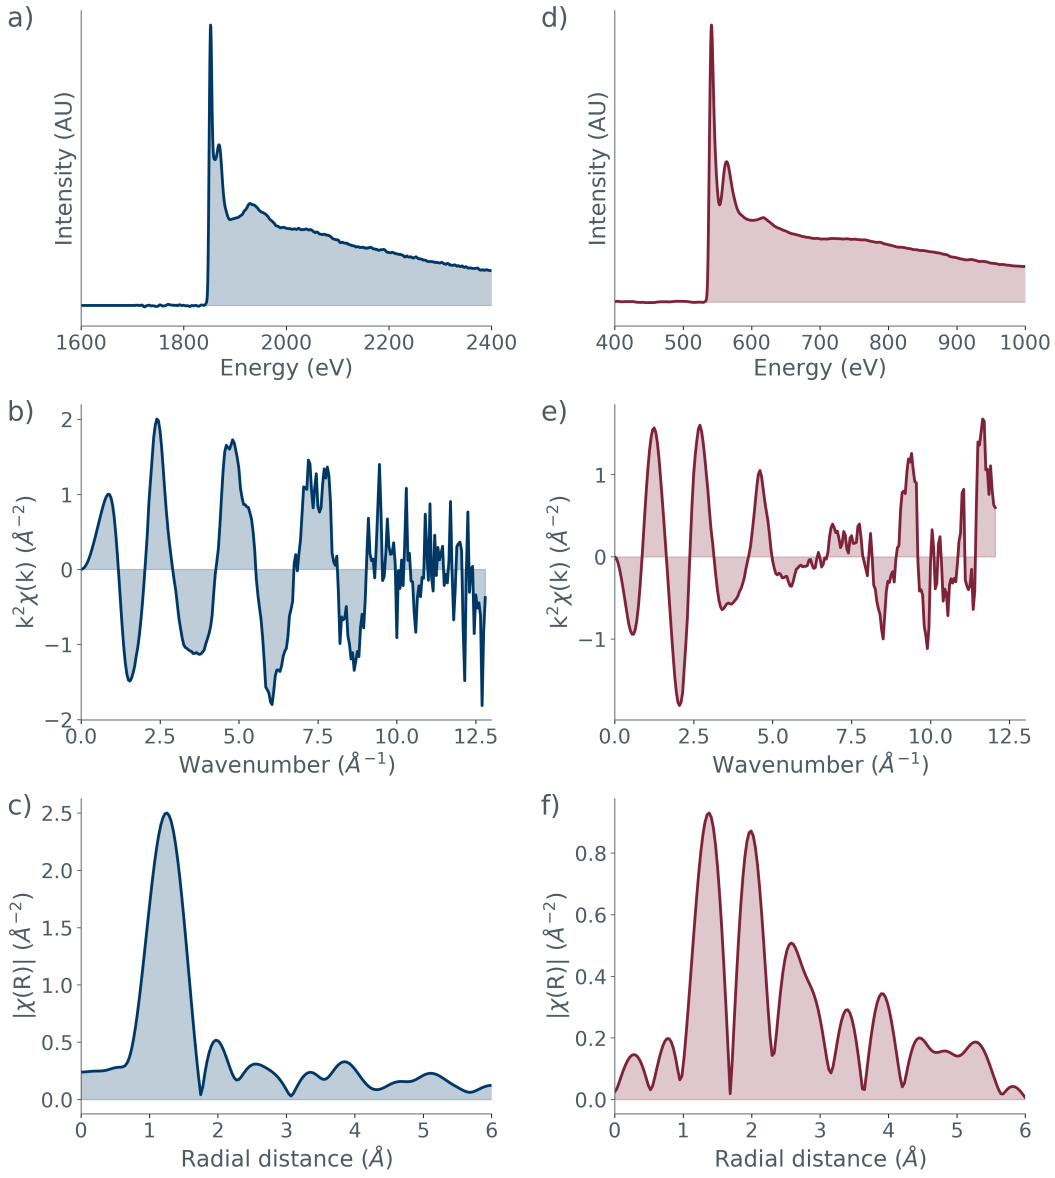

Figure 5.8: Results from the multilayer sample: a) the silicon K-edge; b) the silicon k-space plot, extracted from 150-529.5 post edge; c) the silicon R-space plot made using a k-range of 3-10.5 $Å^{-1}$; d) the oxygen K-edge; e) the oxygen k-space plot, extracted from 150-997.274 post edge; f) the oxygen R-space plot made using a k-range of 3-10.03 $Å^{-1}$.

5.4 Discussion

This section will first examine the initial work, and what was learnt from it, before moving onto discussing the second silicon sample and the silica samples. It will then compare the deposited and substrate silica layers from the aLIGO

test sample, before finally comparing the three silica datasets and putting the results into a wider context.

5.4.1 Initial Work

If success was measured by deriving atomic positions, neither of the first two investigations were particularly succesful. Despite this, there are still things to be learnt from them. The transition metal samples were originally collected with the intention of developing a technique that could acquire complementary data for XANES, which had better spatial resolution in order to investigate micro- and nano-structure. While too noisy and with edges too close together to be useful for EXELFS these spectra will be useful in developing semi-empirical standards, like those used in Chapter 4. Currently, there are no studies of any L-edges above technetium and no published Hartree–Slater cross-sections for EELS L-edges after molybdenum.

Similarly, while the silicon sample did not produce useful data for analysing atomic positions, it shows that improvements to the EELS set-up had a positive effect. The relatively small size of this dataset was prohibitive. More data would be extremely beneficial for reducing the noise and improving potential analysis.

The titanium data, by comparison, was extremely promising. The fit in fig. 5.4 is incomplete, since it should follow the data more precisely across the entire fitting window – but this was deliberate. The intention here is not to create a perfect fit but to show that it is possible to collect EELS data of a high enough quality to perform this kind of analysis. Artemis is designed for XAS data, and uses the EXAFS equation. Thus, a perfect quantitative fit for EELS data is of little use in isolation (i.e. only one dataset). The EXAFS equation needs to be adjusted to quantitatively describe EELS data, and it

is not possible to incorporate this into the Demeter software. The titanium sample does not need to be well characterised. It was a test to see if the silica samples of the next section could be examined using this technique.

This initial work laid the foundations for the data collection and analysis in the rest of this chapter. While the attempt to classify transition L-edges was misguided – the edges are far too close together to perform useful analysis using these techniques – it, alongside the silicon work, showed how crucial collecting multiple datasets were in order to reduce the noise as much as possible. The titanium work proved definitively that it is possible to perform meaningful and quantitative EXELFS analysis. Unfortunately, the specific improvements in the titanium data's SNR could not be replicated for the silica data collection. This is because the set-up at Drexel is different than in Glasgow. For the titanium data, a direct detection camera was used. Rather than the multi-step process involved in a CCD (section 2.2.3), in direct detectors the beam electrons are directly recorded without first being converted to photons. Consequently, they offer greater resolution and SNR than a traditional CCD [210]. The silica data collection focussed on collecting mutliple datasets and using the improvements in set-up discussed at the beginning of this chapter to improve SNR.

5.4.2 Silicon

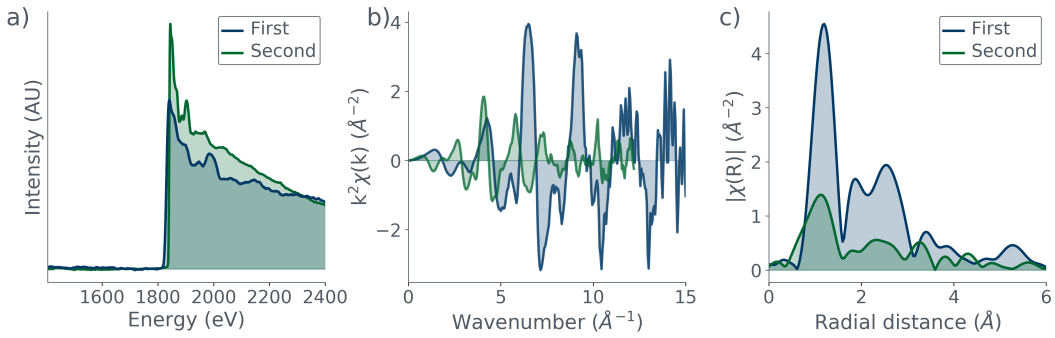

Figure 5.9: Results from both of the silicon samples: a) the silicon K-edge; b) the silicon k-space plot; c) the silicon R-space plot.

The data from the two silicon samples in fig. 5.9 show several differences. The edge shape is different for the two samples; the k-space oscillations are of very different intensities and shape; and the R-space plot shows definite differences in shape. The second silicon R-data has been offset by 0.85 Å, so that a better comparison between the two plots can be made.

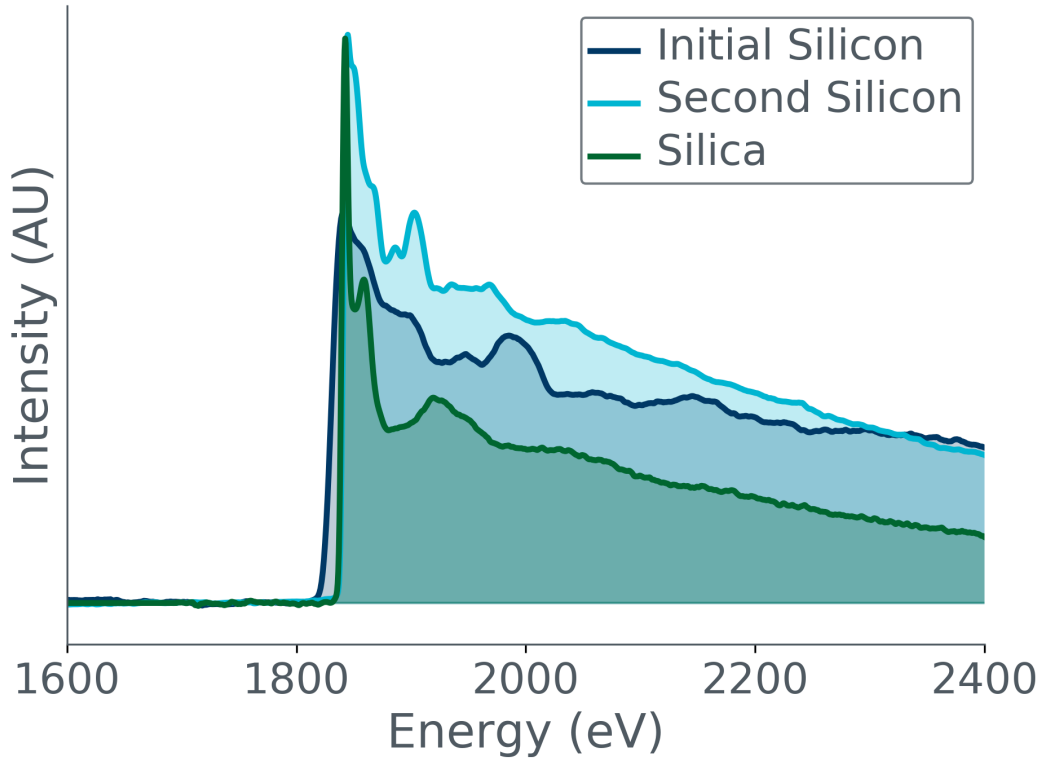

Figure 5.10: The silicon data alongside silica data from the multilayer sample (section 5.3.2).

Adjustments to the data collection were made for the second silicon wafer, based on what was learnt from the first; more datasets were collected and summed together. This can be seen in the smoother post-edge tail of the second attempt. The difference in edge shape can be seen more clearly in fig. 5.10, which also includes the silicon K-edge from the multilayer silica sample. The second silicon dataset bears more similarity to the silica dataset than the initial silicon sample, suggesting the second silicon sample may have begun to oxidise. This casts doubt on the RDF for this sample. Given that the first sample did not match the expected structure of silicon either (see fig. 5.3); it is difficult to draw any conclusions from either dataset. An Artemis fit was not attempted both because of this and because it provides no information in isolation, as discussed above.

5.4.3 aLIGO Test Sample; Inter-Layer Comparison

The data discussed here can be seen in fig. 5.11. The data for the deposited layer is made up from more individual datasets than the substrate layer, hence its slightly lower noise. The two samples have very similar results, with the exception of the edge intensities (fig. 5.11.a&d), which are of no consequence. The oxygen K-edge of the deposited layer seems to decay back into background faster than the substrate layer but otherwise their shapes are similar.

The silicon edge k-space oscillations (fig. 5.11.b) appear consistent, with a slight shift in the position of the substrate oscillations after 5 $Å^{-1}$. The RDF (fig. 5.11.c) is much less consistent, with the substrate and deposited R-space plots varying after the initial peak[2].

The oxygen edge k-space oscillations are very clearly consistent (fig. 5.11.e). The substrate data is noisier at higher k but even then, it still follows the general trendline of the deposited data. This follows through to the R-data (fig. 5.11.f). There are some differences in intensity but the positions of the peaks are very clearly consistent up until 4 Å. They begin to diverge after this and collecting more data, particularly from the subtrate would enable a more definitive answer as to why this divergence occurs. However, given how closely they align up until this point, and how noisy the substrate k-space data is at higher k-values, it is extremely likely that this difference is simply due to the higher noise in the substrate data.

The large similarities between the RDFs of the oxygen K-edge is extremely promising, and suggests that the structure of the two silica types is the same. Admittedly, the silicon K-edge data contradicts this somewhat but this is

[2]A reminder of nomenclature; the large peak around 1 Åis the 'initial' peak, despite there being some small peaks prior to it. This is because these peaks are artefacts in the Fourier transform and not representative of real structure.

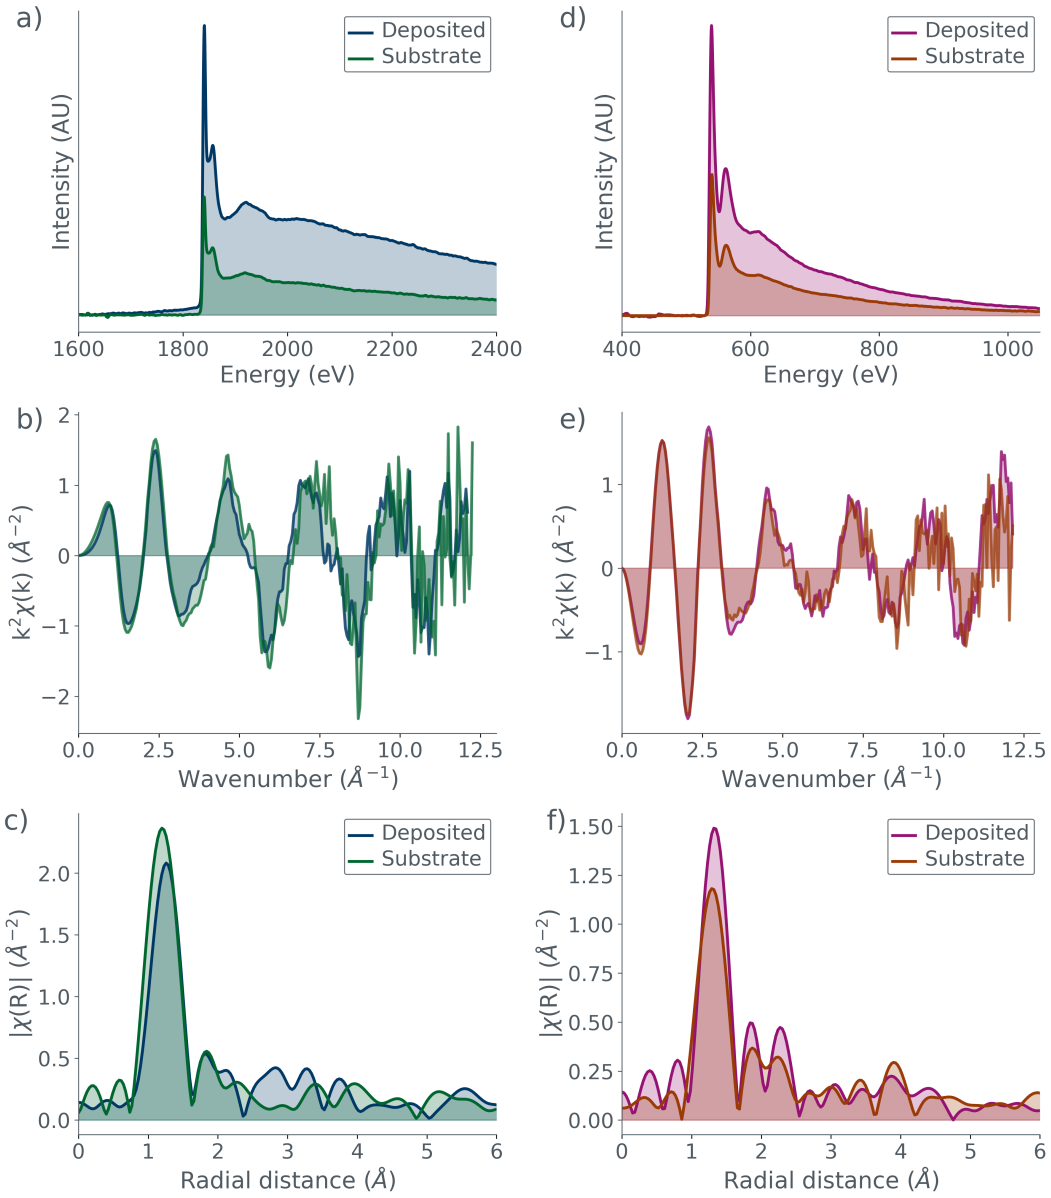

Figure 5.11: Results from both of the aLIGO test sample layers: a) the silicon K-edge; b) the silicon k-space plot; c) the silicon R-space plot; d) the oxygen K-edge; e) the oxygen k-space plot; f) the oxygen R-space plot.

possibly due to the noise. Both of the substrate datasets are much noisier than the deposited data but the silicon data is the noisier of the two. For EELS, noise becomes a bigger issue as the energy increases. Collecting more substrate data, especially over the silicon K-edge would be extremely beneficial, and would confirm if the discrepancy between the two RDFs real or just a result of noise.

5.4.4 Final Discussion

The multilayer sample shows clear differences between the silicon K-edge and the oxygen K-edge (fig. 5.12). The k-data for the oxygen edge (fig. 5.12.e) has an odd and obvious dip in the middle of its oscillations and the differences between the R-plots is stark. The R-plot for the silicon edge is extremely similar to both of aLIGO test samples' but the data from the oxygen edge shows clear and definite structure (fig. 5.12.c&f, respectively). At first glance, this may suggest that the oxygen data is suspect – obviously the two edges contradict each other but only one of them agrees with the other silica data – but it is actually the silicon data that should be considered less reliable. As previously discussed, the silicon K-edge data for the multilayer sample was collected in a different manner than the rest of the data; only spectra were collected rather than the full complement of data. This methodology, alongside only three datasets being collected, was because there were issues with the sample charging excessively. This means that not only could the normal processing and noise reduction not be performed, but the sample size was smaller for the silicon edge. Additionally, the oxygen data were taken over two sessions, several months apart.

Bond distances calculated using CIF were not imposed onto the results visible in fig. 5.12 or fig. 5.11, as was done in fig. 5.3 for two reasons. The first is that there are multiple structures of crystalline silica and it was not known which to choose to best compare with the amorphpous silica in the GW coating samples. Secondly, the intent of this work was not to definitively calculate the length of scattering paths but to make a qualitative comparison between the different silica samples. It is clear from oxygen k-edge plots in fig. 5.12 that the deposited and substrate silica from the aLIGO test sample have very similar structures but that the multilayer's structure is different.

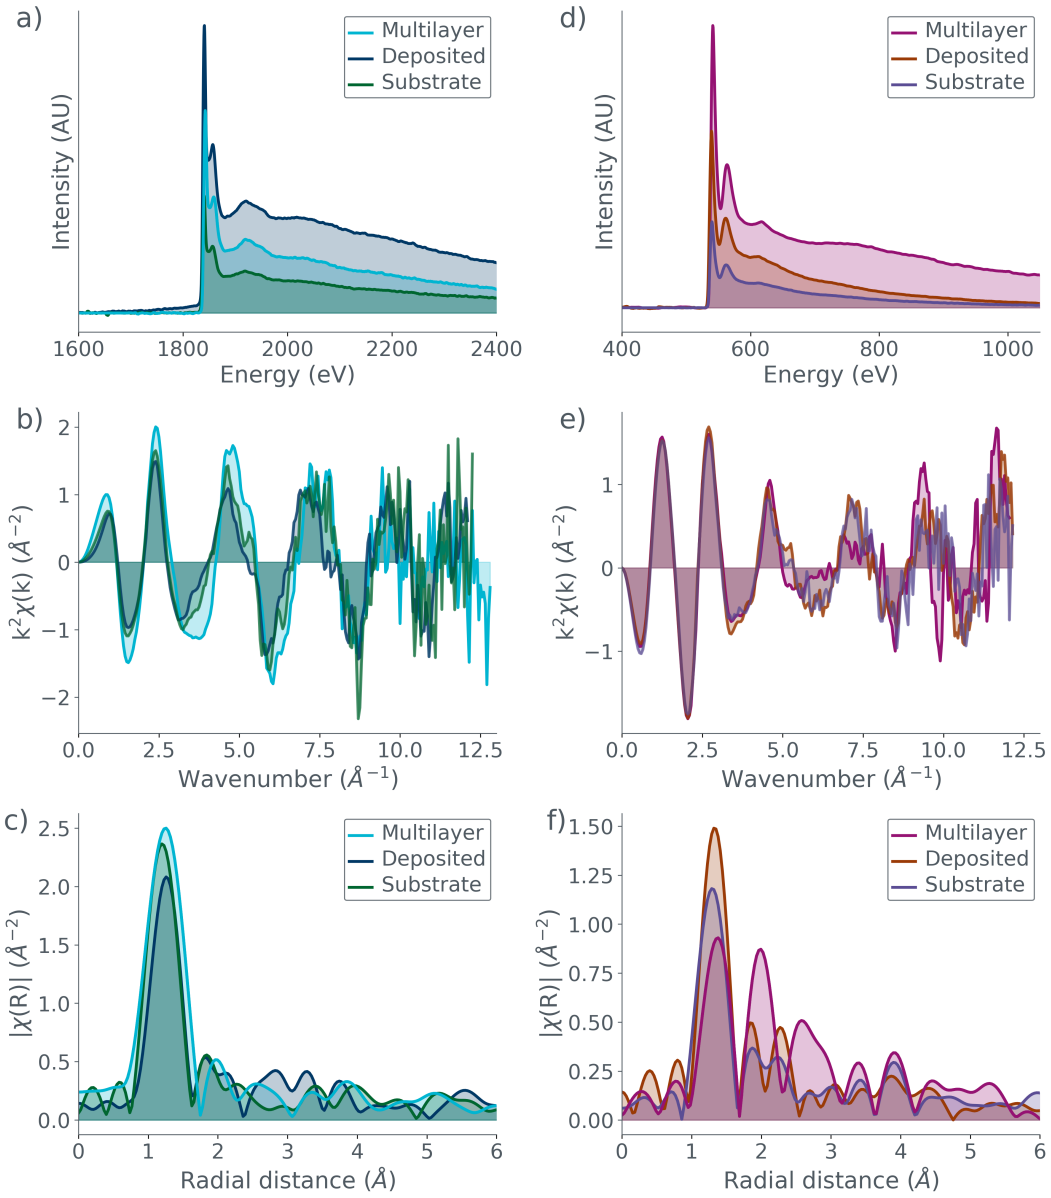

Figure 5.12: Results from all three silica samples: a) the silicon K-edge; b) the silicon k-space plot; c) the silicon R-space plot; d) the oxygen K-edge; e) the oxygen k-space plot; f) the oxygen R-space plot.

The results from the oxygen edge of the multilayer sample mean two things; firstly, that silica made using different methods have differences in their short-range structure and, secondly, that those differences are identifiable in their EXELFS. At this point, the logical next step would seemingly be to move to Artemis and attempt to fit it to a silica CIF. Ultimately, however, this was decided to be redundant. Although the titanium data were fit using Artemis

this was a proof of concept and as discussed, Artemis uses FEFF6, which has no provision for EELS data. There have been provisions for EELS since FEFF8, but it was not optimised until FEFF10, which was only released very recently. It is also the first FEFF code to be open access. Artemis works such that it requires very little interaction with the FEFF6 code itself and it is not possible to implement newer FEFF versions into Artemis. Consequently, running and analysing FEFF10 caclulations on the multilayer oxygen data were felt to be beyond the immediate scope of this chapter.

It is very well established that how amorphous silica is made affects its properties [53, 212, 213, 220] but the relationships between structure and properties; and structure and annealing and synthesis is much less well understood [215]. While other techniques such as X-ray and electron pair distribution function measurements; Raman spectroscopy; nuclear magnetic resonance; and X-ray diffraction can and have been used [215, 212] to investigate the structure of the coatings, much is still unknown. Silica is unlikely to be the low refractive index coating for future GW detectors (amorphous silicon is likely to take its place [58, 57]), understanding the link between these things is still of upmost importance as new coatings are tested and chosen. The technique here is far from limited to amorphous silica, as seen in section 5.4.2. Being able to determine the SRO of the coatings using EXELFS may explain why the process by which amorhpous silica coatings are made affects its properties. This would, in turn, allow the fine tuning of those properties.

The next steps for continuing the work of this chapter are fairly simple. The first is, as always, to collect more data. The ideal candidates for this work are amorphous silica samples, with known deposition techniques and environments; post-deposition annealing; and properties. Multiple datasets should be taken on these samples, focussing on increasing the SNR as much as possible, and

choosing edges with long post-edge tails. Performing k-dependent smoothing might also be helpful, although it is crucial that this is done with great care so as to not remove useful data. Using electron counting and direct detection would also hugely reduce the shot and amplifier noise. A combination of sensible data collection and processing, and using different detectors will hopefully result in oscillations up to 15 Å^{-1}, or even higher with lower noise. This, in turn, should make it possible to identify if there is a link between the structure and these properties/synthesis environments. The second is to incorporate FEFF10 into the project, performing calculations that will allow for characterisation of the structure of the films. This technique is not limited to silica, it can also be used on the tantala layers, as well as potential coatings for the next generation of GW detectors. It has has wide ranging applications for characterisation outside GW detectors. A third avenue for further work is comparing EXELFS results with those from different techniques for investigating amorphous materials, such as electron diffraction and fluctuation microscopy [221].

This results of this chapter also have ramifications far beyond GW mirror coatings. ELFS as a technqiue is used infrequently due to its historically poor SNR when compared with EXAFS. The advances of recent years in EELS signal quality means EXELFS will no longer need to be relegated to second choice but will be a true competitor to EXAFS; truly a 'synchrotron in a microscope' [123].

5.5 Conclusion

The history of EXELFS was discussed, followed by an explanation of the motivation for using EXELFS to investigate GW mirror coatings, specifically the amorphous silica layer. This was followed by a record of the inital work carried out. This initial work focussed on transition metals, silicon and titanium, and was done in order to develop the technique and as a first test of 'usefulness' of EXELFS as an analysis technique following recent developments in improving the SNR of EELS data. The titanium data were collected by Jamie Hart at Drexel University using an alternative data collection technique and was by the far the most succesful of this early data. A second attempt at analysing a silicon wafer was made, following the initial work. While providing better quality EXELFS, the second sample had oxidised. This makes it difficult to compare the two samples. Having established that EXELFS was a viable analysis tool, the silica samples were investigated. The first silica sample was an old aLIGO test sample, with both deposited and substrate silica. The results from this sample showed little difference between the deposited and substrate silica, and showed little in the way of structure. The second sample was a multilayer sample of amorphous silica and zirconium-doped tantalum pentoxide. The oxygen K-edge of this sample showed definite structure.

The results of the different samples were compared and discussed. The results suggest it is indeed possible to see differences in the structure of samples created with different methods using EXELFS and that EXELFS is becoming much more viable as an analysis technique. Potential next steps were discussed, including collecting more data on samples with known properites and synthesis methods to compare their EXELFS, as well as using FEFF10 to analyse the EXELFS.

6

Conclusion

Noise reduction and elimination in GW detectors play a huge role in GW astronomy. Less noise means the detectors will be more sensitive, and therefore able to detect smaller GW. There are many sources of noise in the detectors but the most important for this thesis is the thermal noise in the coatings of the interferometer mirrors. This noise dominates the detector noise at its most sensitive frequency. In order to increase the sensitivity of the detectors, this noise source must be reduced. The first step is fully understanding the structure of the coatings. The coatings are multilayers of amorphous silica (low refractive index) and titania-doped tantalum pentoxide (high refractive index and low dielectric constant) [54].

The first technique used to study the structure of the coatings was a combination of DualEELS and high angle annular dark field (HAADF) imaging to perform absolute spatial quantification. This is detailed in Chapter 4. Seven

different samples were investigated, with the intention to characterise bubbles that appear in GW coatings post-annealing. Five of the samples were doped tantala – two with zirconia and three with titania. One of the zirconia samples was annealed in oxygen but the rest of the tantala samples were annealed in argon. The other two samples were not GW coating samples but rather Zircaloy-4. Both were deliberately ion-implanted with noble gas beams at the MIAMI facility in Huddersfield, in order to simulate the irradiation that happens under reactor conditions. One was ion-implanted with krypton and the other with xenon. The inclusion of these two samples was partly due to the pandemic and making use of the data available during this time, but these samples also provide a useful comparison to the GW samples by showing what might happen in other deposition environments. It was found that the bubbles form upon annealing and increase in size with annealing temperature. The bubbles' pressures and densities decreased with their diameter in all cases.

This work developed a technique that creates semi-empirical absolute EELS standards for noble gases, thereby allowing for counting atoms in bubbles a few nanometres wide. In the samples investigated here, enough bubbles were measured to perform a statistical analysis. It also allowed for the analysis of the properties of the bubbles, which showed that the densest were possibly liquid or even solid. The bubbles likely affect the optical properties of the films, although in what way is unknown. It is known that mechanical loss decreases with annealing but at higher temperatures the benefits plateau [134]. It may be that at this point the larger size of the bubbles causes issues. The inverse may also be true – that the bubbles positively contribute to the decrease in loss but at a critical size their contribution decreases. They may also be affecting the scattering loss [193], as the bubbles will not have the same low-absorption as the surrounding tantala glass.

The second technique investigated the other type of coating (amorphous silica), using EXELFS and is detailed in Chapter 5. There have been many recent advancements in the SNR of EELS at high-losses, but the question remained whether this meant that EXELFS was similarly improved. It was investigated whether EXELFS could be used to analyse the short range order (SRO) of the silica coatings, specifically if silica deposited in different ways had visible differences in its SRO. It was found that EXELFS could be used, and that there are some differences in the EXELFS of silica from different samples. The deposited and substrate silica from the same sample had only subtle differences, suggesting that there were no structural differences between the two. The multilayer sample was drastically different to those two, in the results from its oxygen K-edge. This suggests that samples made in different ways will have differences in their structure. Interpretation of these differences in a structural sense, is however, difficult at this point and will require modelling.

6.1 What's next?

Further work is possible for both the absolute spatial quantification and the EXELFS work, both within and outwith the context of GW mirror coatings. Both techniques would benefit from more data collection and then modelling work, although in slightly different ways.

For the absolute spatial quantification, taking more data from a wider range of samples fabricated in the same manner but annealed at different temperatures and for different lengths of time would be extrememly beneficial. This would help determine the rate at which the bubbles grow on annealing. It would also allow an investigation into the average energy barrier to diffusion for an argon atom in the glass; as well as for the average diffusion distance for a given annealing time and temperature; and what fraction of the total argon will form nanobubbles. Another avenue would be to collect data from samples of different compositions fabricated and annealed in the same way, to see how that affects bubble formation. These investigations could be done alongside ellipsometry and x-ray scattering measurements. To investigate theoretically how the bubbles affect the loss, both atomic modelling and molecular dynamics simualtions could be used provide information on the bubble formation. This technique is already applicable to materials outside GW coatings, as seen by the Zircaloy-4 samples. Further work in the nuclear field could investigate the effects of heavy noble gas deposits in nuclear fuel and reactor elements. It could also be used to examine helium implantation due to alpha irradiation in fusion reactor wall claddings.

For the EXELFS work, more data on silica deposited and annealed in known conditions, as well as with known properties (such as optical absorption) is needed. As many datasets as reasonably possible should be collected on these

samples, with tails of at least 400 eV post-edge. Using electron counting and direct detection would also be of great benefit for this technique to further reduce the noise. This would allow for a direct comparison between the EXELFS, the deposition and annealing conditions, and the properties of the silica. To link these three to the structure would require detailed modelling. Using FEFF10 is likely the best initial route for this, and to characterise the SRO of the silica from the EXELFS would make these results more widely relevant. This technique could also be used on the tantala layers, as well as any thin films made from materials with suitable edges.

6.2 Final Thoughts

This thesis had two, stated goals, the first to provide insight into the atomic structure of GW mirror coatings, and secondly to improve the EELS characterisation and quantification techniques used when analysing the coatings. The hope is that the further understanding of the structure of the coatings gained in this work, will help in the reduction of the thermal noise in detector due to the coatings; and that the techniques would be of use beyond studying the coatings.

The two results chapters tackled this in different ways. The first chapter presents a much more comprehensive technique and more definitive results. The second chapter acts more like the first step in establishing the use of EXELFS as an characterisation tool for the structure of the GW coatings. Although there do seem to be differences in the SRO of the silica samples made under different conditions, the lack of certainty around their fabrication and limited amount of data means these results are only preliminary. Despite this, the work is still promising in that it shows EXELFS is a viable analysis tool, with several avenues for improving it further. The absolute spatial quantification

method provides clear results that the argon known to be in the tantala coatings coalesces into bubbles upon annealing, and that the densities and pressures of the bubbles decrease with their diameter. It also presents a technique for creating a semi-empirical standard for absolute quantification when no standard already exists, as is often the case for noble gases.

Appendix: Annotated Bubble Maps

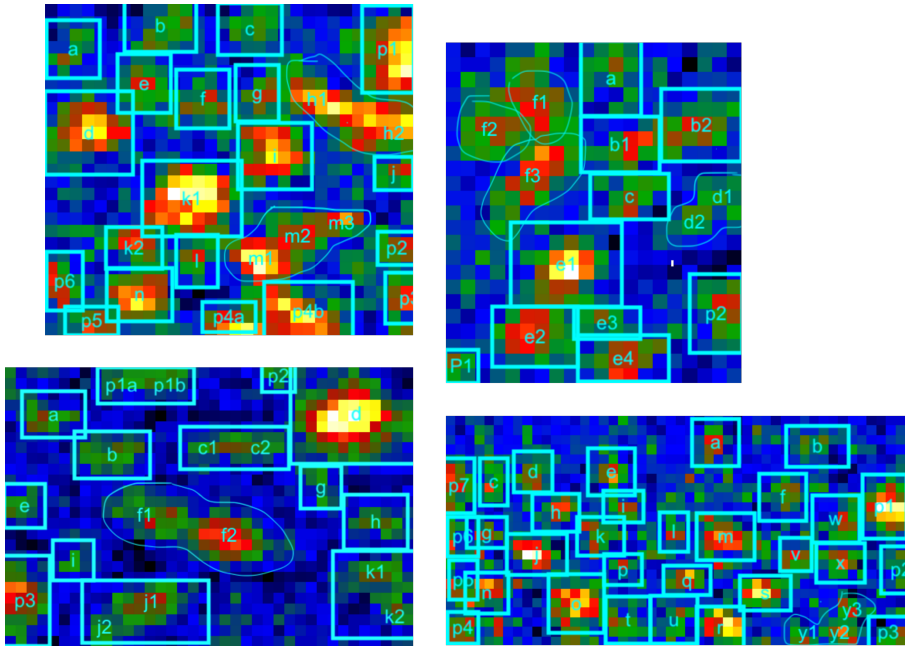

Figure 6.1: Bubble maps from the titania-doped tantala sample annealed at 600 °.

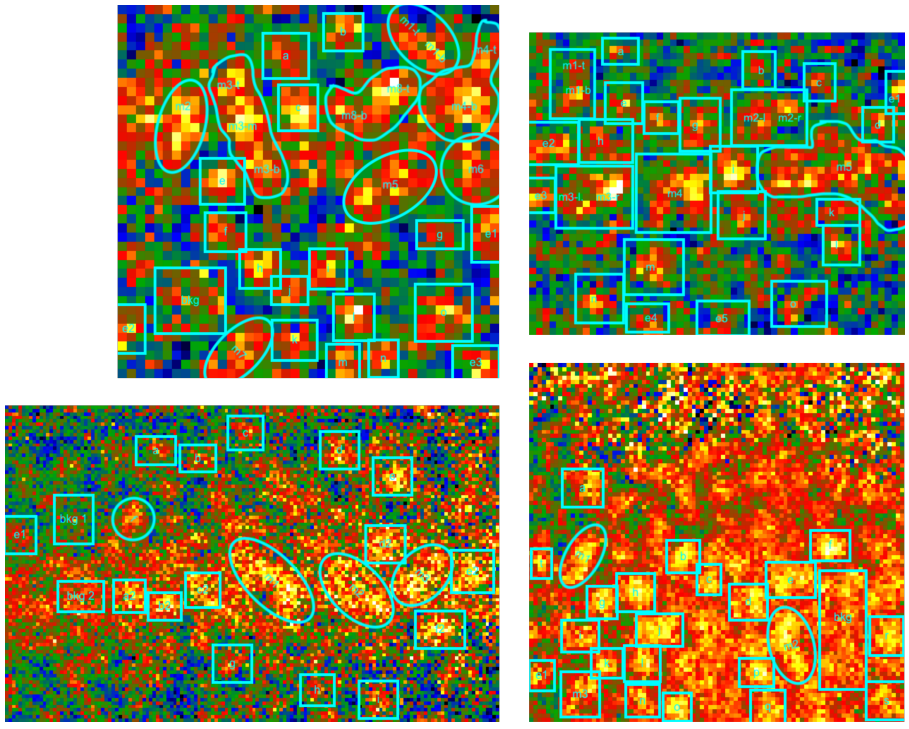

Figure 6.2: Bubble maps from the zirconia-doped tantala sample.

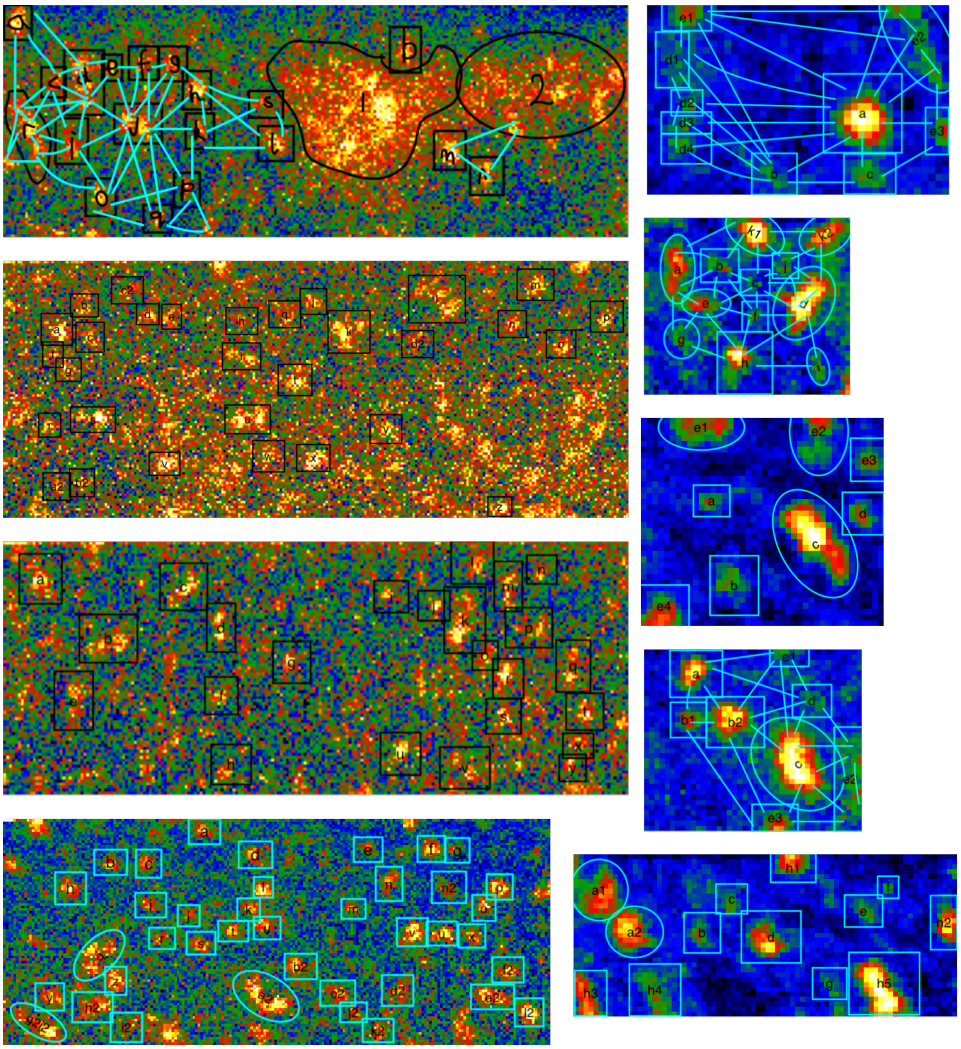

Figure 6.3: Bubble maps from the Zircaloy-4 sample implanted with xenon, with some maps showing nearest neighbour distances.

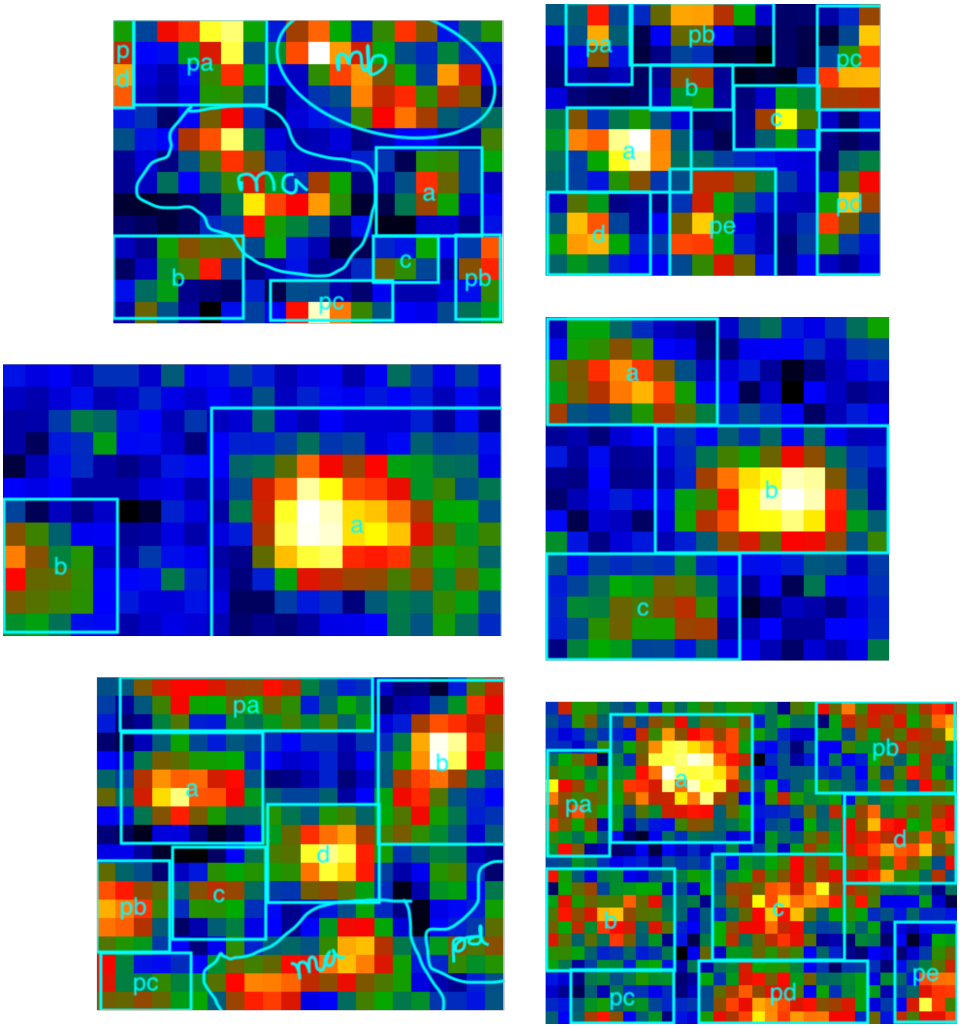

Figure 6.4: Bubble maps from the Zircaloy-4 sample implanted with krypton.

References

[1] I. MacLaren et al. “Chapter Twelve - Performing EELS at higher energy losses at both 80 and 200 kV”. In: *Advances in Imaging and Electron Physics*. Ed. by P. W. Hawkes and M. Hÿtch. Vol. 210. Advances in Imaging and Electron Physics. Elsevier, 2019, pp. 299–355. DOI: https://doi.org/10.1016/bs.aiep.2019.02.001 (cit. on pp. v, 72, 113, 130, 142, 143).

[2] J. L. Hart et al. “Direct Detection EELS at High Energy: Elemental Mapping and EXELFS”. In: *Microscopy and Microanalysis* 25.S2 (2019), pp. 584–585. DOI: 10.1017/S1431927619003659 (cit. on pp. v, 143).

[3] R. B. Cummings, R. Bassiri, I. W. Martin, and I. MacLaren. “Argon bubble formation in tantalum oxide-based films for gravitational wave interferometer mirrors”. In: *Opt. Mater. Express* 11.3 (Mar. 2021), pp. 707–718. DOI: 10.1364/OME.416038 (cit. on pp. v, 89).

[4] R. B. Cummings et al. “Xenon bubbles formed by ion implantation in zirconium alloy films”. In: *Journal of Nuclear Materials* 560 (2022), p. 153497. ISSN: 0022-3115. DOI: https://doi.org/10.1016/j.jnucmat.2021.153497 (cit. on pp. vi, 89).

[5] M. Taheri, I. MacLaren, and R.B. Cummings et al. “EELS Primer”. In: *Nature Methods and Primers* (In Preparation) (cit. on pp. vi, 54, 89).

[6] R. Abbott et al.; The LIGO and Virgo Scientific Collaboration. “Observation of Gravitational Waves from a Binary Black Hole Merger”. In: *Physical Review Letters* 116 (6 Feb. 2016), p. 061102. DOI: 10.1103/PhysRevLett.116.061102 (cit. on pp. 4, 27).

[7] A. Einstein. “Die Grundlage der allgemeinen Relativitaetstheorie”. In: *Annalen der Physik* 354.7 (1916), pp. 769–822. DOI: 10.1002/andp.19163540702. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/andp.19163540702 (cit. on p. 4).

[8] 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 (cit. on p. 4).

[9] L. M. Krauss, S. Dodelson, and S. Meyer. “Primordial Gravitational Waves and Cosmology”. In: *Science* 328.5981 (2010), pp. 989–992. ISSN: 0036-8075. DOI: 10.1126/science.1179541. eprint: https://science.sciencemag.org/content/328/5981/989.full.pdf (cit. on p. 5).

[10] R. Hulse and J. Taylor. “Discovery of a pulsar in a binary system.” In: *Astrophysical Journal* 195 (Jan. 1975), pp. L51–L53. DOI: 10.1086/181708 (cit. on p. 5).

[11] J. Taylor and J. 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 (cit. on p. 5).

[12] J. Weber. “Detection and Generation of Gravitational Waves”. In: *Phys. Rev.* 117 (1 Jan. 1960), pp. 306–313. DOI: 10.1103/PhysRev.117.306 (cit. on p. 5).

[13] J. Weber. “Gravitational Radiation”. In: *Physical Review Letters* 18 (13 Mar. 1967), pp. 498–501. DOI: 10.1103/PhysRevLett.18.498 (cit. on p. 5).

[14] J. Weber. “Gravitational-Wave-Detector Events”. In: *Physical Review Letters* 20 (23 June 1968), pp. 1307–1308. DOI: 10.1103/PhysRevLett.20.1307 (cit. on p. 5).

[15] J. Weber. "Evidence for Discovery of Gravitational Radiation". In: *Physical Review Letters* 22 (24 June 1969), pp. 1320–1324. DOI: 10.1103/PhysRevLett.22.1320 (cit. on p. 5).

[16] M. E. Gertsenshtein and V. I. Pustovoit. "On the detection of low frequency graviational waves". In: *Soviet Physics Jetp* 16.2 (1962), p. 433 (cit. on p. 5).

[17] G. E. Moss, L. R. Miller, and R. L. Forward. "Photon-Noise-Limited Laser Transducer for Gravitational Antenna". In: *Appl. Opt.* 10.11 (Nov. 1971), pp. 2495–2498. DOI: 10.1364/AO.10.002495 (cit. on p. 5).

[18] R. W. P. Drever. "General Relativity and Gravitation: Invited Papers and Discussion Reports of the 10th International Conference on General Relativity and Gravitation, Padua, July 3–8, 1983". In: Dordrecht: Springer Netherlands, 1984. Chap. Laser Interferometer Gravitational Wave Detectors, pp. 397–412. ISBN: 978-94-009-6469-3. DOI: 10.1007/978-94-009-6469-3_23 (cit. on p. 5).

[19] R. W. P. Drever et al. *Laser Interferometer Gravitational-wave Observatory (LIGO) Technical Report*. Tech. rep. LIGO, 1989 (cit. on p. 5).

[20] A. Abramovici et al. "LIGO: The Laser Interferometer Gravitational-Wave Observatory". In: *Science* 256.5055 (1992), pp. 325–333. ISSN: 0036-8075. DOI: 10.1126/science.256.5055.325. eprint: https://science.sciencemag.org/content/256/5055/325.full.pdf (cit. on p. 5).

[21] J. Hough et al. *Proposal for a Joint German-British Interferometric Gravitational Wave Detector*. Other. Document number GWD/137/JH(89). Germany: Max-Planck-Institut, 1989 (cit. on p. 5).

[22] University of Tokyo. *KAGRA gravitational-wave observatory completes construction*. 2020 (cit. on pp. 5, 15).

[23] A Brillet, A Giazotto et al. *The Virgo Project - 1989 Proposal*. Tech. rep. Virgo Scientific Collaboration, 1989 (cit. on p. 5).

[24] J. Aasi, The LIGO Scientific Collaboration, and The Virgo Collaboration. "Advanced LIGO". In: *Classical and Quantum Gravity* 32.7 (Mar. 2015), p. 074001 (cit. on p. 5).

[25] R. Abbott et al.; The LIGO and Virgo Scientific Collaboration. "An upper limit on the stochastic gravitational-wave background of cosmological origin". In: *Nature* 460.7258 (2009), pp. 990–994. DOI: 10.1038/nature08278 (cit. on pp. 5, 11).

[26] R. Abbott et al.; The LIGO and Virgo Scientific Collaboration. "Population Properties of Compact Objects from the Second LIGO-Virgo Gravitational-Wave Transient Catalog". In: *The Astrophysical Journal Letters* 913.1 (May 2021), p. L7. DOI: 10.3847/2041-8213/abe949 (cit. on p. 6).

[27] B P Abbott et al.; The LIGO and Virgo Scientific Collaboration. "LIGO: the Laser Interferometer Gravitational-Wave Observatory". In: *Reports on Progress in Physics* 72.7 (2009), p. 076901. DOI: 10.1088/0034-4885/72/7/076901 (cit. on p. 8).

[28] R. P. Kerr. "Gravitational Field of a Spinning Mass as an Example of Algebraically Special Metrics". In: *Physical Review Letters* 11 (5 Sept. 1963), pp. 237–238. DOI: 10.1103/PhysRevLett.11.237 (cit. on p. 8).

[29] A. Heger et al. "How Massive Single Stars End Their Life". In: *The Astrophysical Journal* 591.1 (July 2003), pp. 288–300. DOI: 10.1086/375341 (cit. on p. 9).

[30] M. Evans and M. Zanolin. "Handbook of Supernovae". In: Cham: Springer International Publishing, 2017. Chap. Detecting Gravitational Waves from Supernovae with Advanced LIGO, pp. 1–12. ISBN: 978-3-319-20794-0. DOI: 10.1007/978-3-319-20794-0_10-1 (cit. on p. 9).

[31] D. Maoz and C. Badenes. "The supernova rate and delay time distribution in the Magellanic Clouds". In: *Monthly Notices of the Royal Astronomical Society* 407.2 (Sept. 2010), pp. 1314–1327. ISSN: 0035-8711. DOI: {10.1111/j.1365-2966.2010.16988.x} (cit. on p. 9).

[32] W. Li et al. "Nearby supernova rates from the Lick Observatory Supernova Search – II. The observed luminosity functions and fractions of supernovae in a complete sample". In: *Monthly Notices of the Royal Astronomical Society* 412.3 (Apr. 2011), pp. 1441–1472. ISSN: 0035-8711. DOI: 10.1111/j.1365-2966.2011.18160.x. eprint: https://academic.oup.com/mnras/article-pdf/412/3/1441/17326749/mnras0412-1441.pdf (cit. on p. 9).

[33] E. Cappellaro et al. *The Rate of Supernovae. II. the Selection Effects and the Frequencies Per Unit Blue Luminosity.* 1993. eprint: astro-ph/9302017 (astro-ph) (cit. on p. 9).

[34] K. Kotake and T. Kuroda. "Handbook of Supernovae". In: Cham: Springer International Publishing, 2017. Chap. Gravitational Waves from Core-Collapse Supernovae, pp. 1671–1698. ISBN: 978-3-319-21846-5. DOI: 10.1007/978-3-319-21846-5_9 (cit. on p. 9).

[35] C. D. Ott. "The gravitational-wave signature of core-collapse supernovae". In: *Classical and Quantum Gravity* 26.6 (2009), p. 063001 (cit. on p. 9).

[36] C. Cutler. "Gravitational waves from neutron stars with large toroidal *B* fields". In: *Phys. Rev. D* 66 (8 Oct. 2002), p. 084025. DOI: 10.1103/PhysRevD.66.084025 (cit. on p. 10).

[37] K. Thorne. "Gravitational Radiation". In: *Annals of the New York Academy of Sciences* 759.1 (1995), pp. 127–152. DOI: 10.1111/j.1749-6632.1995.tb17520.x. eprint: https://nyaspubs.onlinelibrary.wiley.com/doi/pdf/10.1111/j.1749-6632.1995.tb17520.x (cit. on p. 10).

[38] C. Messenger and G. Woan. "A fast search strategy for gravitational waves from low-mass x-ray binaries". In: *Classical and Quantum Gravity* 24.19 (Sept. 2007), S469–S480. DOI: 10.1088/0264-9381/24/19/s10 (cit. on p. 10).

[39] S. Vitale. "Space-borne gravitational wave observatories". In: *General Relativity and Gravitation* 47 (5 2014), p. 1730. DOI: 10.1007/s10714-014-1730-2 (cit. on p. 11).

[40] N. Christensen. "Stochastic gravitational wave backgrounds". In: *Reports on Progress in Physics* 82.1 (Nov. 2018), p. 016903. DOI: 10.1088/1361-6633/aae6b5 (cit. on p. 11).

[41] J. D. Romano and N. J. Cornish. "Detection methods for stochastic gravitational-wave backgrounds: a unified treatment". In: *Living Reviews in Relativity* 20 (1 2017). DOI: 10.1007/s41114-017-0004-1 (cit. on p. 11).

[42] W. Chaibi et al. "Low frequency gravitational wave detection with ground-based atom interferometer arrays". In: *Phys. Rev. D* 93 (2 Jan. 2016), p. 021101. DOI: 10.1103/PhysRevD.93.021101 (cit. on p. 11).

[43] S Kawamura et al. "The Japanese space gravitational wave antenna - DECIGO". In: *Journal of Physics: Conference Series* 122 (July 2008), p. 012006. DOI: 10.1088/1742-6596/122/1/012006 (cit. on p. 11).

[44] M. Banks and S. Tesh. "Ravens at LIGO, stained-glass physics, fake space pics". In: *Phsyics World* (2017) (cit. on p. 13).

[45] K. Evans et al. "Reduced density function analysis of titanium dioxide doped tantalum pentoxide". In: *Journal of Physics: Conference Series* 37.012058 (2012) (cit. on p. 14).

[46] K.-H. Lee et al. "Improved fused silica fibres for the advanced LIGO monolithic suspensions". In: *Classical and Quantum Gravity* 36.18 (Aug. 2019), p. 185018. DOI: 10.1088/1361-6382/ab28bd (cit. on pp. 14, 19).

[47] R. Bassiri et al. “Correlations between the mechanical loss and atomic structure of amorphous TiO2-doped Ta2O5”. In: *Acta Materialia* 61 (2013), pp. 1070–1077 (cit. on pp. 14, 18).

[48] C. M. Caves. “Quantum-mechanical noise in an interferometer”. In: *Phys. Rev. D* 23 (8 Apr. 1981), pp. 1693–1708. DOI: 10.1103/PhysRevD.23.1693 (cit. on p. 14).

[49] J. Miller et al. “Prospects for doubling the range of Advanced LIGO”. In: *Phys. Rev. D* 91 (6 Mar. 2015), p. 062005. DOI: 10.1103/PhysRevD.91.062005 (cit. on pp. 14, 18).

[50] The ET Science Team. “Einstein gravitational wave telescope conceptual design study”. In: *ET technical document number ET-0106A-10* (2011). DOI: 10.1088/0264-9381/29/12/124007 (cit. on p. 15).

[51] LIGO Caltech. “LIGO’s Interferometer”. In: *LIGO Online* (2020) (cit. on p. 16).

[52] S.M. Aston et al. “Update on quadruple suspension design for Advanced LIGO”. In: *Classical and Quantum Gravity* 29.23 (Oct. 2012), p. 235004. DOI: 10.1088/0264-9381/29/23/235004 (cit. on p. 19).

[53] B. S. Lunin and K. V. Tokmakov. “Reduction in internal friction in silica glass with high OH content”. In: *Journal of the American Ceramic Society* 102.6 (2019), pp. 3329–3340. DOI: https://doi.org/10.1111/jace.16187. eprint: https://ceramics.onlinelibrary.wiley.com/doi/pdf/10.1111/jace.16187 (cit. on pp. 18, 167).

[54] R. Bassiri et al. “Order, disorder and mixing: The atomic structure of amorphous mixtures of titania and tantala”. In: *Journal of Non-Crystalline Solids* 438 (2016), pp. 59–66. ISSN: 0022-3093. DOI: https://doi.org/10.1016/j.jnoncrysol.2016.02.009 (cit. on pp. 18, 170).

[55] J. Steinlechner and I. W. Martin. "How can amorphous silicon improve current gravitational-wave detectors?" In: *Phys. Rev. D* 103 (4 2021), p. 042001. DOI: 10.1103/PhysRevD.103.042001 (cit. on p. 18).

[56] K. Craig et al. "Mirror Coating Solution for the Cryogenic Einstein Telescope". In: *Phys. Rev. Lett.* 122 (23 2019), p. 231102. DOI: 10.1103/PhysRevLett.122.231102 (cit. on p. 18).

[57] R. Birney et al. "Amorphous Silicon with Extremely Low Absorption: Beating Thermal Noise in Gravitational Astronomy". In: *Physical Review Letters* 121 (19 Nov. 2018), p. 191101. DOI: 10.1103/PhysRevLett.121.191101 (cit. on pp. 20, 167).

[58] L. Terkowski et al. "Influence of deposition parameters on the optical absorption of amorphous silicon thin films". In: *Phys. Rev. Research* 2 (3 Aug. 2020), p. 033308. DOI: 10.1103/PhysRevResearch.2.033308 (cit. on pp. 20, 167).

[59] G. Vajente et al. "Low Mechanical Loss $TiO_2 : GeO_2$ Coatings for Reduced Thermal Noise in Gravitational Wave Interferometers". In: *Phys. Rev. Lett.* 127 (7 Aug. 2021), p. 071101. DOI: 10.1103/PhysRevLett.127.071101 (cit. on p. 20).

[60] D. T. Wei and A. W. Louderback. "Method for fabricating multi-layer optical films". 4142958. 1979 (cit. on p. 21).

[61] G. Harry, T. Bodiya, and R. DeSalvo. *Optical Coatings and Thermal Noise in Precision Measurement*. Cambridge: Cambridge University Press, 2012. DOI: 10.1017/CBO9780511762314 (cit. on p. 21).

[62] H. R. Kaufman, J. J. Cuomo, and J. M. E. Harper. "Technology and applications of broad-beam ion sources used in sputtering. Part I. Ion source technology". In: *Journal of Vacuum Science and Technology* 21.3 (1982), pp. 725–736. DOI: 10.1116/1.571819. eprint: https://doi.org/10.1116/1.571819 (cit. on p. 21).

[63] J. M. E. Harper, J. J. Cuomo, and H. R. Kaufman. "Technology and applications of broad-beam ion sources used in sputtering. Part II. Applications". In: *Journal of Vacuum Science and Technology* 21.3 (1982), pp. 737–756. DOI: 10.1116/1.571820. eprint: https://doi.org/10.1116/1.571820 (cit. on p. 21).

[64] R. Brown. "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". In: *The Philisophical Magazine and Annals of Philosophy*. new series 4.21 (1828) (cit. on p. 23).

[65] H. B. Callen and T. A. Welton. "Irreversibility and Generalized Noise". In: *Phys. Rev.* 83 (1 1951), pp. 34–40. DOI: 10.1103/PhysRev.83.34 (cit. on p. 24).

[66] G. M. Harry et al. "Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings". In: *Classical and Quantum Gravity* 19.5 (Feb. 2002), pp. 897–917. DOI: 10.1088/0264-9381/19/5/305 (cit. on p. 26).

[67] A. Nowick and B. Berry. *Anelastic Relaxation in Crystalline Solids*. Materials Science Series. Academic Press, 1972. DOI: https://doi.org/10.1016/B978-0-12-522650-9.50001-X (cit. on p. 26).

[68] 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 (Apr. 2010), p. 084030. DOI: 10.1088/0264-9381/27/8/084030 (cit. on p. 26).

[69] A. Ageev et al. "Very high quality factor measured in annealed fused silica". In: *Classical and Quantum Gravity* 21.16 (July 2004), pp. 3887–3892. DOI: 10.1088/0264-9381/21/16/004 (cit. on p. 26).

[70] Y. Levin. "Internal thermal noise in the LIGO test masses: A direct approach". In: *Phys. Rev. D* 57 (2 Jan. 1998), pp. 659–663. DOI: 10.1103/PhysRevD.57.659 (cit. on p. 26).

[71] W. Jia et al. "Point Absorber Limits to Future Gravitational-Wave Detectors". In: *Phys. Rev. Lett.* 127 (24 Dec. 2021), p. 241102. DOI: 10.1103/PhysRevLett.127.241102 (cit. on p. 26).

[72] D. J. Ottaway, P. Fritschel, and S. J. Waldman. "Impact of upconverted scattered light on advanced interferometric gravitational wave detectors". In: *Opt. Express* 20.8 (Apr. 2012), pp. 8329–8336. DOI: 10.1364/OE.20.008329 (cit. on p. 26).

[73] A. Freise et al. *Gravitational wave interferometer noise calculator*. 2020. URL: https://git.ligo.org/gwinc/pygwinc (cit. on p. 27).

[74] A. Diaspro. *Expedition into the nanoworld: an exciting voyage from optical microscopy to nanoscopy*. English. Cham: Springer, 2022. ISBN: 9783030944728 (cit. on p. 29).

[75] S. Bradbury. *The Evolution of the Microscope*. Pergamon, 1967, pp. 1–35. ISBN: 978-0-08-003143-9. DOI: https://doi.org/10.1016/B978-0-08-003143-9.50006-4 (cit. on p. 30).

[76] E. Abbe. "Beiträge zur Theorie des Mikroskops und der mikroskopischen Wahrnehmung". In: *Archiv für Mikroskopische Anatomie* 9 (1873), pp. 413–468. DOI: https://doi.org/10.1007/BF02956173 (cit. on pp. 30, 36).

[77] M. Kang, Y. Ra'di, D. Farfan, and A. Alù. "Efficient Focusing with Large Numerical Aperture Using a Hybrid Metalens". In: *Phys. Rev. Applied* 13 (4 Apr. 2020), p. 044016. DOI: 10.1103/PhysRevApplied.13.044016 (cit. on p. 30).

[78] J. Louten. *Essential Human Virology. Virus Structure and Classification*. Boston: Academic Press, 2016. Chap. 2, pp. 19–29. ISBN: 978-0-12-800947-5. DOI: https://doi.org/10.1016/B978-0-12-800947-5.00002-8 (cit. on p. 30).

[79] J. W. Hsu. "Near-field scanning optical microscopy studies of electronic and photonic materials and devices". In: *Materials Science and Engineering: R: Reports* 33.1 (2001), pp. 1–50. ISSN: 0927-796X. DOI: https://doi.org/10.1016/S0927-796X(00)00031-0 (cit. on p. 30).

[80] B. Huang, M. Bates, and X. Zhuang. "Super-Resolution Fluorescence Microscopy". In: *Annual Review of Biochemistry* 78.1 (2009). PMID: 19489737, pp. 993–1016. DOI: 10.1146/annurev.biochem.77.061906.092014. eprint: https://doi.org/10.1146/annurev.biochem.77.061906.092014 (cit. on p. 30).

[81] Nobel Prize Outreach. *The Nobel Prize in Chemistry 2014; Press Release*. 2014. URL: https://www.nobelprize.org/prizes/chemistry/2014/press-release/ (cit. on p. 30).

[82] L. de Broglie. "XXXV: A tentative theory of light quanta". In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 47.278 (1924), pp. 446–458. DOI: 10.1080/14786442408634378 (cit. on p. 31).

[83] G. P. Thomson and A. Reid. "Diffraction of Cathode Rays by a Thin Film". In: *Nature* 119.3007 (1927), pp. 890–890. DOI: 10.1038/119890a0 (cit. on p. 31).

[84] C. Davisson and L. H. Germer. "Diffraction of Electrons by a Crystal of Nickel". In: *Physical Review* 30 (6 Dec. 1927), pp. 705–740. DOI: 10.1103/PhysRev.30.705 (cit. on p. 31).

[85] H. Busch. "Berechnung der Bahn von Kathodenstrahlen im axialsymmetrischen elektromagnetischen Felde". In: *Annalen der Physik* 386.25 (1926), pp. 974–993. DOI: https://doi.org/10.1002/andp.19263862507. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/andp.19263862507 (cit. on p. 31).

[86] Nobel Prize Outreach. *The Nobel Prize in Physics 1986; Press Release*. 1986. URL: https://www.nobelprize.org/prizes/physics/1986/press-release/ (cit. on p. 31).

[87] E. Ruska. "The Development of the Electron Microscope and of Electron Microscopy (Nobel Lecture)". In: *Angewandte Chemie International Edition in English* 26.7 (1987), pp. 595–605. DOI: https://doi.org/10.1002/anie.198705953. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/anie.198705953 (cit. on p. 31).

[88] Y. Aharonov and D. Bohm. "Significance of Electromagnetic Potentials in the Quantum Theory". In: *Phys. Rev.* 115 (3 Aug. 1959), pp. 485–491. DOI: 10.1103/PhysRev.115.485 (cit. on p. 35).

[89] O. Scherzer. "Das Elektron im Strahlungsfeld". In: *Annalen der Physik* 426.7 (1939), pp. 585–602. DOI: https://doi.org/10.1002/andp.19394260702. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/andp.19394260702 (cit. on p. 39).

[90] O. Scherzer. "Ein elektronenoptischer Apochromat". In: *Zeitschrift für Naturforschung A* 3.8-11 (1948), pp. 544–545. DOI: doi:10.1515/zna-1948-8-1120 (cit. on p. 39).

[91] A. Craven and T. Buggy. "Design considerations and performance of an analytical stem". In: *Ultramicroscopy* 7.1 (1981), pp. 27–37. ISSN: 0304-3991. DOI: https://doi.org/10.1016/0304-3991(81)90020-6 (cit. on p. 42).

[92] C. Ophus. "Four-Dimensional Scanning Transmission Electron Microscopy (4D-STEM): From Scanning Nanodiffraction to Ptychography and Beyond". In: *Microscopy and Microanalysis* 25.3 (2019), pp. 563–582. DOI: 10.1017/S1431927619000497 (cit. on p. 44).

[93] I. MacLaren et al. "Imaging Structure and Magnetisation in New Ways Using 4D STEM". In: *Microscopy and Microanalysis* 24.S1 (2018), pp. 180–181. DOI: 10.1017/S1431927618001393 (cit. on p. 44).

[94] S. Pennycook and D. Jesson. "High-resolution Z-contrast imaging of crystals". In: *Ultramicroscopy* 37.1 (1991), pp. 14–38. ISSN: 0304-3991. DOI: https://doi.org/10.1016/0304-3991(91)90004-P (cit. on p. 45).

[95] D. B. Williams and C. B. Carter. *Transmission Electron Mircroscopy: Diffraction, Imaging, and Spectrometry*. Switzerland: Springer, Cham, 2016. DOI: https://doi.org/10.1007/978-3-319-26651-0 (cit. on p. 45).

[96] R. F. Egerton. *Electron energy-loss spectroscopy in the electron microscope*. 3rd ed. Boston, MA: Springer, 2011 (cit. on pp. 46, 52, 61, 77).

[97] P. Roberts, J. Chapman, and A. MacLeod. "A CCD-based image recording system for the CTEM". In: *Ultramicroscopy* 8.4 (1982), pp. 385–396. ISSN: 0304-3991. DOI: https://doi.org/10.1016/0304-3991(82)90061-4 (cit. on p. 47).

[98] D. B. Williams and C. B. Carter. *Transmission Electron Mircroscopy*. 2nd ed. Springer, 2009 (cit. on pp. 49, 51).

[99] C. Kittel. *Introduction to Solid State Physics*. 5th ed. United States: John Wiley Sons Inc, 1976. Chap. 2, p. 39. ISBN: 0-471-49024-5 (cit. on p. 50).

[100] K. Iakoubovskii, K. Mitsuishi, Y. Nakayama, and K. Furuya. "Mean free path of inelastic electron scattering in elemental solids and oxides using transmission electron microscopy: Atomic number dependent oscillatory behavior". In: *Physical Review B* 77 (10 Mar. 2008), p. 104102. DOI: 10.1103/PhysRevB.77.104102 (cit. on pp. 52, 125, 135).

[101] A. J. Craven, J. Bobynko, B. Sala, and I. MacLaren. "Accurate measurement of absolute experimental inelastic mean free paths and EELS differential cross-sections". In: *Ultramicroscopy* 170 (2016), pp. 113–127.

ISSN: 0304-3991. DOI: https://doi.org/10.1016/j.ultramic.2016.08.012 (cit. on pp. 53, 59, 61, 130).

[102] A. Varambhia et al. "Determining EDS and EELS partial cross-sections from multiple calibration standards to accurately quantify bi-metallic nanoparticles using STEM". In: *Micron* 113 (2018), pp. 69–82. ISSN: 0968-4328. DOI: https://doi.org/10.1016/j.micron.2018.06.015 (cit. on p. 53).

[103] J. D. Hunter. "Matplotlib: A 2D graphics environment". In: *Computing in Science & Engineering* 9.3 (2007), pp. 90–95. DOI: 10.1109/MCSE.2007.55 (cit. on p. 55).

[104] J. Bobynko, I. MacLaren, and A. J. Craven. "Spectrum imaging of complex nanostructures using DualEELS: I. digital extraction replicas". In: *Ultramicroscopy* 149 (2015), pp. 9–20. ISSN: 0304-3991. DOI: https://doi.org/10.1016/j.ultramic.2014.10.014 (cit. on pp. 57, 61, 65, 130).

[105] J. Scott et al. "Near-simultaneous dual energy range EELS spectrum imaging". In: *Ultramicroscopy* 108.12 (Mar. 2008), pp. 1586–1594. DOI: 10.1016/j.ultramic.2008.05.006 (cit. on pp. 57, 92).

[106] A. Gubbens et al. "The GIF Quantum, a next generation post-column imaging energy filter". In: *Ultramicroscopy* 110.8 (2010). Proceedings of the international workshop on enhanced data generated by electrons, pp. 962–970. ISSN: 0304-3991. DOI: https://doi.org/10.1016/j.ultramic.2010.01.009 (cit. on p. 57).

[107] R. Egerton. "Oscillator-strength parameterization of inner-shell cross sections". In: *Ultramicroscopy* 50.1 (1993), pp. 13–28. ISSN: 0304-3991. DOI: https://doi.org/10.1016/0304-3991(93)90087-E (cit. on pp. 59, 92, 130).

[108] A. J. Craven, B. Sala, J. Bobynko, and I. MacLaren. "Spectrum imaging of complex nanostructures using DualEELS: II. Absolute quantification using standards". In: *Ultramicroscopy* 186 (2018), pp. 66–81. ISSN: 0304-3991. DOI: https://doi.org/10.1016/j.ultramic.2017.12.011 (cit. on pp. 59, 61, 90).

[109] R. D. Leapman, P. Rez, and D. F. Mayers. "K, L, and M shell generalized oscillator strengths and ionization cross sections for fast electron collisions". In: *The Journal of Chemical Physics* 72.2 (1980), pp. 1232–1243. DOI: 10.1063/1.439184. eprint: https://doi.org/10.1063/1.439184 (cit. on p. 59).

[110] J. Verbeeck and S. Van Aert. "Model-based quantification of EELS spectra". In: *Ultramicroscop* 101 (2004), pp. 207–224. DOI: doi:10.1016/j.ultramic.2004.06.0046 (cit. on p. 59).

[111] J. Verbeeck, S. Van Aert, and G. Bertoni. "Model-based quantification of EELS spectra: Including the fine structure". In: *Ultramicroscop* 106 (2006), pp. 976–980. DOI: doi:10.1016/j.ultramic.2006.05.006 (cit. on p. 59).

[112] G. Bertoni and J. Verbeeck. "Accuracy and precision in model based EELS quantification". In: *Ultramicroscop* 108 (2008), pp. 782–790. DOI: doi:10.1016/j.ultramic.2008.01.004 (cit. on p. 59).

[113] M. Watanabe and D. Williams. "The quantitative analysis of thin specimens: a review of progress from the Cliff-Lorimer to the new zeta-factor methods". In: *Journal of Microscopy* 221.2 (2006), pp. 89–109. DOI: 10.1111/j.1365-2818.2006.01549.x. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1365-2818.2006.01549.x (cit. on p. 60).

[114] K. MacArthur et al. "Quantitative energy-dispersive X-ray analysis of catalyst nanoparticles using a partial cross section approach". In:

Microscopy and Microanalysis 22.1 (Feb. 2016), pp. 71–81. DOI: 10.1017/S1431927615015494 (cit. on p. 60).

[115] R. Webster et al. "Correction of EELS dispersion non-uniformities for improved chemical shift analysis". In: *Ultramicroscopy* 217 (2020), p. 113069. ISSN: 0304-3991. DOI: https://doi.org/10.1016/j.ultramic.2020.113069 (cit. on pp. 61, 91).

[116] G. Lucas, P. Burdet, M. Cantoni, and C. Hebert. "Multivariate statistical analysis as a tool for the segmentation of 3D spectral data". In: *Micron* 52-53 (2013), pp. 49–56. ISSN: 0968-4328. DOI: https://doi.org/10.1016/j.micron.2013.08.005 (cit. on p. 65).

[117] P. Cueva et al. "Data Processing for Atomic Resolution Electron Energy Loss Spectroscopy". In: *Microscopy and Microanalysis* 18.4 (2012), pp. 667–675. DOI: 10.1017/S1431927612000244 (cit. on p. 65).

[118] D. L. Misell and A. F. Jones. "The determination of the single-scattering line profile from the observed spectrum". In: *Journal of Physics A: General Physics* 2.5 (1969), p. 540. DOI: 10.1088/0305-4470/2/5/007 (cit. on p. 67).

[119] D. W. Johnson and J. C. H. Spence. "Determination of the single-scattering probability distribution from plural-scattering data". In: *Journal of Physics D: Applied Physics* 7.6 (1974), p. 771. DOI: 10.1088/0022-3727/7/6/304 (cit. on p. 67).

[120] E. R. F., W. B. G., and S. T. G. "Fourier deconvolution of electron energy-loss spectra". In: *Proceedings of the Royal Society* 398 (1985), pp. 395–404. DOI: http://doi.org/10.1098/rspa.1985.0041 (cit. on p. 67).

[121] Gatan. *Argon*. 2020. URL: https://eels.info/atlas/argon (cit. on pp. 69, 90).

[122] R. D. Leapman and V. E. Cosslett. "Extended fine structure above the X-ray edge in electron energy loss spectra". In: *Journal of Physics D: Applied Physics* 9.2 (Feb. 1976), pp. L29–L32. DOI: 10.1088/0022-3727/9/2/003 (cit. on pp. 72, 141).

[123] L. M. Brown. "A Synchrotron in a Microscope". In: *Institute of Physics Conference. No 153: Section 1* (Cambridge, 1997). IOP Publishing Ltd., 1997, pp. 17–22 (cit. on pp. 72, 168).

[124] G. Bunker. *Introduction to XAFS: A Practical Guide to X-ray Absorption Fine Structure Spectroscopy*. Cambridge University Press, 2010. DOI: 10.1017/CBO9780511809194 (cit. on p. 75).

[125] J. J. Rehr, J. J. Kas, F. D. Vila, and M. Newville. "Theory and Analysis of XAFS". In: *XAFS Techniques for Catalysts, Nanomaterials, and Surfaces*. Ed. by Y. Iwasawa, K. Asakura, and M. Tada. Cham, Switzerland: Springer International Publishing, 2017, pp. 13–50. ISBN: 978-3-319-43866-5. DOI: 10.1007/978-3-319-43866-5_2 (cit. on p. 76).

[126] S. Calvin. *XAFS for Everyone*. CRC Press, 2013 (cit. on p. 76).

[127] T. Mizoguchi, W. Olovsson, H. Ikeno, and I. Tanaka. "Theoretical ELNES using one-particle and multi-particle calculations". In: *Micron* 41.7 (2010), pp. 695–709. ISSN: 0968-4328. DOI: https://doi.org/10.1016/j.micron.2010.05.011 (cit. on p. 77).

[128] B. Ravel and M. Newville. "*ATHENA*, *ARTEMIS*, *HEPHAESTUS*: data analysis for X-ray absorption spectroscopy using *IFEFFIT*". In: *Journal of Synchrotron Radiation* 12.4 (July 2005), pp. 537–541. DOI: 10.1107/S0909049505012719 (cit. on p. 78).

[129] M. Newville. "IFEFFIT : interactive XAFS analysis and FEFF fitting". In: *Journal of Synchrotron Radiation* 8 (2001), pp. 322–324 (cit. on p. 78).

[130] J. J. Rehr. "Full-Spectrum Calculations of EELS, XAS and Optical Constants". In: *Microscopy and Microanalysis* 10.S02 (2004), pp. 850–851. DOI: 10.1017/S1431927604887105 (cit. on p. 78).

[131] J. J. Rehr and R. C. Albers. "Theoretical approaches to x-ray absorption fine structure". In: *Rev. Mod. Phys.* 72 (3 July 2000), pp. 621–654. DOI: 10.1103/RevModPhys.72.621 (cit. on p. 84).

[132] M. Prato, A. Chincarini, G. Gemme, and M. Canepa. "Gravitational waves detector mirrors: Spectroscopic ellipsometry study of Ta2O5 films on SiO2 substrates". In: *Thin Solid Films* 519.9 (2011). 5th International Conference on Spectroscopic Ellipsometry (ICSE-V), pp. 2877–2880. ISSN: 0040-6090. DOI: https://doi.org/10.1016/j.tsf.2010.12.086 (cit. on pp. 87, 122).

[133] L. Anghinolfi et al. "Optical properties of uniform, porous, amorphous Ta2O5 coatings on silica: temperature effects". In: *Journal of Physics D: Applied Physics* 46.45 (Oct. 2013), p. 455301. DOI: 10.1088/0022-3727/46/45/455301 (cit. on pp. 87, 137).

[134] K. Prasai et al. "High Precision Detection of Change in Intermediate Range Order of Amorphous Zirconia-Doped Tantala Thin Films Due to Annealing". In: *Physical Review Letters* 123 (4 July 2019), p. 045501. DOI: 10.1103/PhysRevLett.123.045501 (cit. on pp. 88, 125, 136, 137, 171).

[135] C. Harthcock et al. "The impact of nano-bubbles on the laser performance of hafnia films deposited by oxygen assisted ion beam sputtering method". In: *Applied Physics Letters* 115.25 (2019), p. 251902. DOI: 10.1063/1.5129454 (cit. on pp. 88, 135, 136).

[136] M. A. Fazio et al. "Structure and morphology of low mechanical loss TiO2-doped Ta2O5". In: *Opt. Mater. Express* 10.7 (July 2020), pp. 1687–1703. DOI: 10.1364/OME.395503 (cit. on p. 88).

[137] S. Donnelly and C. Rossouw. "Lattice images of inert gas bubbles in aluminium". In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 13.1 (1986), pp. 485–489. ISSN: 0168-583X. DOI: https://doi.org/10.1016/0168-583X(86)90552-5 (cit. on pp. 88, 135).

[138] R. Manzke et al. "Quantitative Determination of the Pressure of He in Bubbles in Al and Ni". In: *Solld State Communlcations* 44.4 (1982), pp. 481–484 (cit. on p. 88).

[139] N. Marochov and P. Goodhew. "A comparison of the growth of helium and neon bubbles in nickel". In: *Journal of Nuclear Materials* 158 (1988), pp. 81–86. ISSN: 0022-3115. DOI: https://doi.org/10.1016/0022-3115(88)90157-2 (cit. on p. 88).

[140] J. Rest. "Validation of mechanistic models for gas precipitation in solids during postirradiation annealing experiments". In: *Journal of Nuclear Materials* 168.3 (1989), pp. 243–261. ISSN: 0022-3115. DOI: https://doi.org/10.1016/0022-3115(89)90590-4 (cit. on pp. 88, 135).

[141] M. Song, K. Mitsuishi, and K. Furuya. "Morphologies of metastable inert gas precipitates in aluminum observed with in situ HRTEM". In: *Materials Science and Engineering: A* 304-306 (2001). RQ10, Tenth International Conference on Rapidly Quenched and Metastable Materials, pp. 135–143. ISSN: 0921-5093. DOI: https://doi.org/10.1016/S0921-5093(00)01463-5 (cit. on pp. 88, 135).

[142] H. Pattyn, P. Hendrickx, J. Odeurs, and S. Bukshpan. "Pressure Effects of Krypton Bubbles on the Host Aluminium Lattice as Probed by S_7Co Impurity Atoms". In: *Materials Science and Engineering, A* (1988) (cit. on pp. 88, 135).

[143] W. Jaeger et al. "The density and pressure of helium in bubbles in metals". In: *Radiation Effects* 78.1-4 (1983), pp. 315–325. DOI: 10.1080/00337578308207380 (cit. on p. 88).

[144] K. Stevens and P. Johnson. "Defocussed TEM contrast from small helium bubbles in copper". In: *Journal of Nuclear Materials* 246.1 (1997), pp. 17–29. ISSN: 0022-3115. DOI: https://doi.org/10.1016/S0022-3115(97)00089-5 (cit. on p. 88).

[145] J. H. Evans and D. J. Mazey. "Evidence for solid krypton bubbles in copper, nickel and gold at 293K". In: *Journal of Physics F: Metal Physics* 15.1 (Jan. 1985), pp. L1–L6. DOI: 10.1088/0305-4608/15/1/001 (cit. on pp. 88, 135).

[146] S. E. Donnelly, R. C. Birtcher, C. Templier, and V. Vishnyakov. "Response of helium bubbles in gold to displacement-cascade damage". In: *Phys. Rev. B* 52 (6 Aug. 1995), pp. 3970–3975. DOI: 10.1103/PhysRevB.52.3970 (cit. on p. 88).

[147] K. Alix et al. "Gentle quantitative measurement of helium density in nanobubbles in silicon by spectrum imaging". In: *Micron* 77 (2015), pp. 57–65. ISSN: 0968-4328. DOI: https://doi.org/10.1016/j.micron.2015.05.011 (cit. on p. 88).

[148] E. Oliviero et al. "Synthesis of mesoporous amorphous silica by Kr and Xe ion implantation: Transmission electron microscopy study of induced nanostructures". In: *Microporous and Mesoporous Materials* 132.1 (2010), pp. 163–173. ISSN: 1387-1811. DOI: https://doi.org/10.1016/j.micromeso.2010.02.015 (cit. on p. 88).

[149] F. Rufeh and D.R. Olander and T.H. Pigford. "The Solubility of Helium in Uranium Dioxide". In: *Nuclear Science and Engineering* 23.4 (1965), pp. 335–338. DOI: 10.13182/NSE65-A21069 (cit. on p. 88).

[150] D. Olander and D. Wongsawaeng. "Re-solution of fission gas – A review: Part I. Intragranular bubbles". In: *Journal of Nuclear Materials* 354.1 (2006), pp. 94–109. ISSN: 0022-3115. DOI: https://doi.org/10.1016/j.jnucmat.2006.03.010 (cit. on p. 88).

[151] L. Pagano, A. Motta, and R. Birtcher. "Bubble Formation in Zr Alloys Under Heavy Ion Implantation". In: *MRS Online Proceedings Library volume* 398 (1996) (cit. on p. 88).

[152] Y. Ishiyama et al. "Post-irradiation annealing effects on microstructure and helium bubbles in neutron irradiated type 304 stainless steel". In: *Journal of Nuclear Materials* 239 (1996). Proceedings of the JIM '95 Fall Annual Meeting (117th) on Lattice Defects and Radiation Induced Phenomena, pp. 90–94. ISSN: 0022-3115. DOI: https://doi.org/10.1016/S0022-3115(96)00465-5 (cit. on pp. 88, 135).

[153] A. Jelea. "Molecular dynamics modeling of helium bubbles in austenitic steels". In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 425 (2018), pp. 50–54. ISSN: 0168-583X. DOI: https://doi.org/10.1016/j.nimb.2018.04.008 (cit. on pp. 88, 129).

[154] W. Miekeley and F. Felix. "Effect of stoichiometry on diffusion of xenon in UO_2". In: *Journal of Nuclear Materials* 42.3 (1972), pp. 297–306. ISSN: 0022-3115. DOI: https://doi.org/10.1016/0022-3115(72)90080-3 (cit. on p. 88).

[155] S. Guilbert et al. "Helium behavior in UO_2 polycrystalline disks". In: *Journal of Nuclear Materials* 321.2 (2003), pp. 121–128. ISSN: 0022-3115. DOI: https://doi.org/10.1016/S0022-3115(03)00212-5 (cit. on p. 88).

[156] L. Pagano, A. Motta, and R. Birtcher. "The formation of bubbles in Zr alloys under Kr ion irradiation". In: *Journal of Nuclear Materials* 244.3 (1997). Radiation Materials Science in Technology Applications,

pp. 295–304. ISSN: 0022-3115. DOI: https://doi.org/10.1016/S0022-3115(96)00757-X (cit. on p. 88).

[157] R. Gilbert, K. Farrell, and C. Coleman. "Damage structure in zirconium alloys neutron irradiated at 573 to 923 K". In: *Journal of Nuclear Materials* 84.1 (1979), pp. 137–148. ISSN: 0022-3115. DOI: https://doi.org/10.1016/0022-3115(79)90157-0 (cit. on p. 88).

[158] R. Barnes and D. Mazey. "The nature of radiation-induced point defect clusters". In: *The Philosophical Magazine: A Journal of Theoretical Experimental and Applied Physics* 5.60 (1960), pp. 1247–1253. DOI: 10.1080/14786436008238337. eprint: https://doi.org/10.1080/14786436008238337 (cit. on pp. 88, 124).

[159] R. Barnes and D. Mazey. "The migration and coalescence of inert gas bubbles in metals". In: *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 275.1360 (1963), pp. 47–57. DOI: 10.1098/rspa.1963.0154 (cit. on pp. 88, 124).

[160] G. Greenwood, A. Foreman, and D. Rimmer. "The role of vacancies and dislocations in the nucleation and growth of gas bubbles in irradiated fissile material". In: *Journal of Nuclear Materials* 1.4 (1959), pp. 305–324. ISSN: 0022-3115. DOI: https://doi.org/10.1016/0022-3115(59)90030-3 (cit. on pp. 88, 124).

[161] G. Ran et al. "In situ TEM observation of growth behavior of Kr bubbles in zirconium alloy during post-implantation annealing". In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 307 (2013). The 18th International Conference on Ion Beam Modifications of Materials (IBMM2012), pp. 516–521. ISSN: 0168-583X. DOI: https://doi.org/10.1016/j.nimb.2012.11.085 (cit. on pp. 88, 135).

[162] N. Miyaji, Y. Abe, S. Ukai, and S. Onose. "Post-irradiation creep rupture properties of FBR grade 316 SS structural material". In: *Journal of*

Nuclear Materials 271-272 (1999), pp. 173–178. ISSN: 0022-3115. DOI: https://doi.org/10.1016/S0022-3115(98)00776-4 (cit. on pp. 88, 135).

[163] L. He et al. "Bubble evolution in Kr-irradiated UO_2 during annealing". In: *Journal of Nuclear Materials* 496 (2017), pp. 242–250. ISSN: 0022-3115. DOI: https://doi.org/10.1016/j.jnucmat.2017.09.036 (cit. on p. 89).

[164] A. V. Cumming et al. "Large-scale Monolithic Fused-Silica Mirror Suspension for Third-Generation Gravitational-Wave Detectors". In: *Phys. Rev. Applied* 17 (2 Feb. 2022), p. 024044. DOI: 10.1103/PhysRevApplied.17.024044 (cit. on p. 153).

[165] H. N. Isa. "Towards reliable quantification of mixed oxide coating for LIGO using DualEELS". PhD thesis. University of Glasgow, 2019 (cit. on pp. 95, 97, 100, 107, 125, 127).

[166] M. J. Hart et al. "Medium range structural order in amorphous tantala spatially resolved with changes to atomic structure by thermal annealing". In: *Journal of Non-Crystalline Solids* 438 (2016), pp. 10–17. ISSN: 0022-3093. DOI: https://doi.org/10.1016/j.jnoncrysol.2016.02.005 (cit. on p. 96).

[167] G. M. Harry et al. "Titania-doped tantala/silica coatings for gravitational-wave detection". In: *Classical and Quantum Gravity* 24.2 (2006), pp. 405–415. DOI: 10.1088/0264-9381/24/2/008 (cit. on p. 97).

[168] M. Blackmur et al. "The association of hydrogen with nanometre bubbles of helium implanted into zirconium". In: *Scripta Materialia* 152 (2018), pp. 102–106. ISSN: 1359-6462. DOI: https://doi.org/10.1016/j.scriptamat.2018.04.007 (cit. on p. 112).

[169] R. C. Birtcher et al. "Behavior of Crystalline Xe Nanoprecipitates during Coalescence". In: *Physical Review Letters* 83 (8 Aug. 1999),

pp. 1617–1620. DOI: 10.1103/PhysRevLett.83.1617 (cit. on pp. 119, 135).

[170] R. E. Van Wylen Gordon J.; Sonntag. *Fundamentals of classical thermodynamics*. 3rd. John Wiley Sons, 1985, p. 46. ISBN: 0471829331 (cit. on p. 123).

[171] O. Redlich and J. N. S. Kwong. "On the Thermodynamics of Solutions. V. An Equation of State. Fugacities of Gaseous Solutions." In: *Chemical Reviews* 44.1 (1949). PMID: 18125401, pp. 233–244. DOI: 10.1021/cr60137a013. eprint: https://doi.org/10.1021/cr60137a013 (cit. on p. 123).

[172] R. C. Weast. *CRC handbook of chemistry and physics. A ready-reference book of chemical and physical data*. 53rd ed. Boca Raton, Florida: CRC Press, 1972. ISBN: 0878194541 (cit. on p. 123).

[173] A. Jelea. "An equation of state for xenon/krypton mixtures confined in the nuclear fuels". In: *Journal of Nuclear Materials* 530 (2020), p. 151952. ISSN: 0022-3115. DOI: https://doi.org/10.1016/j.jnucmat.2019.151952 (cit. on pp. 124, 138).

[174] W. Ostwald. "1. Abhandlung: Übersättigung und Überkaltung". In: *Zeitschrift für Physikalische Chemie* 22U.1 (1897), pp. 289–330. DOI: doi:10.1515/zpch-1897-2233 (cit. on p. 130).

[175] L. Ratke and P. W. Voorhees. *Growth and Coarsening: Ostwald Ripening in Material Processing*. Heidelberg: Springer Berlin, 2013. DOI: https://doi.org/10.1007/978-3-662-04884-9 (cit. on p. 130).

[176] Allegheny Technologies Incorporated. "Reactor Grade Zirconium: Technical Data Sheet". In: *1000 Six PPG Place, Pittsburgh, PA 15222-5479 U.S.A.* (2015) (cit. on p. 130).

[177] A. J. Craven, H. Sawada, S. McFadzean, and I. MacLaren. "Getting the most out of a post-column EELS spectrometer on a TEM/STEM by optimising the optical coupling". In: *Ultramicroscopy* 180 (2017).

Ondrej Krivanek: A research life in EELS and aberration corrected STEM, pp. 66–80. ISSN: 0304-3991. DOI: https://doi.org/10.1016/j.ultramic.2017.03.017 (cit. on p. 131).

[178] A. vom Felde et al. "Pressure of Neon, Argon, and Xenon Bubbles in Aluminum". In: *Physical Review Letters* 53.9 (1984) (cit. on p. 135).

[179] C. J. Rossouw and S. E. Donnelly. "Superheating of Small Solid-Argon Bubbles in Aluminum". In: *Physical Review Letters* 55.27 (1985), pp. 2960–2963 (cit. on p. 135).

[180] R. Birtcher and A. Liu. "Temperature dependence of Kr precipitation in Ni". In: *Journal of Nuclear Materials* 165.2 (1989), pp. 101–109. ISSN: 0022-3115. DOI: https://doi.org/10.1016/0022-3115(89)90238-9 (cit. on p. 135).

[181] M. Song et al. "Structure variation of nanometer-sized Xe particles embedded in Al crystals". In: *Applied Surface Science* 241.1 (2005). The 9th International Symposium on Advanced Physical Fields, pp. 96–101. ISSN: 0169-4332. DOI: https://doi.org/10.1016/j.apsusc.2004.09.024 (cit. on p. 135).

[182] K. Furuya et al. "Interfaces of Xe Inclusions Ion-implanted in Al - Ordering in a Fluid Xe and Matrix Oxidation at Surfaces". In: *Microscopy and Microanalysis* 15.S2 (2009), pp. 1360–1361. DOI: 10.1017/S1431927609099577 (cit. on pp. 135, 138).

[183] D. Potter and C. Rossouw. "High Resolution of crystalline xenon and xenon alloys in Aluminium". In: *Journal of Nuclear Materials* 161 (1989) (cit. on p. 135).

[184] S. E. Donnelly and C. J. Rossouw. "Lattice Images of Solid Xenon Precipitates in Aluminum at Room Temperature". In: *Science* 230.4731 (1985), pp. 1272–1273. DOI: 10.1126/science.230.4731.1272. eprint: https://www.science.org/doi/pdf/10.1126/science.230.4731.1272 (cit. on p. 135).

[185] S. Kashibe, K. Une, and K. Nogita. "Formation and growth of intragranular fission gas bubbles in UO_2 fuels with burnup of 6–83 GWd/t". In: *Journal of Nuclear Materials* 206.1 (1993), pp. 22–34. ISSN: 0022-3115. DOI: https://doi.org/10.1016/0022-3115(93)90229-R (cit. on pp. 130, 135).

[186] T. Hojo et al. "Radiation effects on yttria-stabilized zirconia irradiated with He or Xe ions at high temperature". In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 241.1 (2005). The Application of Accelerators in Research and Industry, pp. 536–542. ISSN: 0168-583X. DOI: https://doi.org/10.1016/j.nimb.2005.07.102 (cit. on p. 135).

[187] R. C. Birtcher and W. Jäger. "Precipitation and Phase stability of solid Kr in cavities after room-temperature implantation of Al". In: *Ultramicroscopy* 22 (1987), pp. 267–279 (cit. on p. 135).

[188] A. H. Mir et al. "Xenon solubility and formation of supercritical xenon precipitates in glasses under non-equilibrium conditions". In: *Scientific Reports* 8.1 (2018), pp. 2045–2322. DOI: 10.1038/s41598-018-33556-y (cit. on p. 135).

[189] A. H. Mir, J. A. Hinks, and S. E. Donnelly. "Effect of density and Z-contrast on the visibility of noble gas precipitates and voids with insights from Monte-Carlo simulations". In: *Micron* 126 (2019), p. 102712. ISSN: 0968-4328. DOI: https://doi.org/10.1016/j.micron.2019.102712 (cit. on p. 135).

[190] A. Amato et al. "Optical and mechanical properties of ion-beam-sputtered Nb_2O_5 and $TiO_2-Nb_2O_5$ thin films for gravitational-wave interferometers and an improved measurement of coating thermal noise in Advanced LIGO". In: *Phys. Rev. D* 103 (7 Apr. 2021), p. 072001. DOI: 10.1103/PhysRevD.103.072001 (cit. on p. 135).

[191] A. Paolone et al. "Argon and Other Defects in Amorphous SiO2 Coatings for Gravitational-Wave Detectors". In: *Coatings* 12.7 (2022). ISSN: 2079-6412. DOI: 10.3390/coatings12071001 (cit. on p. 135).

[192] 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 (Feb. 2018), p. 075001. DOI: 10.1088/1361-6382/aaad7c (cit. on p. 136).

[193] M. Granata et al. "Mechanical loss in state-of-the-art amorphous optical coatings". In: *Phys. Rev. D* 93 (1 Jan. 2016), p. 012007. DOI: 10.1103/PhysRevD.93.012007 (cit. on pp. 136, 171).

[194] P. B. Mirkarimi et al. "Improving the laser performance of ion beam sputtered dielectric thin films through the suppression of nanoscale defects by employing a xenon sputtering gas". In: *Opt. Mater. Express* 12.9 (Sept. 2022), pp. 3365–3378. DOI: 10.1364/OME.462407 (cit. on pp. 136, 139).

[195] Z.-G. Mei, L. Liang, and A. M. Yacout. "First-principles study of fission gas incorporation and migration in zirconium nitride". In: *Computational Materials Science* 133 (2017), pp. 175–184. ISSN: 0927-0256. DOI: https://doi.org/10.1016/j.commatsci.2017.03.019 (cit. on p. 138).

[196] C. Cagna et al. "A complementary approach to estimate the internal pressure of fission gas bubbles by SEM-SIMS-EPMA in irradiated nuclear fuels". In: *IOP Conference Series: Materials Science and Engineering* 109 (Feb. 2016), p. 012002. DOI: 10.1088/1757-899x/109/1/012002 (cit. on p. 138).

[197] J. Martin and J. Mansot. "EXELFS analysis of amorphous and crystalline silicon carbide". In: *Journal of Microscopy* 162.1 (1991), pp. 171–178 (cit. on p. 141).

[198] G. Z. Virginie Serin and J. Sévely. "EXELFS as a structural tool for studies of low Z-elements". In: *Microscopy, Microanalysis and Microstructure* (1992), p. 201 (cit. on p. 141).

[199] D. Haskel, M. Qian, E. Stern, and M. Sarikaya. "Development of the EXELFS technique for nanoscale atomic structure determination". In: *Journal de Physique IV* (1996) (cit. on p. 142).

[200] M. Sarikaya, M. Qian, and E. Stern. "EXELFS Revisited". In: *Micron* 27.6 (1997), pp. 449–466 (cit. on p. 142).

[201] F. M. Alamgir, H. Jain, D. B. Williams, and R. B. Schwarz. "The structure of a metallic glass system using EXELFS and EXAFS as complementary probes". In: *Micron* 34.8 (2003), pp. 433–439. ISSN: 0968-4328. DOI: https://doi.org/10.1016/S0968-4328(03)00100-8 (cit. on p. 142).

[202] F. Márquez. "EXELFS characterization of Mg vanadate". In: *Surface and Interface Analysis* 36.8 (2004), pp. 749–751. DOI: https://doi.org/10.1002/sia.1754. eprint: https://analyticalsciencejournals.onlinelibrary.wiley.com/doi/pdf/10.1002/sia.1754 (cit. on p. 142).

[203] P. E. Batson. "Local crystal anisotropy obtained in the small probe geometry". In: *Micron* 39.6 (2008). Proceedings of the Annual Meeting of the Microscopical Society of Canada 2007, pp. 648–652. ISSN: 0968-4328. DOI: https://doi.org/10.1016/j.micron.2007.10.009 (cit. on p. 142).

[204] L. Camilli et al. "Probing the structure of Fe nanoparticles in multiwall carbon nanotubes grown on a stainless steel substrate". In: *Journal of Nanoparticle Research* 15.8 (2013), p. 1846. DOI: https://10.1007/s11051-013-1846-4 (cit. on p. 142).

[205] P. Nandi and J. M. Howe. "Determining the Volume Expansion at Grain Boundaries Using Extended Energy-Loss Fine Structure Analysis". In: *Microscopy and Microanalysis* 25.5 (2019), pp. 1130–1138. DOI: 10.1017/S1431927619014740 (cit. on p. 142).

[206] I. MacLaren, K. Annand, C. Black, and A. Craven. "EELS at very high energy losses". In: *Microscopy* 67.suppl$_1$ (Sept. 2017), pp. 78–85. ISSN: 2050-5698. DOI: 10.1093/jmicro/dfx036. eprint: https://academic.oup.com/jmicro/article-pdf/67/suppl_1/i78/26055171/dfx036.pdf (cit. on p. 142).

[207] C. Hébert-Souche et al. "The orientation-dependent simulation of ELNES". In: *Ultramicroscopy* 83.1 (2000), pp. 9–16. ISSN: 0304-3991. DOI: https://doi.org/10.1016/S0304-3991(99)00168-0 (cit. on p. 147).

[208] P. A. van Aken, Z. Y. Wu, F. Langenhorst, and F. Seifert. "ELNES spectroscopy and XANES calculations of the O K edge: Orientation dependence and effects of protons in $Mg(OH)_2$". In: *Phys. Rev. B* 60 (6 Aug. 1999), pp. 3815–3820. DOI: 10.1103/PhysRevB.60.3815 (cit. on p. 147).

[209] J. C. H. Spence and J. Taftø. "ALCHEMI: a new technique for locating atoms in small crystals". In: *Journal of Microscopy* 130.2 (1983), pp. 147–154. DOI: https://doi.org/10.1111/j.1365-2818.1983.tb04213.x (cit. on p. 147).

[210] J. L. Hart. "Direct Detection Electron Energy Loss Spectroscopy: Characterization and Applications". Doctor of Philosophy. Drexel University, 2019 (cit. on pp. 150, 160).

[211] CrystalMaker Software Ltd. *Generated using CrystalDiffract®: a powder diffraction program for Mac and Windows*. England: Oxford (cit. on p. 150).

[212] K. Prasai, R. Bassiri, H.-P. Cheng, and M. M. Fejer. "Annealing-Induced Changes in the Atomic Structure of Amorphous Silica, Germania, and Tantala Using Accelerated Molecular Dynamics". In: *physica status solidi (b)* 258.9 (2021), p. 2000519. DOI: https://doi.org/10.1002/pssb.202000519. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/pssb.202000519 (cit. on pp. 153, 167).

[213] R. Thielsch, A. Gatto, and N. Kaiser. "Mechanical stress and thermal-elastic properties of oxide coatings for use in the deep-ultraviolet spectral region". In: *Appl. Opt.* 41.16 (June 2002), pp. 3211–3217. DOI: 10.1364/AO.41.003211 (cit. on pp. 153, 167).

[214] V. Jambur et al. "Temperature effects on the structure and mechanical properties of vapor deposited a-SiO2". In: *Journal of Non-Crystalline Solids* 587 (2022), p. 121588. ISSN: 0022-3093. DOI: https://doi.org/10.1016/j.jnoncrysol.2022.121588 (cit. on p. 153).

[215] N. Kim, R. Bassiri, M. M. Fejer, and J. F. Stebbins. "Structure of amorphous silica–hafnia and silica–zirconia thin-film materials: The role of a metastable equilibrium state in non-glass-forming oxide systems". In: *Journal of Non-Crystalline Solids* 429 (2015), pp. 5–12. ISSN: 0022-3093. DOI: https://doi.org/10.1016/j.jnoncrysol.2015.08.030 (cit. on pp. 153, 167).

[216] A. Jeanroy, A. Bouvet, and G. Remillieux. "HRG and marine applications". In: *Gyroscopy and Navigation* 5.2 (2014), pp. 67–74. DOI: 10.1134/S2075108714020047 (cit. on p. 153).

[217] F. Riehle. "Optical clock networks". In: *Nature Photonics* 11.1 (2014), pp. 25–31. DOI: 10.1038/nphoton.2016.235 (cit. on p. 153).

[218] S. Takashima, Z. Li, and T. P. Chow. "Metal–Oxide–Semiconductor Interface and Dielectric Properties of Atomic Layer Deposited SiOsub2/subon GaN". In: *Japanese Journal of Applied Physics* 52.8S (Aug. 2013), 08JN24. DOI: 10.7567/jjap.52.08jn24 (cit. on p. 153).

[219] H. K. Raut et al. "Porous $SiO_2 anti - reflective coatings on large - area substrates by electrospinning and their application to solar modules$". In: *Solar Energy Materials and Solar Cells* 111 (2013), pp. 9–15. ISSN: 0927-0248. DOI: https://doi.org/10.1016/j.solmat.2012.12.023 (cit. on p. 153).

[220] C. R. Billman et al. "Origin of the second peak in the mechanical loss function of amorphous silica". In: *Phys. Rev. B* 95 (1 Jan. 2017), p. 014109. DOI: 10.1103/PhysRevB.95.014109 (cit. on p. 167).

[221] D. Cockayne. "The Study of Nanovolumes of Amorphous Materials Using Electron Scattering". In: *The Annual Review of Materials Research* 37 (2007), pp. 159–187. DOI: 10.1146/annurev.matsci.35.082803.103337 (cit. on p. 168).

I have hated the words and I have loved them, and I hope I have made them right.

— Markus Zusak, *The Book Thief*.