Design, build and characterisation of a prototype single crystalline silicon cryogenic suspension for 3rd generation gravitational wave detectors

Graeme Arthur Robert Eddolls BSc (Hons)

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

Institute for Gravitational Research
School of Physics and Astronomy
College of Science and Engineering
University of Glasgow

November 2022

Summary

Over 100 years ago Albert Einstein predicted the existence of gravitational waves through his General Theory of Relativity. These waves propagate through spacetime resulting from collisions of extremely dense objects, such as neutron stars and black holes, that are travelling at significant fractions of the speed of light just prior to their final merger. At the time of publishing this theory, the direct detection of these gravitational waves was thought impossible since even the most extreme of these events were expected to cause a strain amplitude of 10^{-21} by the time they reached Earth.

After 100 years of research and development, from using resonant bar detectors, to modern day laser interferometry, the direct detection of gravitational waves was made by the aLIGO detectors in September 2015. This detection was made possible due to advanced laser interferometry techniques combined with low mechanical loss materials and careful experimental design. aLIGO detections were enabled by the final stage fused silica suspensions that hang the test mass mirrors, reflecting the interferometer beam. Up to 90 such gravitational wave events have been detected so far, enabled by the upgrades that are allowing the detectors to approach design sensitivity, with the Virgo collaboration joining the hunt in 2017 and working alongside aLIGO, to give three operational 2^{nd} generation gravitational wave detectors around the world.

In order to significantly increase the sensitivity of current gravitational wave detectors, a step change is required to realise their next generation, allowing for a sensitivity increase of up to one order of magnitude, and beyond. A specific change that is being considered for next generation detectors will be the cryogenic operation of the final-stage suspension. Operating the detector at lower temperatures will ultimately reduce a dominant low frequency source of noise in the detector, thermal noise. A fourth detector based in Japan, called KAGRA, has been the first full scale demonstration of the beginnings of this next generation of detector. KAGRA differs from the aLIGO and Virgo detectors by being built underground, to reduce seismic noise, as well as making the final stage suspension from sapphire which has excellent material properties at low temperatures, as compared to fused silica. Further, next generation detectors like the Einstein Telescope and Cosmic Explorer are in development. The low frequency, cryogenic component of these detectors proposes using silicon, instead of sapphire, for the final stage suspension due to similar superior cryogenic properties.

The work in this thesis takes the first steps towards designing, building and characterising the

very first cryogenic silicon prototype suspension that needs to be developed in order to realise these full scale suspensions. Silicon, like sapphire, requires a significant amount of research and development in order to exploit these superior cryogenic properties in a practical sense in future detectors.

Work was conducted to investigate the ultimate tensile strength of silicon with an average median value of 60.3 MPa measured for the *control new* laser cut silicon sample set. Experiments were carried out to determine if strengthening of silicon was possible through three techniques; chemomechanical polishing, argon ion etching and wet chemical etching. All surface treatments were found to increase the strength of the silicon samples, with most approximately doubling the median strength of the silicon sets, due to the removal of surface damage, of which the majority was induced by laser cutting. Increasing the strength of silicon suspensions is necessary in order to safely suspend the test mass mirror of next generation detectors. Furthermore, reducing the cross-sectional area of the suspension geometry in order to improve dilution and reduce thermoelastic loss requires higher ultimate tensile strengths of silicon. Consideration was also given to the different surface treatments ability to shape silicon in order to control energy distributions and hence thermal noise within the suspension system.

The first hydroxy-catalysis bonded silicon suspension has been successfully designed, built and characterised as part of this work. Furthermore, this silicon suspension has survived a full cryogenic thermal cycle from room temperature to 10 K and back with the suspension mass being heated to 20 K and 120 K, two proposed detector operating temperatures. This work emphasises the difficulty in handling as well as the levels of tooling and tolerancing required when assembling silicon suspensions, highlighting the engineering challenges of building such a detector. The first estimations of thermal conductivity of this cryogenic silicon suspension were also made.

Tolerancing work for the hydroxy-catalysis bonded silicon suspension was conducted in finite element analysis software before actual assembly of the suspension. This work also highlights the subtleties to be considered for the correct modelling of silicon in finite element analysis software. Furthermore, the first robust modelling of hydroxy-catalysis bonds within complex realistic silicon suspension geometries has been demonstrated. This allows one to compare the hydroxy-catalysis bond loss to thermoelastic loss, surface loss and bulk loss and their respective contributions to the overall suspension thermal noise. The assembled hydroxy-catalysis bonded silicon suspension was modelled in finite element analysis and scaled to a next generation-like detector. This, combined with the previous strength work, allowed for the calculation of suspension thermal noise which can be compared to the Einstein Telescope's proposed design suspension thermal noise target. It was found that the suspension thermal noise of this design of silicon suspension, inclusive of hydroxy-catalysis bonds to join the suspension together, was too high when compared to the design sensitivity target. If one was to remove the bond losses, the suspension thermal noise was reduced by a factor of 15 and found to be below the target,

however the design target includes the entire thermal noise budget of the Einstein Telescope-low frequency detector, whereas the work in this thesis only considers suspension thermal noise arising from the test mass mirror and final stage suspension elements. Furthermore, by previously demonstrating the strength of silicon can be at least doubled through surface treatments, one can run the suspension at double the stress, yielding an improvement in suspension thermal noise by a factor of 1.3. During the course of this work, it was noted that one has to consider higher stresses pushing the bounce mode further down into the frequency band where these cryogenic low frequency detectors are meant to be most sensitive. The high frequency part of these next generation detectors are designed to have superior sensitivity >30 Hz. This has not been a previous consideration for broadband detectors utilising fused silica suspensions, where the aim was always to push the bounce mode frequency as low as possible out of the detection frequency of interest.

The work in this thesis shows one needs to increase the ultimate tensile strength of silicon in order to operate next generation cryogenic silicon detectors at higher operating stresses to reduce suspension thermal noise. Any surface treatment applied must keep the surface loss of the treated silicon to a minimum to avoid this becoming a dominant loss mechanism. Furthermore, this must be developed in parallel with low loss jointing techniques which may be a significant challenge, particularly if suspension geometries are limited from the inability to have full control over the desired manufacturing and shaping of silicon suspension elements which has yet to be demonstrated. This demonstrates the difficulty in the balancing of the many factors required to obtain an optimal cryogenic silicon final stage suspension.

Ultimately, the construction of monolithic silicon suspensions remains a crucial ongoing research area, building on the work demonstrated in this thesis. This constructional challenge may ultimately define the materials, geometries and construction methodology used in these next generation instruments.

Contents

List of Figures

List of Tables

"I'll bet my PhD I wired that correctly"
measures open circuit

- Graeme Eddolls, 2019

Acknowledgements

First and foremost, my thanks go to Dr Alan Cumming. Alan supported me heavily through this PhD, especially during the pandemic period which I found particularly tough. Without his constant supervision, guidance and friendship, this thesis would never have been written. Similarly, my thanks also go to Dr Karen Haughian. Her breadth of guidance and knowledge also supported me immensely and I will always be grateful to her for the steadfast support of my public outreach over the years.

I am also grateful to Dr Marielle van Veggel as my original first supervisor for taking me on as her PhD student. I am particularly thankful for all of the initial advice and support received as a fledging PhD student. Coming out of industry after 3 years was a daunting task and she made me feel like I was more than capable. While I was sad to see her leave academia, it was great to see Marielle pursue her passion for teaching the younger generation.

As part of the breadth of work covered, I am extremely appreciative of my larger academic supervisory team who all provided valuable input to my thesis and who I was very lucky to have access to. Thanks to Prof Giles Hammond for taking over as my primary supervisor and for sharing his extensive knowledge on everything. My thanks also go to Dr Iain Martin who, like Karen and Alan, was always there to chat to scientifically or socially.

It has also been an immense privilege to work closely with Prof Sir James Hough. Jim has been nothing but patient with me, sharing his amazing knowledge in the field and beyond, alongside other assorted musings and stories (particularly discussions on boats). I have also enjoyed many long and interesting discussions with another Jim, Prof James Faller, when he was a regular visitor to Glasgow as part of our research group. I was always entranced listening to the stories of his time working on the Apollo missions and his invention of the Lunar Laser Ranging Program. Similarly, my thanks also go to Mr Alastair Grant for his advice for testing experimental methods and ideas. It never ceases to amaze me to see Alastair, at 90+ years old, still contributing to the field and I can only hope to follow in his footsteps. I also wish to thank the many, many people I have had the pleasure of collaborating with in many countries all over the world. It has been the honour of my life to be a member of the LIGO-Virgo-KAGRA Scientific Collaboration, showing what humanity can achieve when it works together.

I also want to thank, in no particular order, the following people in IGR who have significantly helped me over the years. Thanks to Angus Bell, Borja Sorazu, Stuart Reid, Ross Birney,

Mark Barton, Andrew Spencer, Alasdair Taylor, Bryan Barr, Gregoire Lacaille, Simon Tait, Karl Toland, Richard Middlemiss, Mariela Masso Reid, Jack Callaghan, Jennifer Wright, Peter Dupej and Victoria Graham for all of their help scratching heads over experiments. On the coding side, thank you to Michael Williams, Simon Tait, Christian Chapman-Bird, Daniel Williams, Kristian Anastasiou, Joe Bayley and Peter Dupej for their infinite patience with my bad code. Thanks especially to Liam Cunningham for allowing me to constantly pester him for ANSYS help.

The work in this thesis would not be possible were it not for the support of our amazing engineering and technical staff, Russell Jones, Mikey Perreur-Lloyd, Colin Craig, Stevie Craig and Stephen O'Shea - I have always said I tend to get on better with engineers than physicists and they have all proven this point! Thanks also to Jamie Scott for putting up with my relentless requests for more RAM and HDD memory to sate ANSYS' appetite. Thanks to our research secretaries Karen Hegyi and Jenny Anderson for always fulfilling last minute travel requests for conferences and always being there for some sound advice. An honourable mention must also go to Jim McIver in Chemistry and Andrew Rathbone at Cryomech for helping me, in the middle of a pandemic, build a chemical etching set up and rewire and fix a cryogenic compressor respectively.

My thanks also go to Prof Sheila Rowan, not just for her advice, or for her leadership of the IGR, but as a role model for scientists. To have worked as Chief Scientific Advisor to the Scottish Government through the COVID-19 pandemic, and later serve as the President of the Institute of Physics, all while doing public outreach and promoting diversity in physics is an impressive feat. She has demonstrated that to be a scientist is more than doing and communicating the science, and that we can and should have a role in shaping our society. She has inspired me to spend more of my free time studying and doing public outreach aimed at adults, scientific policy and the public's perception of science. I know I am a better scientist today because of the people mentioned above.

I want to thank everyone in IGR for the many fun social times we have had together from flat parties to conferences, I will miss this aspect of IGR the most. A special shout-out to Karl who stuck with me since we started undergraduate together. We have both gone from (I maintain) *collectively* breaking a Michelson interferometer in undergraduate labs to working for the largest, most sensitive interferometer in the world, LIGO. Furthermore, thank you to my colleagues in the wider School of Physics and Astronomy, particularly those who have supported me since my undergraduate days, specifically Morag Casey, Lynsdey Fletcher and Nic Labrosse and to those in the School who strive to make physics and astronomy more diverse and inclusive to the benefit of everyone.

I want to thank Dr Deborah McNeill for taking me on in 2010 to take part in my first Glasgow Science Festival. Having helped almost every year since, this has allowed me to springboard into work with the Institute of Physics, STEM Ambassadors, Glasgow Skeptics, British Science Association and Festival, Science Ceilidh, InterSci, and so many more science communication

organisations. My 12 years of science communication with these organisations has made me both a better scientist, better communicator and all round better individual as a result, so thanks to them and all the science communicators I have had the pleasure of working with.

Moving away from the physics, I must thank my pals from back home in Troon; Barnaby, Peter, Paul, Jamie, Graham, Ally and Marc. From late night gaming to early morning clubbing, I have really appreciated and continue to value their friendship. I have to give a shout out to a group I affectionately call "the artists" (even though some are lawyers, bakers, or roasters); Sean, Paul, David, Eve, Nick, Grant and Agne and, for similar reasons, to Gem, Leonie, and Maggie, for all collectively introducing me to my artistic and philosophical side. Looking at the world from a social scientist's or artist's perspective through reading books, going to exhibitions, long philosophical and political chats, watching films, going to gigs, etc has massively opened up the world to me, unveiling hidden parts of life and society I never thought I could enjoy.

In this vein, I want to thank some musical artists, a few of whom I have had the pleasure of meeting and chatting to. Without them, I would not have spent many long hours in the labs or writing this thesis so thank you to The Juan Maclean, Man Power, SOPHIE, LCD Soundsystem, Octo Octa and Eris Drew. Thanks for creating some of the best nights of my life so far and for welcoming me into your spaces. Thanks too to Barnaby and Paul for being great DJ partners in Bloc over the years.

My science communication work has also allowed me to link my passion for science communication into my activism (and vice versa) in local communities in and around Glasgow which has been the greatest source of enjoyment, learning and sometimes escape from the PhD. It has been a joy to work with charities, groups and campaigns such as Kinning Park Complex, Refuweegee, Get Glasgow Moving, No Evictions Network, Glasgow Autonomous Space and Woodlands Community Development Trust, to name a few. All of these organisations and the people working within them make me proud to be a Glaswegian. A special shout out to my dear friend Yvonne, who on top of her own studies and work, tirelessly advocates for refugees and asylum seeking people as part of Migrants Organising for Rights and Empowerment. She has taught me so much and allowed me to be part of one of the most amazing communities in this city. It has been a privilege to work with her and other anti-racist activists throughout the city, particularly when COP26 came to Glasgow. I want to thank some of my fellow climate activists, COP26 survivors and pals, Quan, Tess, Bryce, Nick, Lykke, Kate, Mim and Olly for fighting the good fight, creating mischief and being all round good people.

Thanks too to my many University friends over the years, particularly Newman and Nancy, as well as the Glasgow University Skydivers and the Glasgow University Rifle Club (as well as St Andrews, Robert Gordon's, Aberdeen and, unfortunately, Edinburgh). On this note, thanks to Richard, Hannah, Heather and Fraser for the many competition weekends and collective complaining about our PhDs. Thanks to those friends I have made from the PhD Society for creating a massive Whataspp chat for us to bond over our PhD stress. Thanks also to the Glasgow Bike

Polo people for making me feel good, even though my improvement in skill has been somewhat lacking over the years.

I appreciate these acknowledgements are far too long, just like my thesis, and I know I have missed out many people. I am so lucky to have so many friends, from all walks of life, who I love dearly and who enrich my life every day, so thank you to all of you.

Finally, my thanks go to my partner, Michalina, who has supported, loved and cared for me for all my faults and stresses throughout this PhD. I could not have finished this without your support - kocham cię.

To my sister, Kirsty, thank you for all of your support, especially in caring for our parents in recent years. I know this has been extremely difficult on both of us but I hope you are proud of the family you have raised with Dean and I hope you know how proud I am of my two nephews, Alex and Harris.

To my Mum and Dad, thank you for letting me follow my dreams. I could not have asked for a better upbringing and I hope I have made you proud by taking as many opportunities presented to me as possible. As I pursue my next dream of becoming a pilot, know that these dreams are all achieved thanks to your love and support.

Preface

This thesis documents work undertaken by the author spanning February 2018 to August 2022, in order to fulfil the requirements for the degree of Doctor of Philosophy. This work took place within the Institute for Gravitational Research, School of Physics and Astronomy, University of Glasgow.

Chapter 1 is an introduction to gravitational waves. This chapter details astrophysical sources and historic detection attempts, alongside modern techniques, which made the first direct detection of gravitational waves, and their associated noise sources. The chapter goes on to discuss the future of gravitational wave detection. This chapter is based on published literature reviewed by the author.

Chapter 2 discusses the concept of using silicon as a detector material for next generation cryogenic gravitational wave detectors. The work covers the properties of silicon and compares to the current detector material, fused silica, from published literature reviewed by the author. This also includes discussion covering suspension dynamics as well as challenges that will be encountered with silicon suspensions. Work conducted on welding silicon was carried out with Dr Alan Cumming.

Chapter 3 details the surface treatments applied to silicon in order to increase its ultimate tensile strength. Strength testing was conducted on a strength tester originally designed by Mr Russell Jones, with further modification made specific to this work. The initial strength testing set up work and training for the author was conducted by Dr Alan Cumming. Dimensional characterisation of samples was carried out using the LIGO fibre profiler, previously constructed by Dr Alan Cumming and the thin fibre profiler previously constructed by Dr Karl Toland. Surface treatments were discussed with and proposed by the author, Dr Marielle van Veggel, Dr Karen Haughian, Dr Alan Cumming, Dr Gregoire Lacaille, Prof Giles Hammond, Prof Sir James Hough, Prof James Faller and Prof Sheila Rowan. Chemomechanical polishing was conducted by the author with advice from Dr Marielle van Veggel and Mr Tom McGroggan from Logitech Limited. Fellow PhD student, Ms Victoria Graham, measured the pH of the polishing solution. Argon ion etching was conducted in equipment at the University of Strathclyde in collaboration with Mr Gavin Thomson, Dr Paul Hill, Dr Ross Birney and Prof Stuart Reid. Chemical etching for the initial KOH and HNA JWNC etch samples was conducted with Dr Matthew Smith. Chemical etching was conducted at ITM Leven with significant advice from Mr Iain McDonald.

Further chemical bath design and advice was provided by chemical technician Mr James McIver from the School of Chemistry, University of Glasgow. All fuse ends were sandblasted and drilled by Mr Steven Craig.

Chapter 4 contains work designing, building and assembling the prototype cryogenic silicon suspension. The design concepts, including the bonding jig, were discussed with and proposed by the author, Mr Russell Jones, Mr Michael Perreur-Lloyd, Dr Marielle van Veggel, Dr Karen Haughian, Dr Alan Cumming, Dr Iain Martin, Prof Giles Hammond, Prof Sir James Hough and Prof Sheila Rowan. Building and assembly of the suspension was carried out with Dr Alan Cumming with significant advice from Mr Russell Jones, Dr Karen Haughian, Dr Gregoire Lacaille, Dr Alan Cumming, Dr Iain Martin, Prof Giles Hammond and Prof Sir James Hough. The 4 HC bonded silicon suspensions were assembled and bonded by Dr Karen Haughian and Dr Gregoire Lacaille. The cryostat platform was designed and assembled by Mr Russell Jones. The aluminium and stainless steel masses were made by Mr Colin Craig and Mr Steven Craig. Mr Andrew Rathbone, of Cryomech, offered significant advice, remotely, on the maintenance and repair of the Cryomech compressor. The LabView program used was originally developed by Dr Karen Haughian, before separate modifications were made by Dr Alan Cumming, Dr Mariela Masso Herrera and the author. Dr Naren Hoovinakatte and Prof Giorgio Frossati, of Leiden Cryogenics, offered significant advice on the operation and maintenance of the cryostat system. Dr Iain Martin, Dr Alan Cumming, Dr Karen Haughian and Prof Sir James Hough helped significantly with proposed experimental heat steps applied and subsequent theoretical analysis of the cryogenic suspension, contained in appendix B. Professor Sir James Hough significantly contributed to the error analysis of this work, contained in appendix B.

Chapter 5 covers finite element analysis using ANSYS Mechanical APDL software. Significant advice, guidance and training on modelling techniques was provided by Dr Alan Cumming. Further significant advice on modelling techniques was provided by Dr Liam Cunningham and Dr Karen Haughian. Software-specific advice was given by Mr Luke Reszke from Wilde Analysis.

The code in appendix E was written and developed by the author.

The code in appendix F was written and developed by the author with advice from Dr Alan Cumming and Prof Giles Hammond.

All of the above work requiring design and manufacturing of mechanical components that were not built by the author, or ordered externally, were designed by Mr Russell Jones and Mr Michael Perreur-Lloyd, and manufactured by Mr Colin Craig and Mr Stephen Craig.

Two Python plots were produced in this thesis with the aid of Mr Christian Chapman-Bird. All images used in this thesis that are not the original work of the author have been obtained with permission and relevant licenses, where appropriate, which are available upon request.

The author's PhD was generously funded by the Royal Society.

The work in this thesis has previously been presented at:

- LIGO-Virgo Collaboration meeting March 2019, Geneva, Switzerland (talk, remote)
- Gravitational Wave Advanced Detector Workshop (GWADW) May 2019, Elba, Italy (poster, in person)
- Gravitational Wave Science and Technology Symposium (GRASS) October 2019, Padova, Italy (talk, in person)
- LIGO-Virgo-KAGRA Collaboration meeting March 2021, online only (talk, remote)
- GWADW May 2021, online only (talk, remote)
- LIGO-Virgo-KAGRA Collaboration meeting March 2022, online only (talk, remote)
- GWADW May 2022, online only (poster, remote)
- LIGO-Virgo-KAGRA Collaboration meeting September 2022, Cardiff, UK (talk and poster, in person)
- Einstein Telescope Instrument Science Board meeting October 2022, L'Aquila, Italy (plenary talk and chairing of two workshop parallel sessions, in person)

Chapter 1

Gravitational wave detection

1.1 Introduction

In 1915, Albert Einstein published his General Theory of Relativity [1] [2]. This novel theory led to the concept of curvature of spacetime. This curved spacetime is caused by the presence of mass which warps the spacetime around it, having a measurable effect on other objects. This variation in warping manifests itself in the phenomenon we call gravity. As objects with mass evolve through spacetime, the local spacetime curvature changes to reflect this.

Einstein predicted these changes in spacetime curvature propagate outwards, at the constant, c, the speed of light, in a wave-like form. The General Theory of Relativity states that energy is carried away in the form of gravitational radiation. An event that emits gravitational radiation is so intense that it has the potential to outshine all of the stars in the universe due to the energy emitted [3]. These astronomical events produce localised changes, or ripples, in the fabric of spacetime which are known as gravitational waves.

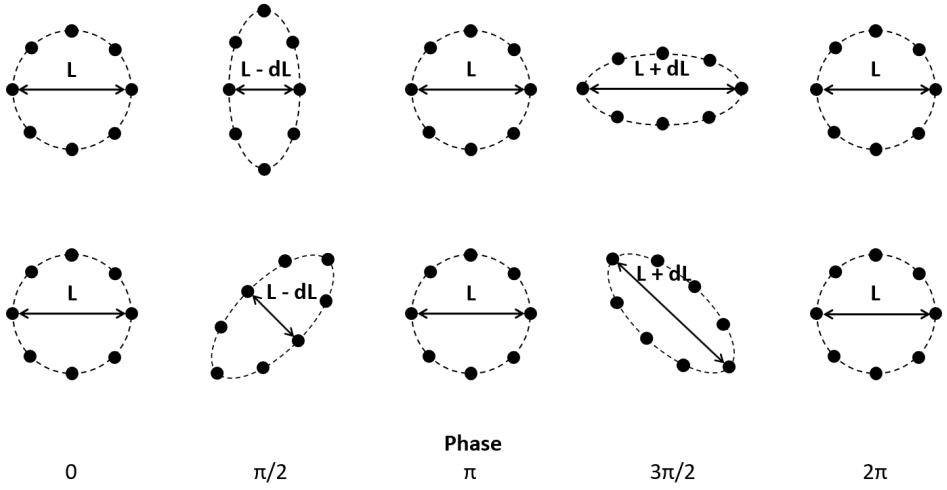

Figure 1.1: Black dots represent a ring of free particles with mass, distorted by each polarisation, $+$ (top) and $\times$ (bottom), of a gravitational wave travelling perpendicular to the page.

Gravitational waves are a gravitational analogue to electromagnetic radiation, where an accelerated charge generates electromagnetic waves, but instead asymmetrical accelerated masses generate gravitational waves [4]. Gravitational waves are fundamentally quadrupolar in nature with two polarisation states, $+$ and $\times$, as shown in figure 1.1 [5].

The strength of a gravitational wave is characterised by the strain amplitude, h, a dimensionless variable, which is given by [5]:

$$h = \frac{2dL}{L} \tag{1.1}$$

where dL is the change in separate between two test massses separated by the distance, L.

1.2 Gravitational wave sources

Every object with an asymmetry of mass that accelerates produces gravitational waves, however the strain amplitudes produced from sources on typical human scales are far too small to be detected [6]. This is due to gravity's relative weakness compared to the other three fundamental forces; the weak nuclear, strong nuclear and electromagnetic forces. Furthermore, the fabric of spacetime is enormously stiff. If one could measure spacetime's Young's modulus in a classical sense, it would be found to be approximately 10^{20} times greater than that of steel ($E_{steel} \approx 200\,\text{GPa}$) [7].

This combination of factors means one requires large amounts of mass, undergoing a large acceleration, in order to produce detectable strain amplitudes, of the order 10^{-21} [8]. There are a number of astrophysical sources throughout the universe that meet these criteria.

1.2.1 Compact binary coalescence

Compact binary coalescences (CBCs) occur due to massive bodies orbiting one another in space. Such massive bodies able to produce detectable strain amplitudes are black holes and neutron stars due to their extreme density. As such, CBCs are typically split into 3 categories; binary black hole (BBH), neutron star black hole (NSBH) and binary neutron star (BNS) mergers. An artists impression of a NSBH is shown in figure 1.2(a).

Just like planets orbiting stars, black holes and neutron stars can orbit one another around a common centre of mass. As their orbit decays, the energy lost is emitted in the form of gravitational radiation. In the final death throes, called the inspiral, these masses are orbiting at relativistic speeds with enough energy to produce strong gravitational radiation that can be detected before they merge together.

Gravitational wave strain amplitudes of CBC inspirals are expected to be [9]:

$$h = 10^{-23} \left(\frac{M}{M_\odot} \right)^{\frac{5}{3}} \left(\frac{f}{100\,\mathrm{Hz}} \right)^{\frac{2}{3}} \left(\frac{100\,\mathrm{Mpc}}{r} \right) \tag{1.2}$$

where $M_\odot$ represents the solar mass, f is the frequency of the gravitational wave, r is the distance from the source and M is the sum of the mass of the two merging objects, $m_1 + m_2$.

The characteristic signal of two inspiralling masses for CBC would be an oscillating sinusoid as the two objects orbit one another [10]. As they inspiral, both the frequency and amplitude would increase to a maximum at the point where the two objects merge into one. This produces a signal that, if shifted into the human auditory frequency band, would sound like a "chirp" noise [11]. There is a very short period in the newly formed singular mass, called the ringdown, as its vibrates and settles, after which no production of gravitational waves occurs due to the new symmetry of the singular mass [12].

Detection of CBCs allows scientists to test some predictions of General Relativity as well as providing natural test beds to constrain other theories of gravity. Furthermore, being able to measure black hole spins will help to understand the dynamics of black hole formation [8].

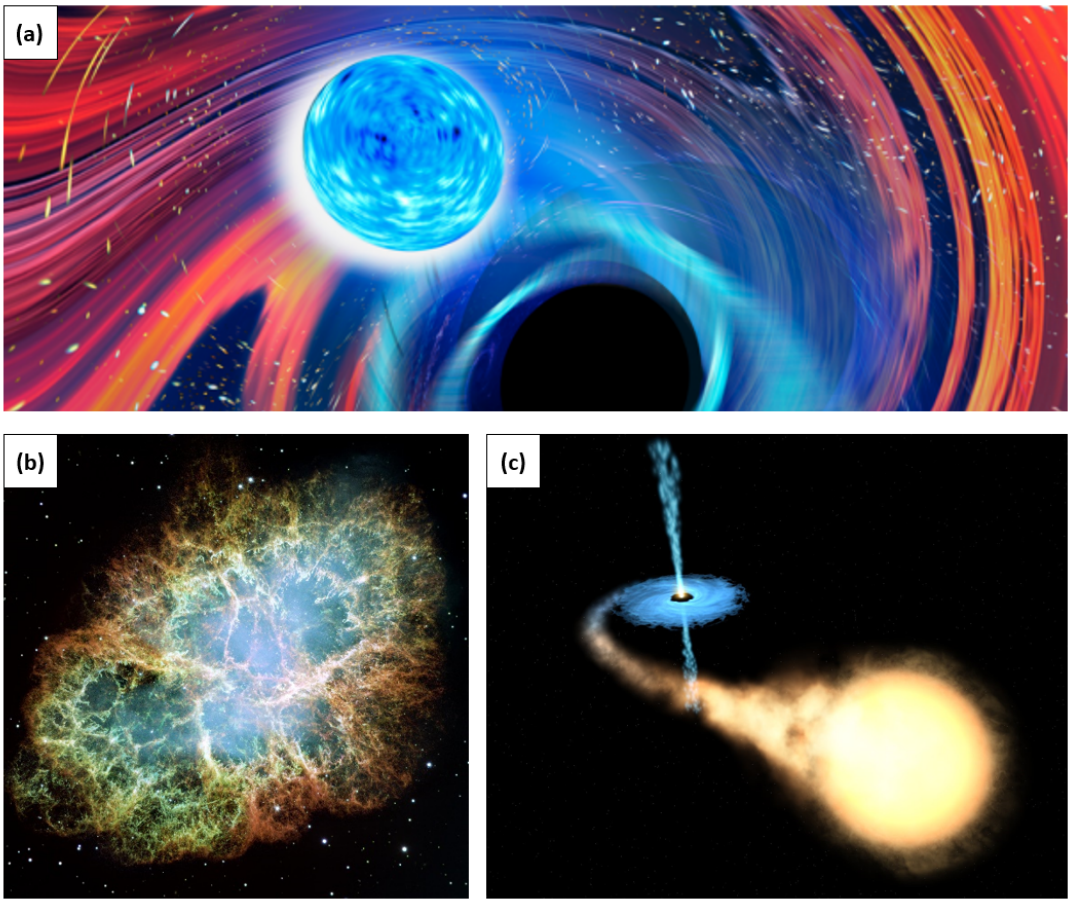

Figure 1.2: (a) Artist impression of a NSBH merger [13]. (b) Image of the Crab Nebula, formed through core-collapse, obtained with the Hubble Space Telescope [14]. (c) Artist impression of a LMXB system [15].

1.2.2 Burst events

Burst events are defined as signals with large amplitudes that occur over a short duration, on the order of a few wave periods [16]. Astronomical events that may produce burst-like signals could be from core-collapse supernovae; an example of the aftermath of such a collapse is shown in figure 1.2(b). If the core-collapse of a supernova was entirely symmetric, no gravitational waves would be produced. Where there is an asymmetry during the supernova, gravitational waves would be produced with a strain amplitude given by [8]:

$$h = 6 \times 10^{-21} \left(\frac{E}{10^{-7} M_{\odot}} \right)^{\frac{1}{2}} \left(\frac{1\,\mathrm{ms}}{T} \right)^{\frac{1}{2}} \left(\frac{1\,\mathrm{kHz}}{f} \right) \left(\frac{10\,\mathrm{kpc}}{r} \right) \tag{1.3}$$

where E is the total energy radiated and T is the time to collapse.

CBCs and burst events represent the majority of transient gravitational wave signals one would expect to detect, although other sources are believed to produce transient signals such as collapsars and isolated neutron star oscillations or instabilities [17].

Detection of burst events will significantly increase our understanding of supernova collapse as well as provide excellent accuracy with which to measured the speed of gravitational waves [8].

1.2.3 Continuous waves

Source of continuous gravitational wave emissions come from a variety of astrophysical sources. Singular pulsars are believed to produce sources of continuous gravitational waves. The first discovered pulsar was observed by Jocelyn Bell in 1967 [18]. Pulsars are rapidly rotating neutron stars which emit radio pulses periodically, with detected pulsars ranging in spin frequencies from 42 mHz to 716 Hz [19] [20]. These radio pulses are emitted from the magnetic poles of the pulsar which, when orientated in the direction of the observer, are seen as a periodic pulsing radio source, akin to a lighthouse beam sweeping past.

Pulsars are some of the most dense objects in the universe, with radii only 10s of kilometres across but containing up to 3 solar masses [21]. Even though this produces extremely strong localised gravitational fields which should provide an extremely uniform spherical star, non-axisymmetric distortions, velocity perturbations and free precession can lead to the emission of gravitational waves from such sources. These emissions lead to continuous, periodic gravitational wave emissions with gradually shifting frequency as energy is lost through this emission mechanism [22].

For a single pulsar, emitting gravitational waves at twice its spin frequency, the strain amplitude can be expressed as [23]:

$$h_0 = 10^{-24} \left(\frac{\varepsilon}{0.237} \right) \left(\frac{1\,\mathrm{kpc}}{r} \right) \left(\frac{\nu}{1\,\mathrm{Hz}} \right)^2 \left(\frac{I_{zz}}{10^{-38}\,\mathrm{kg\,m^2}} \right) \tag{1.4}$$

where ε is the pulsar's equatorial ellipticity, ν is its spin frequency and I_{zz} is its principle moment of inertia. The spin frequency is twice the pulsar's spin frequency due to specific search techniques used. These techniques relate to the fact that emission near the rotational frequency for precessing stars is likely to be offset from the observed frequency by some factor dependent on unknown details of the star structure [23] [24].

Another continuous source of gravitational waves is believed to come from low mass x-ray binarys (LMXBs), an example of which is shown in figure 1.2(c). Neutron stars, such as pulsars, or black holes can also exist in a binary system with a low mass donor star. This donor star can be stripped of its matter by the companion neutron star or black hole due to the companion's strong gravitational field. This process generates what is known as an accretion disk, which causes an increase in the angular momentum of the companion object, and can be one of the brightest x-ray sources in astronomy [25]. This system is known as a LMXB and can reach a point of instability known as the Chandrasekhar-Friedman-Schutz (CFS) instability where the rotation of the companion object becomes non-axisymmetric, leading to the production of continuous gravitational waves.

The stain amplitude for a LMXB system can be expressed as [26]:

$$h = 3 \times 10^{-27} \left(\frac{1\,\mathrm{kHz}}{m f_s} \right)^{\frac{1}{2}} \left(\frac{F_x}{10^{-8}\,\mathrm{erg\,cm^{-1}\,s^{-2}}} \right)^{\frac{1}{2}} \tag{1.5}$$

where $m = 4$ is the mode number, f_s is the spin frequency and F_x, is the x-ray flux measured at the observer.

The detection of continuous waves will allow scientists further insight of neutron star interiors, especially the crust physics and elastic properties at low temperatures [27].

1.2.4 Stochastic background radiation

The stochastic background gravitational radiation is another potential source of gravitational waves. A wide variety of astrophysical and cosmological sources are expected to make up this stochastic gravitational wave background signal [28]. These signals are defined as "a random gravitational-wave signal produced by a large number of weak, independent, and unresolved sources" [29]. Background sources of these stochastic signals may include the gravitational waves resulting from the Big Bang or asymmetric phase transitions in the early universe. These signals results in the superposition of the combined sources' gravitational radiation that would be isotropic in nature and range over an extremely wide frequency [30]. The stochastic background gravitational radiation can be thought of as an analogue to the cosmic microwave background radiation (CMBR) in the electromagnetic spectrum.

Detection of the stochastic background could yield new information on the fundamental nature of the laws of physics at energies higher than one could ever hope to achieve in a laboratory [8].

1.3 Experimental detection of gravitational waves

Einstein himself had no confidence that the detection of gravitational waves was ever possible [30]. Indeed, from equations 1.2, 1.3, 1.4 and 1.5, one can see the typical strain sensitivities required to detect gravitational wave sources on Earth range in the order of $h \approx 10^{-21}$–10^{-27} From equation 1.1, even for the strongest gravitational wave signals, one would require the ability to detect length changes of a thousandth the diameter of an individual proton for two masses separated hundreds of metres apart. This is clearly an extremely difficult experimental challenge and explains why Einstein thought the direct detection of gravitational waves was impossible.

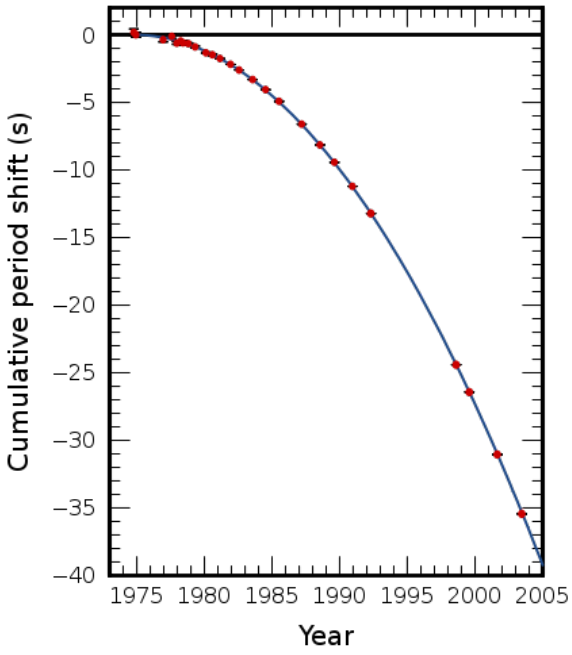

Figure 1.3: Orbital decay of PSR B1913+16 measured by Hulse and Taylor compared to the prediction of the General Relativity [31].

Gravitational waves were first *indirectly* detected by Hulse and Taylor [32] [33]. Their work observed the orbital decay of a binary pulsar system, PSR B1913+16, through its radio signal emissions. As shown in figure 1.3, the orbital decay of this binary pulsar system was consistent with the prediction of the emission of gravitational radiation. This discovery led to the award of the Nobel Prize in Physics in 1993 [34]. Furthermore, this led to another validation of Einstein's General Theory of Relativity [33] and encouraged the scientific community to pursue the *direct* detection of gravitational waves.

1.3.1 Resonant bar detectors

One of the first attempts at direct gravitational wave detection was made by Joseph Weber. His proposal was to install two large aluminium cylinders (150 cm, 61 cm wide, weighing 1.5 t) separated by a large distance. Weber proposed that a passing gravitational wave would excite and resonate these aluminium cylinders or bars at a narrowband frequency, with the vibration measured with piezoelectric crystal transducers [35] [36]. The separation of the cylinders was necessary in order negate excitation from local noise sources at one bar detector; a coincident excitation at two separate detectors sufficiently distant would imply an astrophysical source. Initially, the two detectors were separated by 2 km [37], later increasing to 1000 km separation with detectors located at Argonne National Laboratory and at the University of Maryland [36].

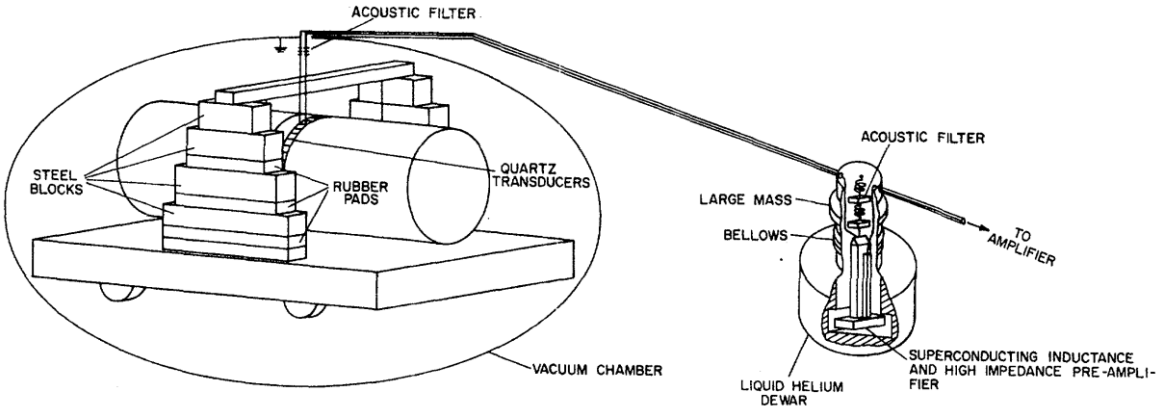

Figure 1.4: Schematic of a resonant bar detector [38].

These bars were designed to detect gravitational waves at 1660 Hz as Weber expected this frequency would be swept through during the emission in a supernova collapse [36]. A schematic of the detector is shown in figure 1.4.

In 1969, Weber claimed to have directly observed gravitational waves [36]. One year later, Weber claimed to have detected 311 such coincident events of gravitational waves [39]. A study of the distribution of these events led Weber to conclude the signals were emanating from the galactic centre. Analysis suggested the sensitivity of Weber's resonant bars to be $\approx 10^{-16}$ making the detection of strain amplitudes from astrophysical events unlikely [40].

Other resonant bar detectors were built around the world at Moscow State University, the IBM Thomas J. Watson Research Center (Yorktown Heights), the University of Rochester (New York), Bell Labs (New Jersey), the Max-Planck-Institut für Physik und Astrophysik (Munich) and the University of Glasgow [41]. The bar detector built in Glasgow by Ronald Drever and James Hough, shown in figure 1.5, varied from Weber's design where the original bar detector was split in half longitudinally to create two aluminium bars with a gap between them. Piezoelectric transducers were placed into this gap between the bars in order to monitor any length change between them. None of these resonant bar detectors built verified Weber's results [42].

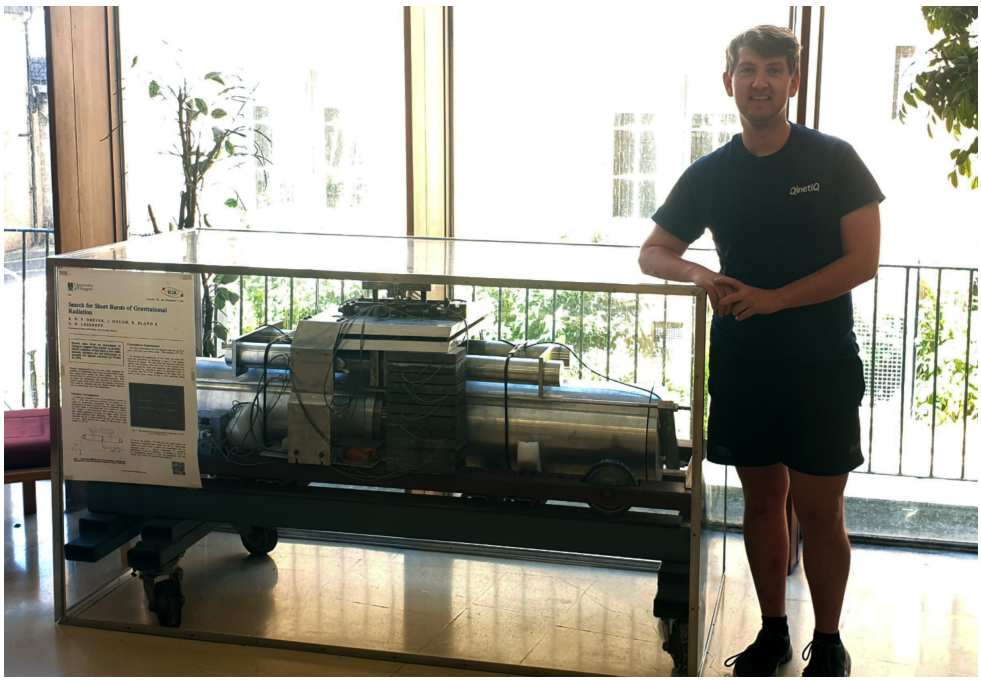

Figure 1.5: The author with Drever and Hough's resonant bar detector situated in the University of Glasgow's Kelvin building, metres from where the work in this thesis was conducted. Photograph taken by Mr Thejas Seetharamu.

Nonetheless, improvements in resonant bar detectors were made. The main limitations to sensitivity of resonant bar detectors arose from thermal excitation of the normal modes and electronic noise [43]. The most obvious way to reduce detector noise level was to reduce the temperature of the resonant bar [40]. This led to the development of multiple cryogenic resonant bar detectors such as the Allegro [44], AURIGA [45], Explorer [46], NAUTILUS [47] and Niobe [48] detectors alongside other cryogenic resonance detectors that are spherically shaped, such as MiniGRAIL [49] and the Brazilian spherical detector [50]. Such detectors have achieved impressive strain sensitivities of $10^{-21}\ \sqrt{\text{Hz}}$ but are however limited to narrow bands of 1–30 Hz [51] meaning only the strongest gravitational events, at just the right frequency, have a chance of being detected.

These resonant bar detectors have since been superseded by a different detection technique which can not only match and improve on this $10^{-21}\ \sqrt{\text{Hz}}$ sensitivity but offers broadband detection capabilities - laser interferometry.

1.3.2 Laser interferometry

In the 1970s, a laser interferometric gravitational wave detector was built by Moss, Miller and Forward [52]. The detector was a Michelson interferometer, the basic set up of which is shown in figure 1.6 [53].

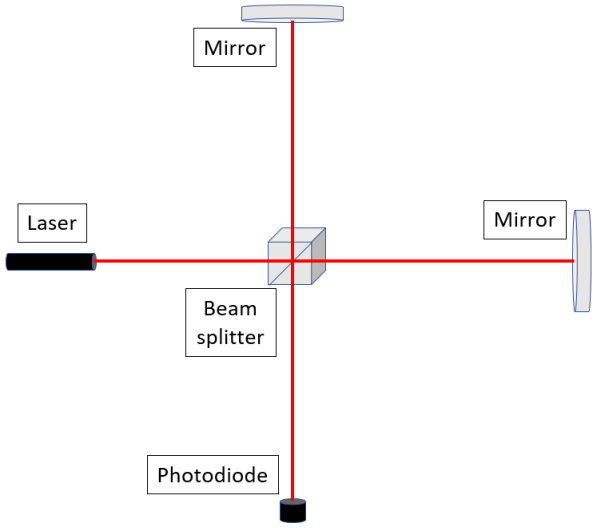

Figure 1.6: Schematic diagram of a basic Michelson interferometer set up.

The interferometer splits the input laser beam in to two perpendicular paths or "arms". The two separate beams reflect off of mirrors situated at the end of the detector arms. This reflected light is then recombined at the beam splitter with the resultant combination measured using a photodiode.

If a gravitational wave were to pass through this apparatus, due to the polarised nature of the gravitational radiation, it would cause a differential change in the length of the arms due to distortion of the localised spacetime (provided the wave did not travel parallel to the plane of the beam paths). For gravitational wave detection, these reflective mirrors are commonly referred to as "test masses" since they are the masses that are deflected when local spacetime curvature changes due to a gravitational wave. The recombination of the two laser beams after reflecting off the test mass mirrors creates an interference pattern, with a differential change in arm length resulting in a change in this pattern which can be recorded at the photodiode.

One can set up the interferometer to keep the beam path length the same - this is called locking. To reduce laser noise while maximising signal, one can lock this set up so a dark fringe is observed at the photodiode. Locking is achieved using a feedback system which applies forces on the mirrors to keep the output interference pattern constant. The error signal in this feedback loop is often used as the output signal of the detector [54].

The Moss, Miller and Forward interferometer had arm lengths of 2 m. This interferometer achieved a sensitivity of $6 \times 10^{-15}\ \sqrt{\text{Hz}}$ at frequencies near 1 kHz [40]. One can increase the sensitivity of a gravitational wave detector by increasing the length of the interferometer arms. This is reflected in equation 1.1, where for a given strain h, increasing L will also increase dL and therefore making the arm length change more easily detectable. The optimal sensitivity of a detector in this configuration is when the time it takes for the beam to travel the length of the arm, reflect and recombine at the beam splitter is half the period of a gravitational wave [30]. Typical sources, such as CBC or burst events, are expected to range in frequency from 1–1000 Hz.

Therefore the optimal interferometer arm length for this frequency range is 75–75 000 km. These distances are either unfeasible or extremely costly since the curvature of the Earth, which limits the straight line distance of a laser to ≈4 km, would require the significant engineering of large parts of the detector arms to be built underground [30].

It is therefore necessary to look at other interferometry techniques that allow for an increase in storage time for the light in the arms.

1.3.3 Delay line interferometry

Alternative interferometry techniques use a second set of mirrors added into the traditional Michelson design. This is called delay line interferometry, a technique proposed by Weiss for the use of gravitational wave detection [55]. Two such delay line inteferometer design schemes are shown in figure 1.7.

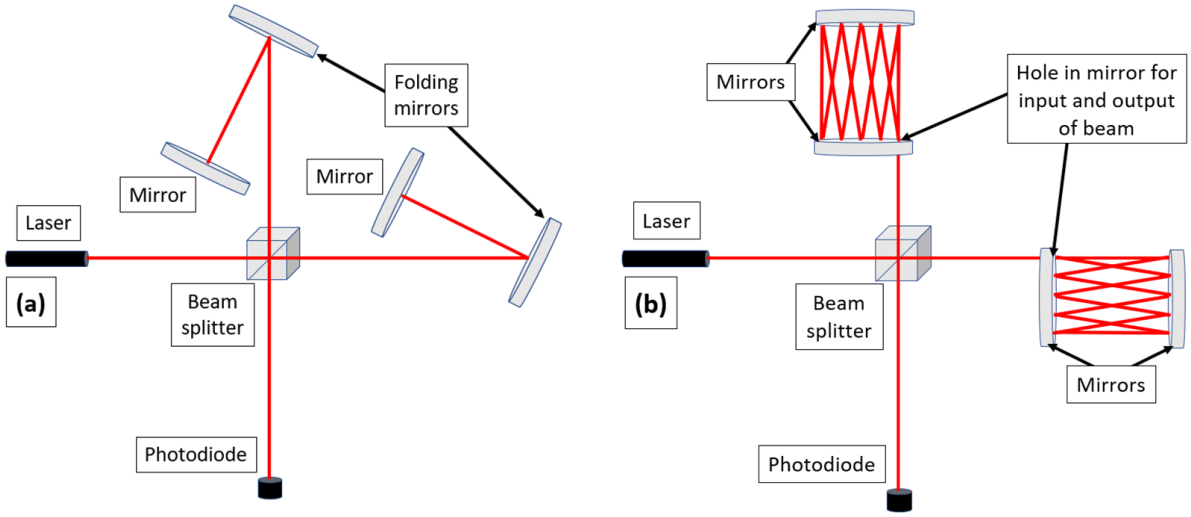

Figure 1.7: (a) Folded arm interferometer schematic. (b) Herriott delay line interferometer schematic.

Figure 1.7(a) shows the use of folded mirrors which can effectively increase the optical path length without significant increasing the physical length of the interferometer. A second scheme, shown in figure 1.7(b), uses curved mirrors with a small gap in the inner mirrors that allows for the beam to be reflected multiple times before returning the same input path [56].

1.3.4 Fabry-Pérot interferometry

Delay line interferometry can be improved yet again to further increase the sensitivity of the detector using what is known as a Fabry-Pérot cavity [57]. This technique for gravitational wave detection was first developed at the University of Glasgow in the 1980s [58]. A schematic of a Fabry-Pérot cavity is shown in figure 1.8.

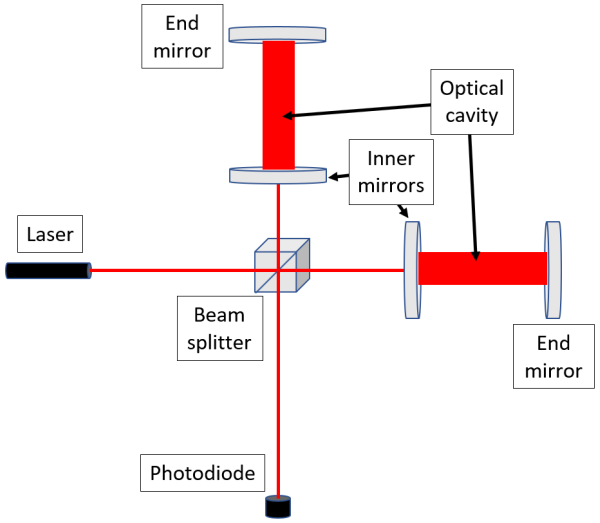

Figure 1.8: Schematic of an interferometer with Fabry-Pérot arm cavities.

This interferometer set up utilises partially transmissive inner mirrors that allow light coming from the beam splitter to pass through the rear of the inner mirror into the optical cavity where it is reflected back and forth multiple times between the highly reflective inner and end mirrors. This set up allows for smaller, simpler mirrors to reflect the light along the optical cavity path. The cavity requires the positions of the mirrors to be an integer of half the wavelength of the laser used to build up a resonance. This allows for energy to build up within the cavity and due to the multiple reflections, significantly increases the effective arm length, and hence sensitivity, of the detector.

1.3.5 Power and signal recycling

One can make further improvements to this detector by optical techniques, such as power recycling and signal recycling.

Power recycling is used to increase the circulating laser power by adding a partially transmissive mirror at the laser input of the interferometer, as shown in figure 1.9(a). Similar, in principle, to the Fabry-Pérot cavities themselves, the power recycling mirror (PRM) creates a cavity with the interferometer output [59]. Since the interferometer is typically operated so a dark fringe appears at the photodiode, this means most of the light is reflected back towards the laser. This light is reflected by the PRM and is recycled back into the interferometer, increasing the stored laser power and hence sensitivity (discussed in section 1.4.3).

Similarly, signal recycling works by reflecting light outputted from the interferometer towards the photodiode, using a partially transmissive mirror, called the signal recycling mirror (SRM), as shown in figure 1.9(b). With the output to the photodiode held on a dark fringe, a passing gravitational wave only produces a small amount of light. By forming a cavity between the SRM and interferometer output, this small signal is then amplified. By determining the

position of the SRM, one can shift the frequency at which the detector is most sensitive [60].

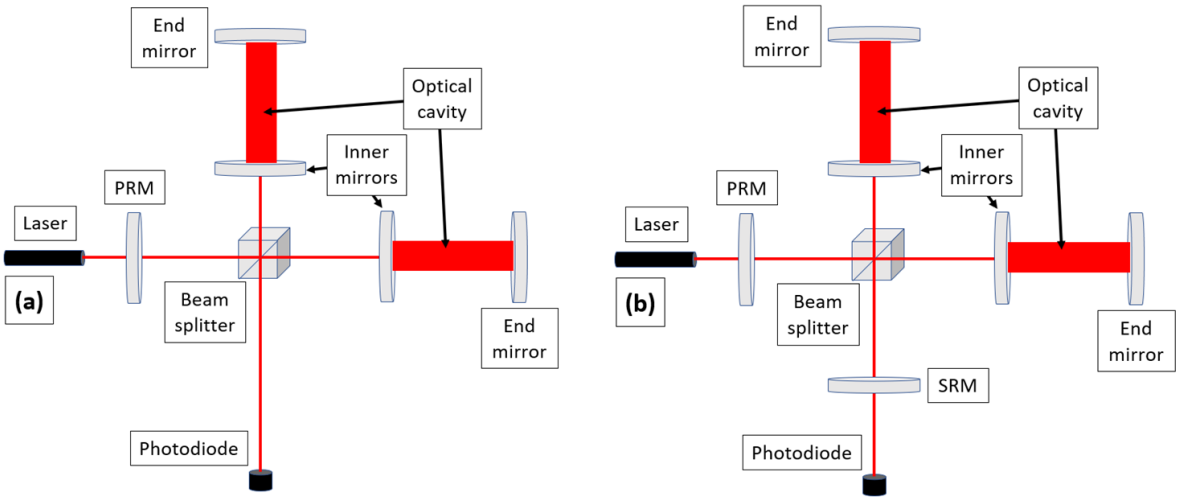

Figure 1.9: (a) Schematic of a Fabry-Pérot Michelson interferometer with a PRM. (b) Schematic of (a) with a SRM, creating a Fabry-Pérot dual-recycled Michelson interferometer.

By using both power and signal recycling, this creates what is called a Fabry-Pérot dual-recycled Michelson interferometer.

1.4 Noise sources in gravitational wave detectors

This section discusses the main sources of noise arising in such a detector system.

1.4.1 Seismic noise

Seismic noise arises from natural causes as well anthropogenic activity. Natural seismic events can include earthquakes anywhere on Earth, Earth's own microseismic tremors and weather and oceanic events. Human activity such as traffic, local footfall, maintenance and operational equipment in and around the detector itself such as air conditioners, pumps, etc also create localised vibrations that can couple into the detector mirrors [61] [62].

Seismic noise varies by location and can be affected by local geology or geographic location. This noise source is most dominant in the low frequencies (<10 Hz) of the detector bandwidth with a typical acceleration spectral density of $10^{-7}\,\mathrm{f}^{-2}\mathrm{ms}^{-2}/\sqrt{\mathrm{Hz}}$ for a quiet geological site [63]. This is well above the required horizontal isolation for gravitational wave test mass mirrors by an order of $\approx 10^{9}$ [40]. This horizontal seismic noise is reduced by taking advantage of the physics of a simple pendulum. For a pendulum system, the transfer function from the suspension point to the suspended mass falls off as f^{-2} above the pendulum resonance, therefore hanging the test mass mirrors on a series of pendulums can help mitigate this noise source [40]. Furthermore, coupling of vertical noise through the horizontal axis is typically around 0.1 %, as discussed in

section 2.4.3 [64] [65]. Therefore, vertical isolation is also required and is provided in the form of cantilever spring designs from which the suspension point to the suspended mass is affixed [40]. This is shown in more detail in section 1.5.4.

One can also reduce seismic noise by hanging the suspension from both passive and active seismic isolation platforms where an array of accelerometers, geophones and seismometers can monitor the seismic motion providing a feedback for the active isolation systems [66].

Another method for reducing seismic noise is to build the detector underground. The aim of building the detector underground is to increase the detectors sensitivity at even lower frequencies and to push the seismic noise limit to <1 Hz [67].

1.4.2 Gravitational gradient noise

Gravitational gradient, also known as Newtonian, noise results from the test mass mirrors directly coupling to local fluctuations in the gravitational field. These can arise from changes in mass or density in and around the detector. Such changes in density can come from seismic surface waves travelling near the detector or a localised density traversing the detector, such as a vehicle or aircraft in proximity to a test mass.

There is no way to shield the detector from gravitational gradient noise. However, models can be developed, using arrays of seismometers placed around the detector site, similar to those used for seismic isolation, in order to subtract the measured gravitational gradient noise from the signal [68] [69]. Similar to seismic noise, this low frequency noise source can be mitigated by building the detector underground [70].

For both seismic and gravitational gradient noise, a more radical solution would be to place the detector in space, free from terrestrial noise, with smaller gravitational gradients, giving detector sensitivities able to probe well below 1 Hz. This is discussed more in section 1.7.

1.4.3 Quantum noise

Quantum noise can be split into radiation pressure noise and photon shot noise. These two sources of noise combined limit the detector by what is known as the standard quantum limit (SQL).

Radiation pressure noise arises from fluctuations in the number of photons reflected from the interferometer mirrors. As the mirror reflects these photons, some amount of momentum, and hence force, is imparted into the reflecting mirrors. This fluctuation in the number of photons leads to a fluctuation in the force and hence position of the mirrors resulting in noise. An example of radiation pressure noise is shown on the left of figure 1.10. The amplitude spectral density of radiation pressure noise can be expressed as [30]:

$$h_{RP} = \left(\frac{1}{mf^2L}\right)\left(\frac{\hbar P}{2\pi^3 c\lambda}\right)^{-\frac{1}{2}} Hz^{-\frac{1}{2}} \tag{1.6}$$

where m is the mass of the mirror, f is the frequency of the gravitational wave, L is the arm length of the interferometer, $\hbar$ is the reduced Planck's constant, P is the input power of the laser, c is the speed of light and λ is the wavelength of the interferometer laser. From this equation, one can see that radiation pressure noise can be reduced by increasing the mass of the test mass mirror.

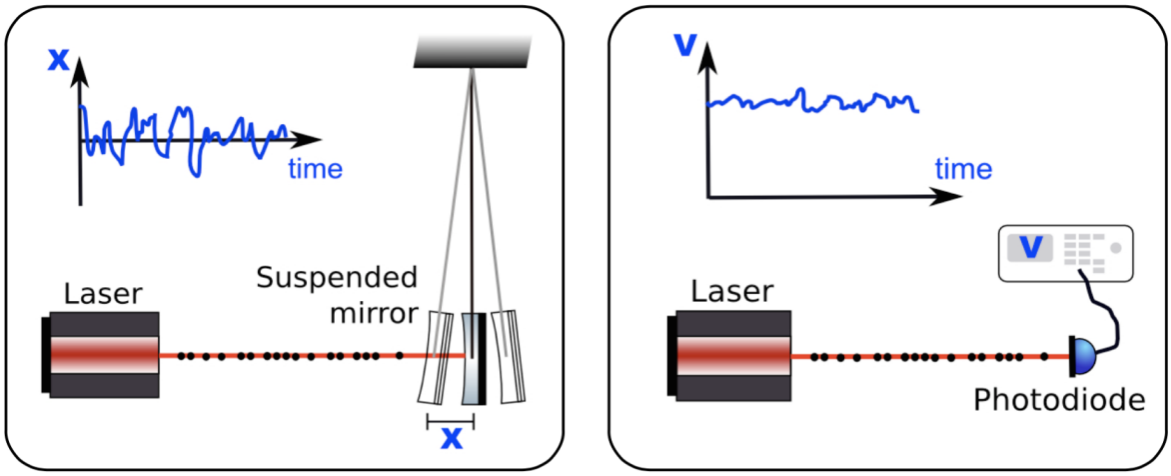

Figure 1.10: Example of radiation pressure noise impinging on interferometer mirrors (left). Example of photon shot noise as measured at the photodiode (right) [71].

Photon shot noise, measured directly at the photodiode, arises as a result of the unequal distribution of photons exiting the input laser which follows a Poisson distribution [72]. An example of photon shot noise is shown on the right of figure 1.10. The uncertainty in the number of photons, N, is $\sqrt{N}$ over a time period, t. The amplitude spectral density of photon shot noise can be expressed as [30]:

$$h_{SN}(f) = \frac{1}{L}\left(\frac{\hbar c\lambda}{2\pi P}\right)^{-\frac{1}{2}} Hz^{-\frac{1}{2}} \tag{1.7}$$

where all terms are as defined for equation 1.6. From this equation, one can see that photon shot noise can be reduced by increasing the laser input power.

Increasing the laser power to reduce photon shot noise will increase the radiation pressure noise. Therefore, a balance must be struck to find the noise minimum, where $h_{SN}(f) = h_{RP}(f)$, which is the SQL of the detector, the best sensitivity one can get at a certain frequency. The SQL is a manifestation of the Heisenberg Uncertainty Principle whereby there is a trade off between higher frequency sensitivity from reduction in photon shot noise against lower frequency sensitivity due to the increase in radiation pressure noise.

Aside from increasing the laser power, novel methods can bypass the SQL. Techniques include optical squeezing [73], optical springs [74] and speed-meters [75].

1.4.4 Thermal noise

Thermal noise arises from the fact the interferometer runs at some finite temperature, T, and therefore all of the atoms that make up the detector itself have some thermal energy and thus some thermal motion, known as Brownian motion [76]. The equipartition theorem states the average energy of each atom is $\frac{1}{2}k_BT$ per degree of freedom, where k_B is Boltzmann's constant and T is the temperature in Kelvin [30]. This thermal motion couples into excitation of resonant modes of the detector system which contributes to noise in the detector. The test mass mirror's reflective coating material and substrate material suffer coating Brownian and substrate Brownian noise that can be reduced by the use of novel low mechanical loss materials [77] [78]. Mechanical loss is discussed in detail in section 2.2.2 of chapter 2 but for this chapter it can be simply described as internal frictional damping of a material. Thermal noise noise is an important noise source which can dominate from 10s to 100s of Hz.

Materials experience statistical temperature fluctuations throughout their volume. This temperature variation can lead to fluctuations in the refractive index of the material. This can affect both the test mass mirror and beam splitter materials, as well as the reflective coating materials on the mirrors, leading to what is called thermorefractive noise. While not currently a limiting noise source, as detectors become more sensitive, thermorefractive noise will become more of a significant noise source [79].

Thermoelastic noise is another form of thermal noise that results from finite temperature fluctuations in a material. These variations lead to small mechanical motions related to the coefficient of thermal expansion of the material leading to noise in the detector. This is discussed in more detail in section 2.3.3.2 of chapter 2.

Thermal noise mitigation through the use of low mechanical loss materials is discussed in detail in the context of current detectors in section 1.5.4. The work in this thesis is primarily focused on suspension thermal noise, discussed in detail in chapter 2, arising from thermal motion and excitation of the detector suspension system that hangs the test mass mirrors. This thesis looks at further thermal noise mitigation for future detectors by reducing the temperature the detector operates at, hence reducing the thermal energy, primarily to minimise Brownian and thermoelastic noise.

1.5 Current detectors and detections

This section discusses the current interferometric gravitational wave detectors around the world as well as their historic detections and contributions to physics. This section also highlights a novel suspension isolation system used for the test mass mirrors that enables these detectors to reach the strain sensitivities required for detection.

1.5.1 2nd generation detectors

Figure 1.11: LIGO Hanford, Washington (left), LIGO Livingston, Louisiana (centre) and Virgo, near Pisa (right), gravitational wave detectors [80] [81].

The Laser Interferometer Gravitational-Wave Observatory (LIGO) and Virgo interferometer are probably the most well known gravitational wave detectors. These detectors are classed as "2nd generation" gravitational wave detectors. LIGO is officially referred to as Advanced LIGO (aLIGO) due to its upgraded status from Initial LIGO and Enhanced LIGO, which are classed a 1st generation detectors [82] [83] [84]. Likewise Virgo is now referred to as Advanced Virgo (AdV) due to similar upgrades [85] [86].

All three 2nd generation detectors work on the same principal interferometer design. One can see each of the three detectors in figure 1.11. These images neatly illustrate the L-shaped interferometer design with the Fabry-Pérot arms extending away from the laser source, beam splitter and detection photodiode that are housed in the corner control buildings. Both LIGO detectors have interferometer arms lengths of 4 km whereas Virgo has interferometer arms lengths of 3 km. Both detectors' arms are housed within ultra-high vacuum (UHV) systems, with LIGO having the second largest UHV system in the world.

Many of these upgrades to the detectors were, and are still, tested in the Gravitational European Observatory 600 (GEO600) detector in Ruthe, Germany. While this detector is technically operational, with an arm length of only 600 m, it is not nearly as sensitive as 2nd generation detectors [87]. However, this detector has been invaluable in prototyping the technologies that have allowed 2nd generation to reach the sensitivities at which they could finally detect gravitational waves [88].

There are plans to build a network of more 2nd generation and future generation detectors with increasingly advanced techniques to improve sensitivity, improve sky localisations and better determine information about the sources. One such example is LIGO India [89] [90].

In order for 2nd generation detectors to work, a number of noise mitigations, discussed in section 1.4, are implemented.

Seismic isolation in the form of both passive and active isolation was implemented into aLIGO and AdV [66] [91]. The passive isolation is in the form of the suspension systems where aLIGO's is specifically discussed in detail in section 1.5.4. aLIGO's active isolation system consists of sensors and actuators that seek to keep all large vacuum chambers and optical

benches as still as possible [66]. AdV uses a slightly different system of passive and active suspension isolation called a superattenuator [92].

As discussed in section 1.4.2, gravitational gradient noise is not currently mitigated in either detector although it is not currently a dominant noise source [86] [93]. Novel techniques are being explored to allow for detectors becoming more sensitive at lower frequencies where this noise may start to limit the detector [68] [69].

Photon shot noise is reduced with the use of high power lasers. For example, in the aLIGO detectors, the input laser power is 125 W. Since aLIGO uses dual recycling mirrors, this power is increased to 5.2 kW by the PRM. This power is then amplified within the Fabry-Pérot arm cavities resulting in a power of 750 kW [84]. This increase in laser power will increase radiation pressure noise, however this is mitigated by the use of heavy 40 kg test masses. The laser wavelength is 1064 nm [94].

Further efforts to reduce the quantum noise saw the employment of optical squeezing in aLIGO for the last observing run. Squeezing can reduce the quantum noise in the amplitude or phase quadrature of the laser at the expense of the other. If one wishes to increase the sensitivity of the phase of the laser light, one can inject "squeezed" vacuum states into the dark port of the detector. This then reduces the phase quadrature, at the expense of increasing the amplitude quadrature [95]. For the last observing run, the type of optical squeezing was frequency-independent [96].

Mirror thermal noise is reduced by careful selection of the test mass aspect ratio, chosen to ensure resonant modes of the test mass mirror are sufficiently high in frequency (>10 kHz). Due to the low mechanical loss of the test mass mirror substrate, these high frequency modes store the majority of the Brownian substrate thermal energy and hence thermal noise, well above the frequency band of interest for detections. For 2nd generation detectors, the mirror substrate thermal noise is significantly lower than the coating thermal noise [63].

Coating thermal noise arises for the multilayer dielectric coating stacks used on the mirrors to reflect the interferometer beam. Increasing the diameter of the interferometer beam can reduce the noise by averaging this contribution over a larger area. Furthermore, careful selection of substrate and coating material combinations can be used to optimise thermal noise through low absorption and low mechanical loss materials [78].

The final method specifically for reducing suspension thermal noise is detailed below in section 1.5.4 and forms the basis for much of the work in this thesis.

1.5.2 LIGO and the first direct detection

The direct detection of gravitational waves by LIGO would not have been possible were it not for the employment of noise source mitigation measures and subsequent upgrades to the detectors. Two detectors make up the LIGO detector, one in Hanford, Washington, the LIGO Hanford Observatory and one in Livingston, Louisiana, the LIGO Livingston Observatory. In 2015, 100

years after Einstein's General Theory of Relativity, the LIGO detectors and the LIGO Scientific Collaboration (LSC) of 1200+ scientists made the first direct detection of gravitational waves, an event known as GW150914 [97].

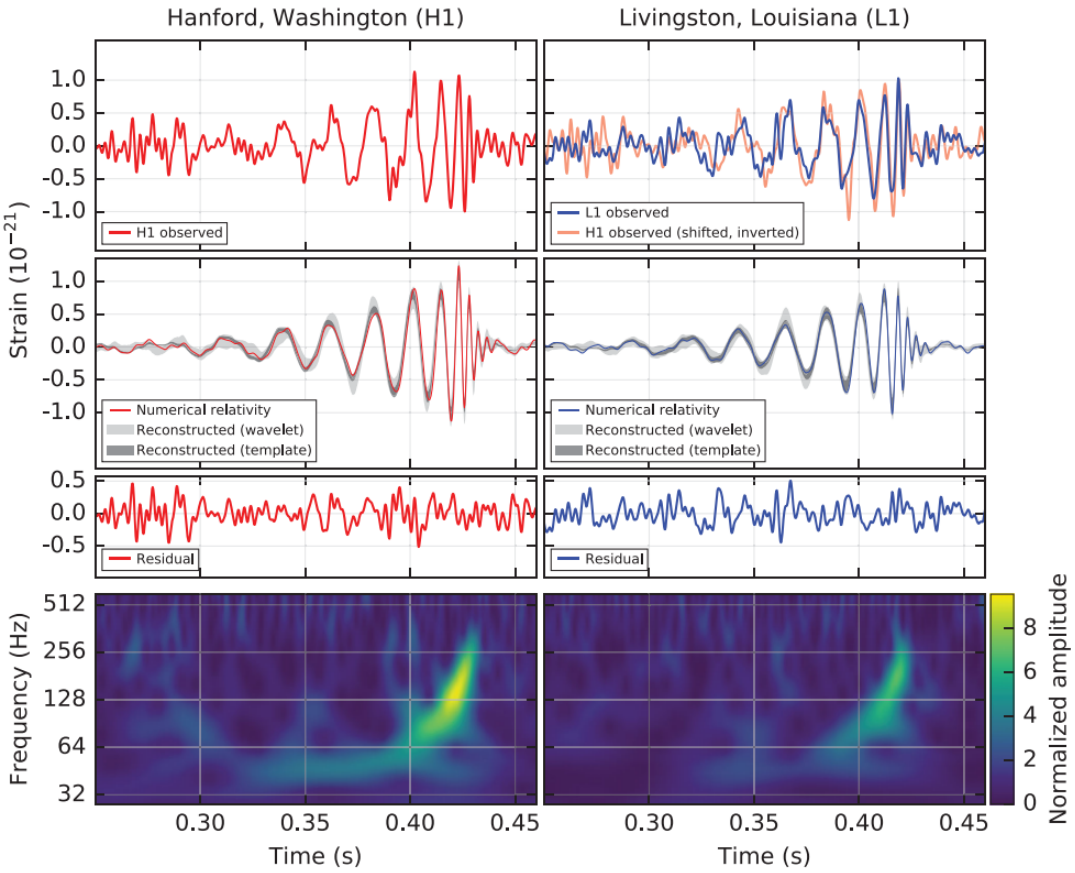

Figure 1.12: GW150914 gravitational wave signal observed in the LIGO detectors. The first row shows the processed signal in each detector. The second row shows the gravitational-wave strain projected onto each detector in the 35–350 Hz band. Solid lines show a numerical relativity waveform for a system with parameters consistent with those recovered from GW150914. The third row shows the residual between the first and second rows. The fourth row shows a time-frequency representation of the strain data, showing the characteristic signal frequency and amplitude increase over time as part of the inspiral phase [97].

GW150914 was produced by a BBH system comprising of two black holes of 36 and 29 solar masses originating approximately 1.3 billion light years away. The resultant merger formed a 62 solar mass black hole, radiating 3 solar masses worth of energy in the form of gravitational radiation. This spectacular event was over 50 times more energetic than the combined luminosity of *every* star in the observable universe. This event opened a new window into the universe, allowing scientists to probe our universe for the first time using gravitational observations.

By the time these gravitational waves reached Earth, 1.3 billion years later, they produced a peak strain amplitude in the detector of only 1×10^{-21}. The LIGO detector had successfully detected a length change in the interferometer arms equivalent to detecting a change in the

distance between the solar system and the nearest star, 4 light years away, to the accuracy of the thickness of a human hair. This discovery led to the 2017 Nobel Prize in Physics for LIGO members Rainer Weiss, Barry Barish and Kip Thorne [98] and the Special Breakthrough Prize in Fundamental Physics for the LIGO Scientific Collaboration members [99].

The first direction detection signal is shown in figure 1.12. As can be seen in the bottom row of figure 1.12, the characteristic increase in frequency and amplitude as the black holes inspiral towards one another is observed before merging into one black hole with a short ringdown and ceasing to produce gravitational waves due to the removal of the asymmetry of mass.

At the time of writing, there have been 90 confident detections of gravitational waves made through 3 observing runs (O1-O3) [100].

Two separate detectors form LIGO due to the same principle Weber enacted for resonant bar detectors where a second detector is located sufficiently far away, in this case $\approx$1000 km, so that any coincident signal observed in the detectors is not due to local noise but of astrophysical origin. Furthermore, two detectors allow for sky localisation of the gravitational wave source. However, for only two detectors this would typically be limited to a ring pattern across the entire sky although, due to the detectors' directional sensitivity, this can determine areas of probability of the source within this ring [101]. Having more detectors located around the world will improve sky localisation and parameter information of gravitational wave sources [102].

1.5.3 Virgo

The Virgo gravitational wave detector, located near Pisa, Italy, joined observations with the LIGO detectors in May 2017 [103].

A notable detection by the LIGO-Virgo detectors was discovered during O2, called GW170817, shown in figure 1.13. This event was equally special as it was the first detection of a BNS system, located approximately 130 million light years away [104]. Furthermore, because it involved the collision of two neutron stars, this event should create a kilonova, visible to the rest of the astronomical community using telescopes that observe across the more familiar electromagnetic spectrum.

Having the third Virgo detector allowed for greater sky localisation. An example of this is shown on the left hand side of figure 1.13 where one can see the light green source estimation areas arising from the two LIGO detectors. Including Virgo's observational data, this source could be better localised, as shown by the dark green area in figure 1.13. This gave astronomers a much smaller (but still significantly large) area of the sky to search for the electromagnetic counterpart to the gravitational wave event.

Indeed, astronomers across the world observed a kilonova emanating from the region predicted by the 3 gravitational wave detectors. The Fermi and INTEGRAL gamma ray observatories, which can do all-sky and wide-sky observations of gamma rays detected a burst of gamma rays from the localised region 1.7 seconds after the gravitational wave event [105]. This short

delay is attributed to complex astrophysical effects but fundamentally confirmed Einstein's prediction that gravitational waves travel at the speed of light [106]. This detection gave birth to the era of multi-messenger astronomy where the scientific community can now observe astrophysical processes in both gravitational and electromagnetic forms [105].

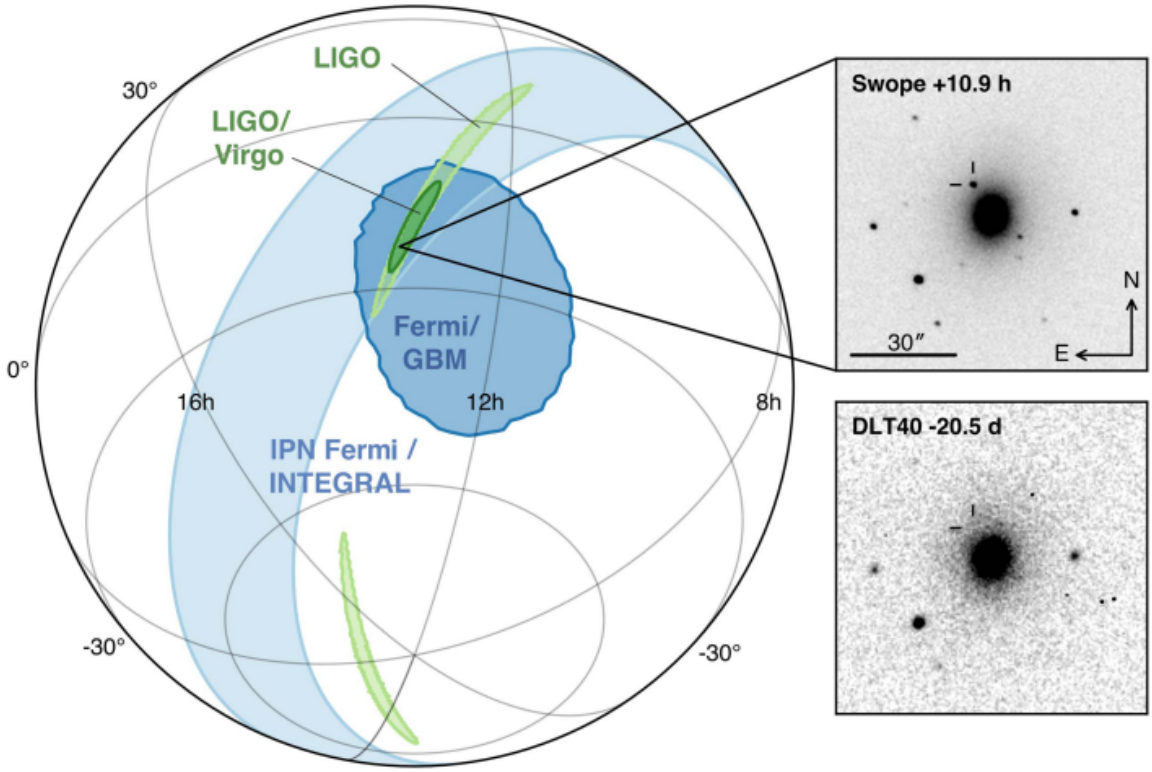

Figure 1.13: The left hand image shows localisation of the GW170817 gravitational wave signal with corresponding corroboration from the Fermi and INTEGRAL gamma ray observational satellites. The right hand images show the appearance of an electromagnetic counterpart (the kilonova) from the SWOPE observatory 10.9 hours after, where the DTL40 observatory shows no electromagnetic counterpart 20.5 days prior to the event signal [105].

This was the first time any source had been observed in gravitational waves and electromagnetic radiation with detections ranging over X-ray, ultraviolet, optical, infrared, and radio wavelengths with the peak emission shifting through the variety of wavelength bands over the course of hours, days and weeks. The electromagnetic observations allow for greater interpretation of the nature of the binary neutron star system and offer a comprehensive and sequential description of the physical processes related to the merger of a such a system [105].

1.5.4 aLIGO suspension system and future upgrades

In order to keep the test mass mirrors as still as possible, or reduce displacement noise, one requires superior seismic and thermal isolation.

The development of novel fused silica suspension systems for aLIGO at the University of Glasgow allowed for the direct detection of gravitational waves [107].

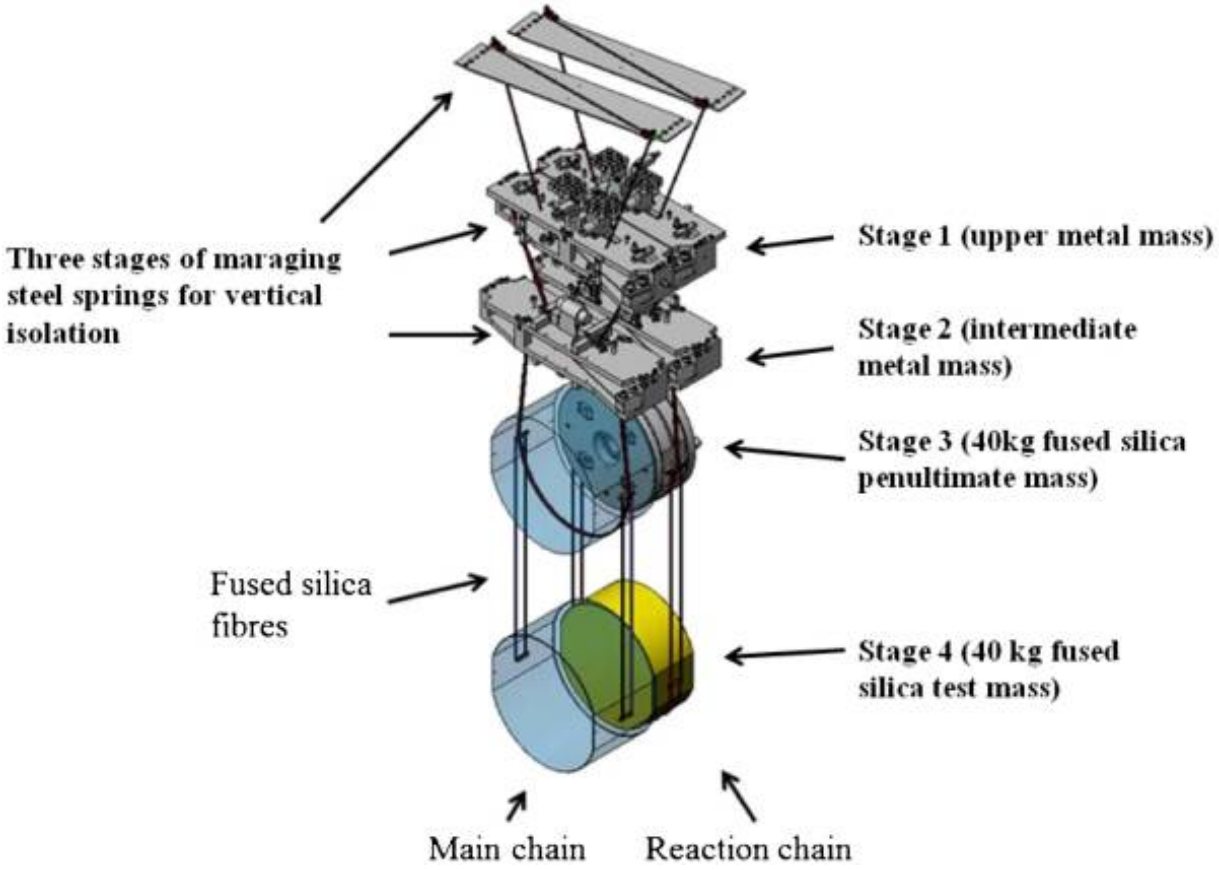

Figure 1.14: Schematic of the aLIGO quadruple pendulum suspension with 40 kg test mass [63].

A drawing of the aLIGO quadruple stage pendulum suspension that allowed for the direct detection of gravitational waves is shown in figure 1.14. Many of the methodologies discussed in this suspension system were first pioneered in the GEO600 detector [65].

A parallel suspension system, called the reaction chain, sits behind the quadruple pendulum system. The top 3 pendulums contain a series of coils and magnets while the bottom stage contains an electrostatic drive for the purposes of low-noise suspension control and interferometer alignment.

The top stage of this entire system is affixed into large active isolation chambers discussed previously in section 1.5.1. As shown in the figure 1.14, there are 3 stages of maraging steel springs designed to improve vertical isolation which can couple into the detector's horizontal motion, as discussed in more detail in section 2.4.3. The four pendulum stages each provide passive seismic isolation, as discussed in section 1.4.1. The final stage suspension, which consists of the penultimate mass, test mass and fused silica fibres provide the remaining significant noise reduction to allow the detector to reach the necessary sensitivity. With a typical seismic noise level of $10^{-9}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$, the suspensions system achieve an extremely impressive seismic

and thermal noise reduction of order 10 billion in order to achieve the target displacement noise sensitivity for the test mass mirror of $10^{-19}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ [108].

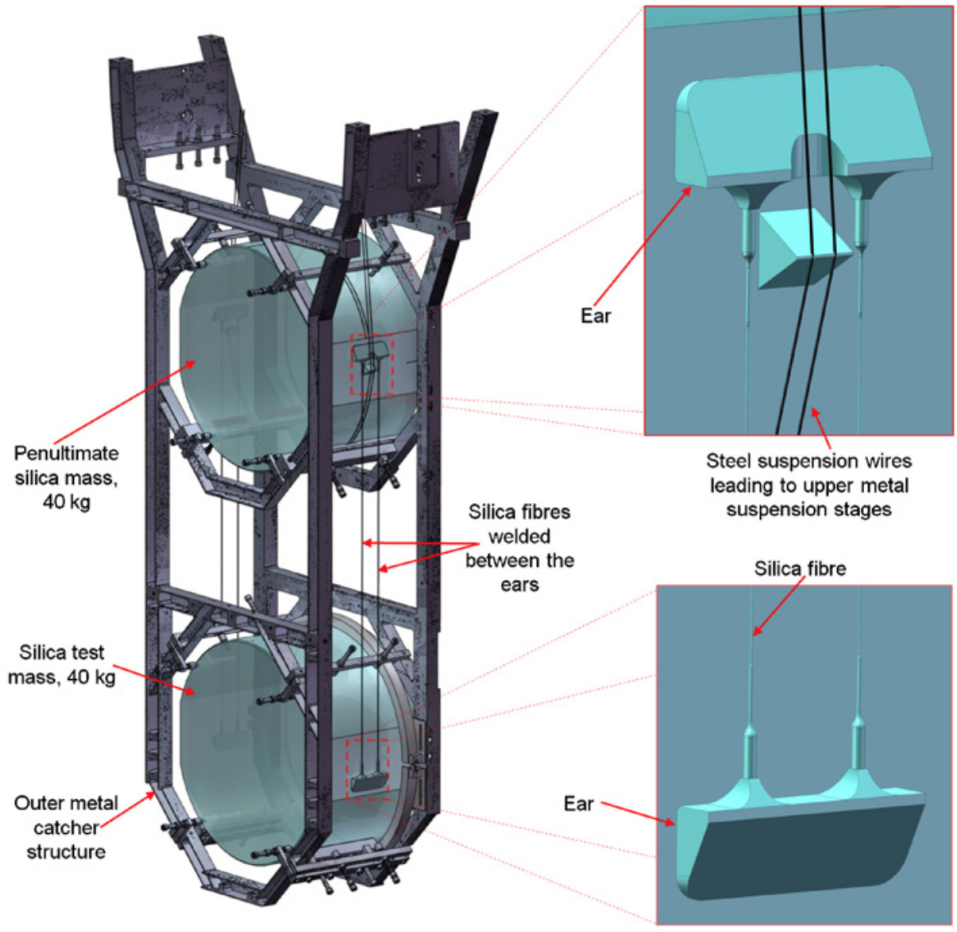

Figure 1.15: The aLIGO final stage suspension consisting of two fused silica 40 kg masses. The coating to reflect the interferometer beam is not shown but is present on the bottom test mass. Insets show details of fused silica attachment ears and suspension fibres. Image modified from [77].

Figure 1.15 shows the final stage suspension of aLIGO. The suspension consists entirely of fused silica due to its superior strength and thermal noise properties, specifically its low mechanical loss. The top and bottom masses are 40 kg with flats on either side to allow for the attachment of "ears". These ears are fused silica prisms that are chemically bonded using hydroxy-catalysis (HC) bonding. These ears allow for the welding of the fused silica fibres which are used to suspend the bottom test mass. The bonded and welded elements produce what is known as a quasi-monolithic suspension due to the final stage suspension effectively being made entirely of fused silica.

The fused silica fibres are $\approx$400 µm in diameter at their thinnest section and 60 cm long. It is important to note here, the fused silica fibres are not uniform along their length, and are specifically shaped at either end to become thicker. While fused silica has superior mechanical

loss at room temperature compared to other materials, the HC bonds are a factor of $\approx 10^7$ worse in mechanical loss [109] [110] [111]. By thickening the fibres at either end and welding the fibres to the ears, one confines the bond only to the ear and mass interface, thus keeping energy in the suspension away from these high loss bond regions. This highlights the importance of minimising areas of high noise, by keeping energy in the suspension system away from areas of high mechanical loss, in order to realise the required detector sensitivity.

To improve non-suspension aspects of aLIGO further, upgrades, called A+, involve improved coating mechanical loss, larger beam size diameter with a correspondingly larger optic, and the introduction of frequency-dependent squeezing [112]. Frequency-dependent squeezing will allow one to select the trade off between phase or amplitude noise at high or low frequencies, permitting simultaneous minimisation of shot noise and radiation pressure noise respectively. Frequency-dependent squeezing has not yet been demonstrated at gravitational wave interferometer laser frequencies but work is ongoing to enable this [113] [114].

Further to A+ upgrades, a set of upgrades are proposed to be implemented post observing run 5 (O5+). Specific suspension upgrades for O5+ will utilise heavier 100 kg test masses, running at fibre tensions of 1.6 GPa and longer 100 cm suspensions with reduced surface loss and lower mechanical loss components [115] [116].

1.5.5 KAGRA

The Kamioka Gravitational-Wave (KAGRA) detector is similar in design to the 2nd generation detectors, however it is unique in that it operates at cryogenic temperatures, reducing thermal noise, and is built underground, reducing seismic and Newtonian noise [117]. For these reasons KAGRA is sometimes regarded as a 2.5 generation detector and is an asset for developing the necessary cryogenic payloads, built underground, for 3rd generation (3G) detectors [118].

As shown in the illustration in figure 1.16(a), the KAGRA detector is situated in Mount Ikenoyama, as part of the Kamioka Observatory, in Japan. The interferometer has arms lengths of 3 km [118].

As shown in figure 1.16(b), the "type A" suspensions in KAGRA suspend the test mass mirrors, weighing 23 kg, as part of a long suspension chain. As shall be discussed in detail in chapter 2, fused silica is not suitable for operation at cryogenic temperatures and so the KAGRA Collaboration decided to use sapphire for the final stage suspension test mass mirror and fibres. The test mass mirror is suspended from 4 sapphire fibres, 1.6 mm diameter and 350 mm long. The sapphire fibres are attached to "nail heads" at either end, using an adhesive called SUMICERAM [119]. These nail heads then slide into ear-like prisms on the side of the test mass. These ear-like prisms are HC bonded onto the sapphire test mass sides. There is no penultimate mass in the KAGRA suspensions and so the nail heads at the top of the fibre are affixed into sapphire blade springs with slots. To keep the components together, the nail head-slot interfaces are contacted using gallium bonding with ultrasonic soldering techniques at the

top and bottom of the suspension [119].

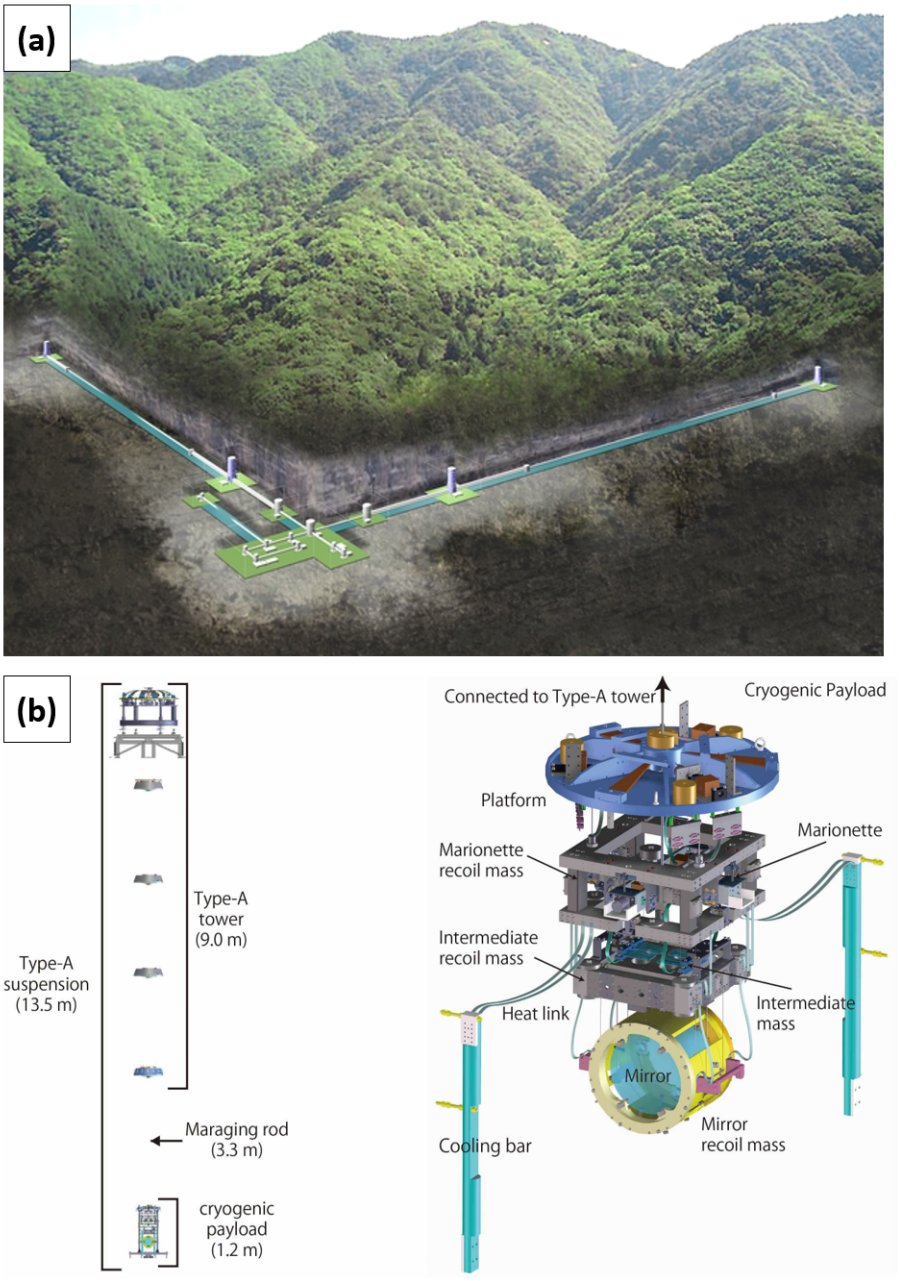

Figure 1.16: (a) Illustration of the KAGRA detector located underground with 3 km arms [120]. (b) KAGRA cryogenic payload "type A" suspension located at the end of each interferometer arm [119].

Cooling these KAGRA test mass mirrors is through both radiation and conduction in the respective high and low temperature regimes. Pulsed-tube cryocoolers are used to maintain outer and inner radiations shields at 80 K and 8 K respectively. Cooling bars, shown in figure 1.16(b), provide conductive cooling through ultra-pure aluminium heat links to the marionette and intermediate mass above the test mass. At low temperatures, final cooling of the sapphire

test mass is purely through conduction of a sapphire fibres suspending the test mass, bringing the final operational temperature of the test mass to 20 K [119].

The KAGRA detector intended to join O3 alongside aLIGO and AdV however, due to the COVID-19 pandemic, these detectors had to cut the O3 observing run short. Instead, the KAGRA and GEO600 detectors jointly observed for a period of two weeks in April 2020. At the time of observation, the KAGRA detector had a similar sensitivity as the GEO600 detector meaning the detection of gravitational waves was unlikely unless the sources were particularly strong and nearby. KAGRA was limited at low frequencies (<100 Hz) by local control noise of the mirror suspensions. It was also limited at higher frequencies (>400 Hz) by shot noise. At intermediate frequencies the noise is not well-modelled but shows some coherence with environmental acoustic noise, which may arise from scattered light coupling [121].

This highlights the difficulty in the implementation of cryogenic detectors, which will be discussed further in section 1.6. Nonetheless, the KAGRA Collaboration have achieved impressive engineering and scientific feats by building the detector underground as well as successfully operating a cryogenically cooled interferometer. Future upgrades, still to be specified, are planned to meet design sensitivity allowing KAGRA to fully compliment the current gravitational wave detector network [102].

1.6 3rd generation detectors

It is believed the current generation of detectors will reach the thermal and quantum limits of their designs within a few years [122], with the LIGO A+ upgrade expected to enhance sensitivity by up to 50 %. 3G detectors seek to be able to detect gravitational waves emanating from the very edge of our observable universe by creating more sensitive detectors, with improvements of at least half to one order of magnitude better than aLIGO and expanding the frequency range which can be observed with good sensitivity [122] [123].

The current plans for 3G detectors can broadly be split into two categories, upgrades to the current LIGO facilities, namely LIGO Voyager and Cosmic Explorer (CE) and entirely new detectors, the Einstein Telescope (ET).

1.6.1 LIGO Voyager and Cosmic Explorer

The proposal for LIGO Voyager should yield up to half an order of magnitude improvement over aLIGO by tackling several noise sources; quantum noise, mirror thermal noise, suspension thermal noise and Newtonian noise [122].

The primary quantum and mirror thermal noises sources in this design will be addressed by a material change of the test mass mirrors and final stages from fused silica to silicon, the justification for which is discussed in detail in chapter 2. Quantum noise reduction arises from

an increase in optical power stored in the arms permitted by higher thermally conductive silicon. Lowering the mirror test mass temperature will address thermal noise alongside better silicon coatings to reduce optical absorption in the test mass. Newtonian noise will be improved by using seismometer arrays in conjunction with adaptive noise regression [122].

Suspension thermal noise is mitigated through an overall reduction in temperature of the final stage suspension. This method is addressed within this thesis and discussed in detail in the following chapters.

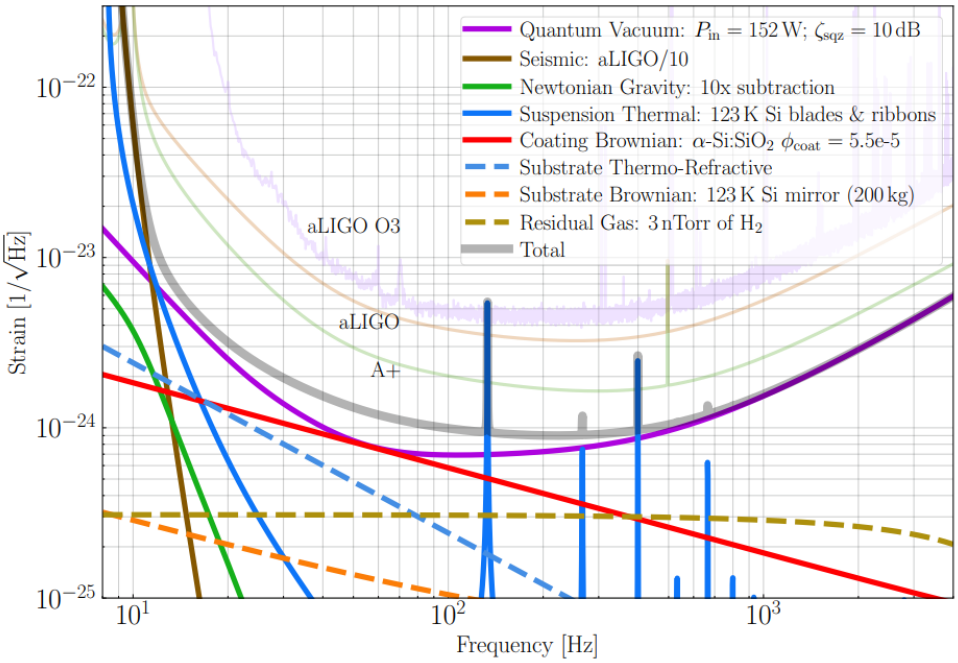

Figure 1.17: LIGO Voyager noise curve compared to aLIGO's measured noise curve during O3, aLIGO's design sensitivity curve and A+ upgrade design goal [122].

The overall LIGO Voyager design proposes a Fabry-Pérot dual-recycled Michelson interferometer similar to aLIGO with the additional upgrades of silicon suspensions supporting 200 kg silicon test masses with amorphous silicon coatings, all cooled to 123 K. This temperature regime relies on a combination of conductive and radiative cooling. The change in test mass material requires a change to a longer laser wavelength, of 1550 nm or 2000 nm, which will reduce optical scattering losses from the mirrors. Optical squeezing, combined with these heavier test masses and higher power stored in the arm cavities will reduce both aspects of quantum noise. The benefits of these upgrades translate to the LIGO Voyager noise curve shown in figure 1.17.

Much of the work of LIGO Voyager will feed in to a separate detector, Cosmic Explorer (CE) [122] [123] [124]. The CE detector aims to push the reach of gravitational wave astronomy to the edge of the observable universe, as shown in figure 1.18. The CE detector evoluton can be split into two stages, CE1 and CE2.

The CE1 detector aims to scale up aLIGO technologies with interferometer arms closer in

length to the wavelengths targeted by ground-based detectors. This requires up to 40 km long arm cavities which would mean tunnelling underground in order to create a straight beam-line through the ground with the end stations remaining on the surface. CE1 upgrades will included much heavier fused silica suspensions, up to 320 kg, with lengths up to 1.2 m. Upgraded fused silica suspensions are currently being developed in Glasgow [125]. Circulating laser power will also be increased up to 1.4 MW, utilising optical squeezing techniques [123].

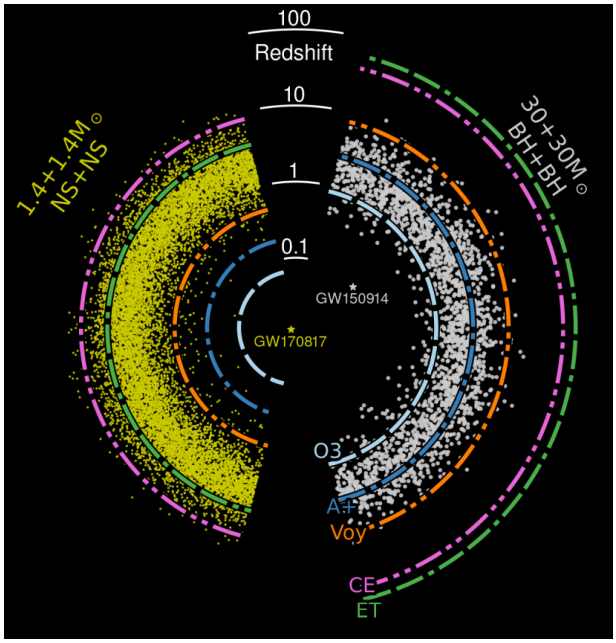

Figure 1.18: Astrophysical horizon of current and proposed detectors for compact binary systems. The yellow dots represent a simulated population of binary neutron star mergers and the white dots represent binary black hole mergers [123].

The CE2 detector upgrade will replace the test mass mirrors with 320 kg silicon optics, 1.2 m in length run at a proposed cryogenic temperature of 123 K. As with voyager, radiative and conductive cooling are required for operating in this temperature regime. The laser power would also be increased further, up to 2 MW [123].

1.6.2 Einstein Telescope

As shown in figure 1.19(a), the Einstein Telescope (ET) will be an entirely new detector that aims to achieve an order of magnitude better sensitivity than current 2nd generation detectors, with a sensitivity ranging from as low as 1 Hz up to 10 kHz. The detector will have interferometer arm lengths of 10 km built entirely underground to reduce seismic and Newtonian noise [126] [127].

As shown in figure 1.19(b), the proposed interferometer configuration consists of 3 separate interferometers, overlapping and arranged to form an equilateral triangle. This configuration provides improved sensitivity and information over a range of gravitational wave polarisations. It also provides redundancy by being able to operate with only 2 of the 3 interferometers active.

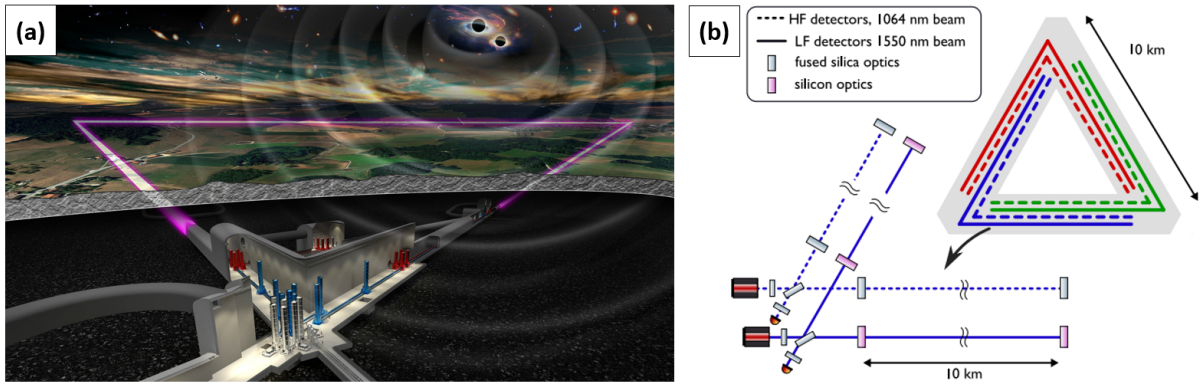

Figure 1.19: (a) Artist's impression of the Einstein Telescope, built underground [128]. (b) General overview of the Einstein Telescope interferometer layout [129].

Further to this, the interferometers will be split into two separate detection bands, tackling issues that arise in single broadband detectors. For higher frequency detection, one requires higher optical powers to reduce the quantum noise, however this makes cooling of test masses to cryogenic temperatures, by reducing thermal noise for low frequency detection, extremely difficult. The ET detector design proposes a high power, room temperature interferometer, referred to as ET-high frequency (HF). In parallel, for low frequency detection is a cryogenic interferometer, referred to as ET-low frequency (LF) [127]. This technique is sometimes referred to as xylophone detection. Both ET-HF and ET-LF will take advantage of frequency dependent optical squeezing.

ET-HF will operate at room temperature with a circulating laser power of 3 MW in the arm cavities and a laser wavelength of 1064 nm. The fused silica mirror diameter will be 62 cm with a mass of 200 kg. The fused silica suspension will consist of fibres 620 µm diameter at their thinnest points, with a length of 1.5 m [127], with such fibres having been already demonstrated in Glasgow [125].

ET-LF will operate at either the low temperature, conductive-only, 10–20 K regime or the high temperature, conductive and radiative cooling, 120 K regime, with a circulating laser power of 18 kW in the arm cavities and a proposed laser wavelength of 1550 nm. The crystalline silicon mirror diameter will be 45 cm with a mass of 211 kg. The crystalline silicon suspension will consist of fibres 4.4 mm diameter, with a length of 1.5 m [127].

The ET-LF suspension specifically raises a number of challenges to be overcome, primarily the development of fabrication techniques to produce long, thin silicon fibres or ribbons that are strong enough to support the 211 kg test mass. Work is conducted in chapter 3 of this thesis to investigate this issue. Furthermore, small-scale prototype suspensions must be designed, built and cooled in order to investigate feasibility and this work is covered in chapter 4. Finally, techniques to joint suspensions together without contributing to a significant suspension thermal noise increase must be investigated and the work and chapters 4 and 5 explores this [127].

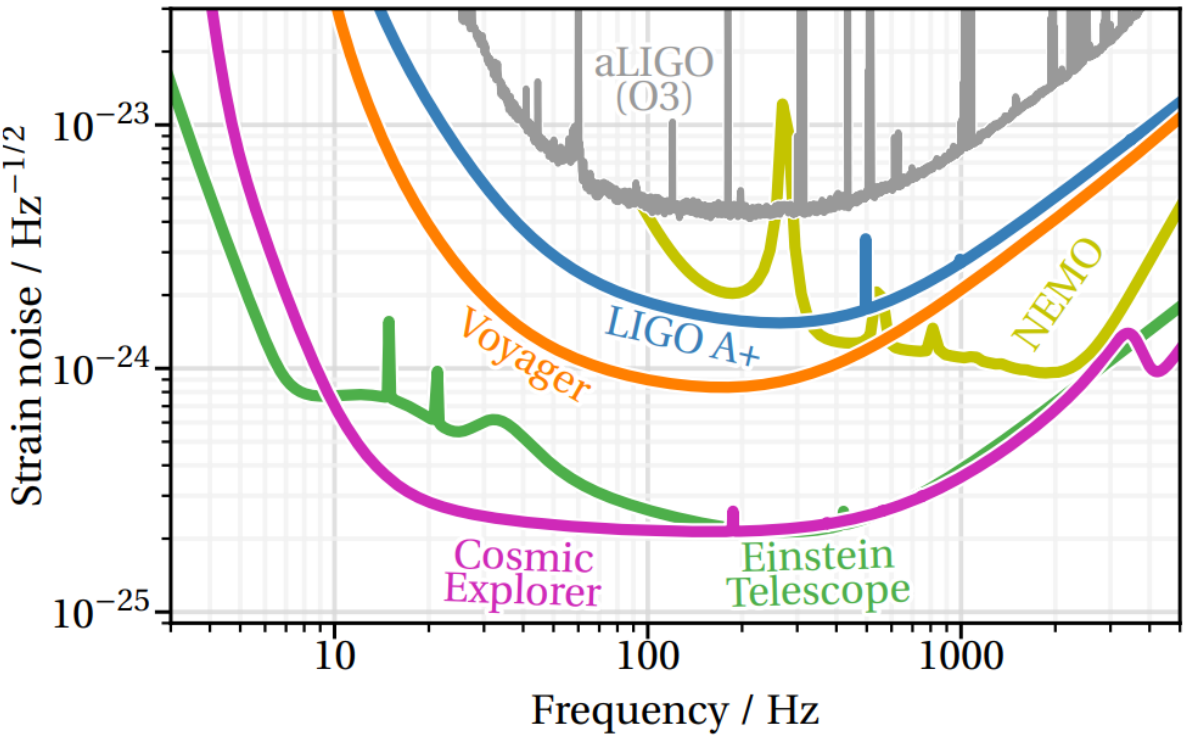

Figure 1.20: Noise curves comparing aLIGO's measured noise curve during O3 to the A+ upgrade design goal, LIGO Voyager, Cosmic Explorer and the Einstein Telescope [123]. The Neutron Star Extreme Matter Observatory (NEMO) detector is also listed here as a proposed high-frequency (>1 kHz gravitational wave detector based in Australia [130].

Figure 1.20 compares all of the above gravitational wave detector noise curves, highlighting the significant improvements expected with the advent of 3G detectors. With multiple research efforts working in parallel all over the world, the realisation of 3G detectors should allow for the detection of gravitational waves emanating from the very edge of the observable universe.

1.7 Laser Interferometer Space Antenna

Figure 1.20 highlights the fundamental limit experienced by all ground-based detectors where the strain sensitivity decreases significantly below 1 Hz. This low frequency limit arises from unavoidable terrestrial seismic and Newtonian noise on Earth [67].

Many interesting astrophysical sources of gravitational waves at frequencies <1 Hz are thought to exist. Such sources include the first seed black holes, supermassive black holes and extreme mass-ratio inspirals as well as more speculative signals from proposed cosmic strings and primordial gravitational waves generated by cosmological inflation [131] [132] [133] [134].

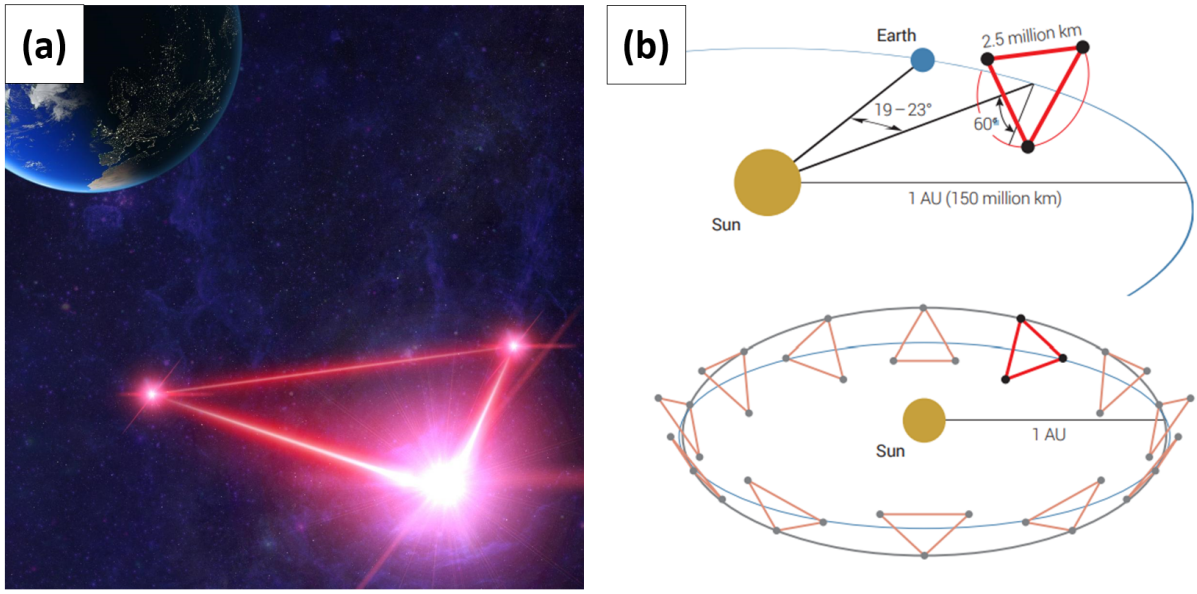

Figure 1.21: (a) Artist's impression of the LISA space-based gravitational wave detector [135]. (b) Proposed orbit of LISA [131].

To probe these low frequency sources, the space-based Laser Interferometer Space Antenna (LISA) detector was conceived. This detector is based on laser interferometry between free flying test masses inside drag-free spacecraft, as shown in figure 1.21(a). Three spacecraft will orbit in a triangular configuration producing laser interferometry separated by a distanced of 2.5 million km. These spacecraft will maintain this separation as the orbit in Earth's Lagrangian point L3, as shown in figure 1.21(b). This design will allow for all-sky monitoring of sources from 10^{-4}–10^{-1} Hz [131].

1.8 Conclusion

This chapter has highlighted the fundamental astrophysical origins and mechanisms of how gravitational waves are generated, including the early attempts to detect them. Significant advances in physics and experimental techniques have allowed for the direct detection of gravitational waves through laser interferometry. This technique allows for detection over a wide frequency band, ensuring a plethora of astrophysical sources can be detected and better understood, particularly when used in conjunction with more traditional electromagnetic techniques, in what is termed multi-messenger astronomy.

The fundamental sensitivities of these detectors are limited by an array of physical noise sources that must be understood and mitigated for. This thesis focuses particularly on suspension thermal noise and its improvement by the use of novel low mechanical loss materials alongside direct thermal noise reduction through cryogenic cooling. The work in this thesis aims to provide the foundation for designing, assembling, characterising and commissioning the first cryogenic

single crystalline silicon suspensions for 3G gravitational wave detectors.

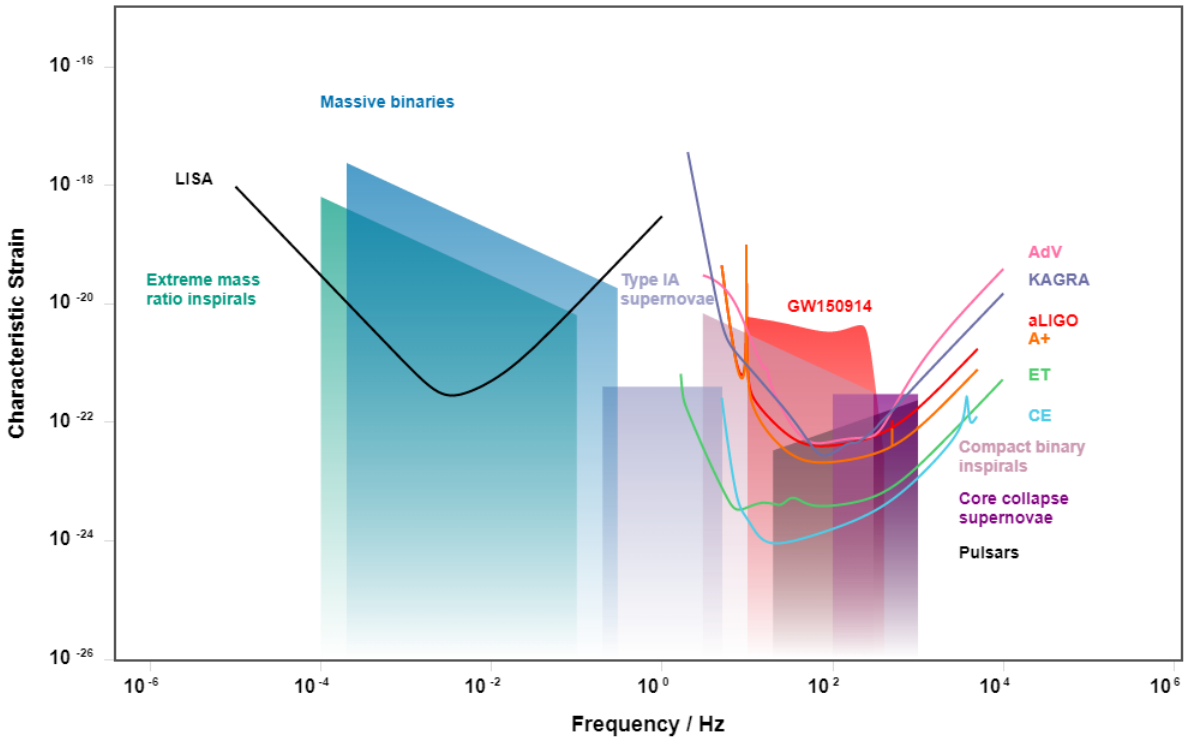

Figure 1.22: Summary of gravitational wave sources and the relative sensitivities of all detectors discussed in this chapter [16].

By increasing the number of detectors around the world, one can better localise the source of gravitational waves alongside attaining more information on the properties of the source itself. Furthermore, by building bespoke high or low frequency detectors one can further improve the detection sensitivity across the full bandwidth of interest. For low frequency detection, space-based detectors must be employed to overcome the <1 Hz wall of noise experienced by ground-based detectors. Figure 1.22 shows the extensive and diverse range of gravitational wave detectors in operation and in proposed design that will allow humanity to better understand the universe in which we live.

Chapter 2

Silicon gravitational wave detectors

2.1 Introduction

Current gravitational wave detectors utilise fused silica for the suspension elements and test mass mirrors, bonded together to create a quasi-monolithic suspension [64]. These suspensions afford excellent noise performance across the detector spectrum as shown in figure 2.1. However, noise from the thermal motion of the test masses and their suspension elements still limits detector performance within the $\approx$ 10–200 Hz region [136]. For suspension elements, this noise is particularly significant within the 10–30 Hz range [5]. This reduces the signal-to-noise ratio (SNR) of the detector.

This chapter discusses key factors in reducing the thermal noise for future 3G detectors and the motivation for moving towards cryogenic silicon suspension elements.

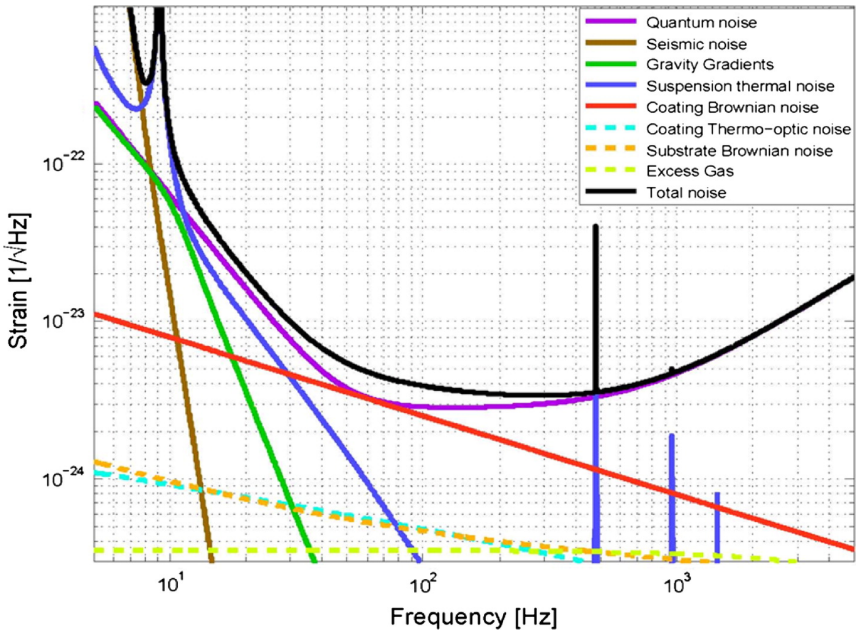

Figure 2.1: Noise contributions to the aLIGO detector in the high power, broadband mode of operation [63].

2.2 Thermal noise in gravitational wave detectors

Suspension thermal noise is encapsulated in the equation for the thermal displacement noise amplitude spectral density, $x_{rms}(\omega)$, derived from the Fluctuation-dissipation theorem [137] [138]:

$$x_{rms}(\omega) = \sqrt{\frac{4k_BT}{\omega m}\left(\frac{\omega_o^2\phi(\omega)}{\omega_o^4\phi^2(\omega)+(\omega_o^2-\omega^2)^2}\right)} \tag{2.1}$$

where k_B is the Boltzmann constant, T is the temperature, m is the pendulum mass, ω_o is the resonant angular frequency, ω is the angular frequency of interest and $\phi(\omega)$ is the combined mechanical loss angle of the final-stage suspension elements such as the attachment pieces, called ears, the fibres and the joints.

Since k_B is a constant and the resonant angular frequency, ω_o, is determined by suspension element geometries, this leaves 3 factors that can directly reduce suspension thermal noise; increasing the mass, reducing the mechanical loss and reducing the temperature. These factors are discussed individually in the following sections.

2.2.1 Increasing the mass

As shown in equation 2.1, the thermal displacement noise is proportional to $m^{-\frac{1}{2}}$. Therefore, increasing the mass will decrease the thermal displacement noise of the detector.

This is an approach already taken and demonstrated to reduce suspension thermal noise for previous and current fused silica room temperature suspensions. The GEO600 detector uses test masses that are 5.6 kg in mass [139]. Initial LIGO increased this mass to 10 kg [140], eventually increasing to 40 kg for aLIGO and AdV [64] [141]. Further increases in mass are planned as part of future aLIGO upgrades and for ET-HF suspensions [125].

Furthermore, this increase in mass will have the effect of reducing radiation pressure noise which is inversely proportional to the mass [142]. This will also lead to a larger optic diameter, allowing for larger interferometer beam size, reducing the coating noise [143].

The combination of factors arising from increasing the mass, shows the thermal displacement noise decreases, as shown in figure 2.2.

Two further improvement factors are proposed that look at increasing the stress in the fibres and changing the geometry of the suspension elements by lengthening the final stage suspension, in order to further reduce the thermal displacement noise. These factors are covered in chapters 3 and 5 respectively.

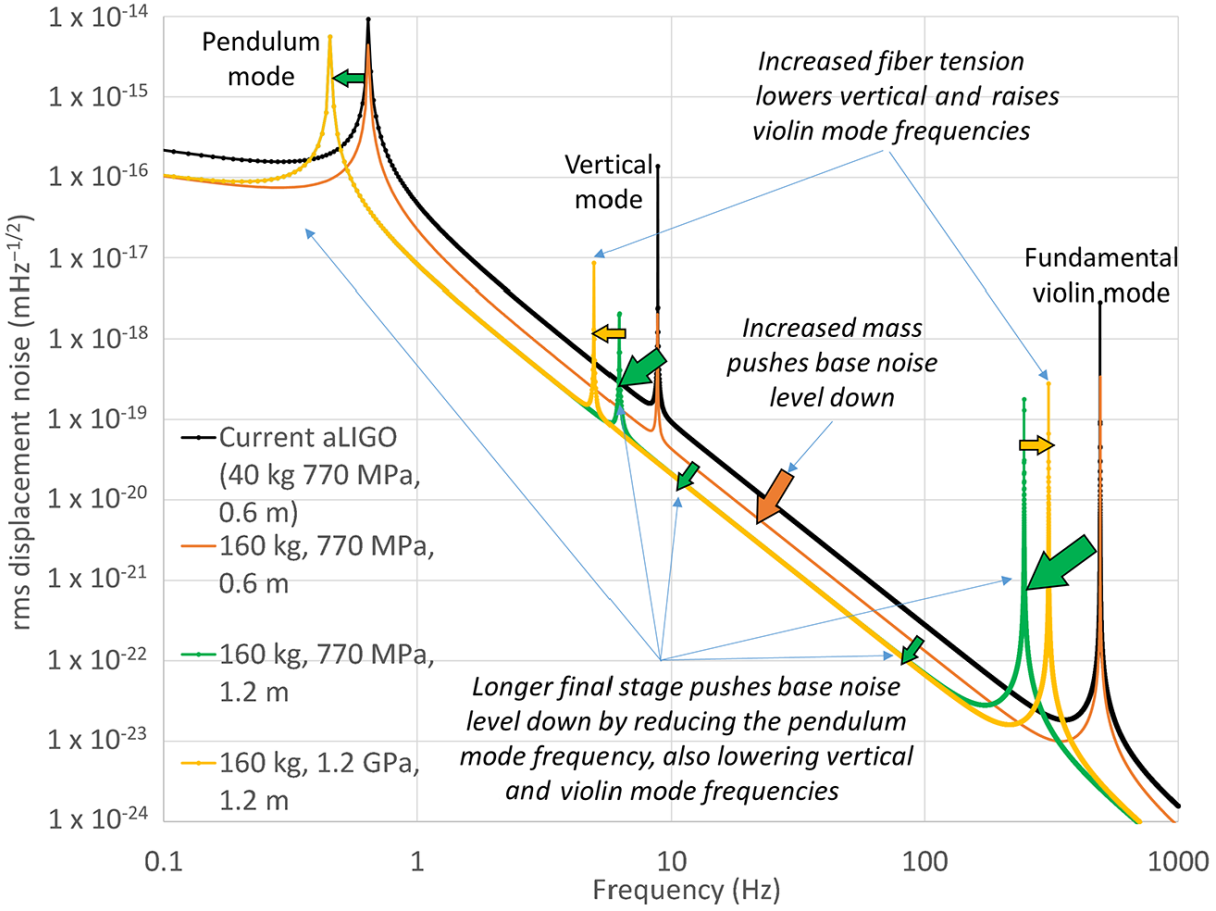

Figure 2.2: Calculation of thermal noise of a single mirror fused silica suspension showing the improved performance by increasing the mass, increasing the final-stage length and increasing the fibre's suspended stress [125].

2.2.2 Mechanical loss

As can be seen from equation 2.1, reducing mechanical loss, $\phi(\omega)$, is another method of reducing thermal displacement noise.

Mechanical loss is a dimensionless factor that is an inherent material property. If one were to apply a force to one end of a material, one would expect this force to transmit through the object over a finite time. This delay arises from the time it takes for the force to translate atom to atom, through anelastic spring-like atomic bonds, within the material as it deforms to move to a new equilibrium state. This is a source of dissipation known as mechanical loss, where $\phi(\omega)$ represents the phase angle at which the strain response of the material lags behind the driving force [144].

If one assumes no anelasticity within the system, a perfect elastic material can be described by Hooke's law:

$$F = -kx \tag{2.2}$$

where F is the restoring force, k is the spring constant and x is the linear distance the spring is displaced from equilibrium.

To account for anelasticity, Hooke's law can be written in its complex form to include the phase lag of the system:

$$F(\omega) = -k(1 + i\phi(\omega))x \tag{2.3}$$

where $\phi(\omega)$ is the mechanical loss of the material.

Initial LIGO used steel wires to suspend optics [82]. The steel wire had a mechanical loss value on the order of 1×10^{-4} [145]. GEO600 pioneered the advanced technology of using fused silica for the suspension material, which had a superior mechanical loss of 1×10^{-8} [146] [147]. An up-scaling of this technology was subsequently implemented in the LIGO detectors which have recently been demonstrated to have the lowest measured mechanical loss of any gravitational wave detector [148].

For extremely low mechanical loss suspension elements, such as fused silica, mechanical loss is experimentally difficult to measure at frequencies of interest for gravitational wave detectors since the dynamics of the system do not depend on the loss to a very high degree [149]. For a constant temperature, or constant amount of thermal energy, low loss materials integrate more of that energy into resonance peaks which reduces the overall baseline loss of the system off-resonance. This is clear to see from figure 2.2. These sharp-peaked resonances can then be moved around spatially in frequency with careful mechanical design to keep them out of detection band frequencies.

Instead, due to these sharp-peak resonances, it is easier to measure the quality factor, Q, at the resonant frequency, $Q(\omega_0)$ [149]. The quality factor can be obtained from measuring the decay of the amplitude of motion of an excited resonant mode of a material [150]. The quality factor is defined as [78] [151]:

$$Q(\omega_0) = 2\pi \frac{E_{stored}}{E_{dissipated}} = \frac{1}{\phi(\omega_0)} \tag{2.4}$$

where E_{stored} is the total energy stored in the oscillating system and $E_{dissipated}$ is the energy dissipated in each cycle of oscillation.

This method has been used extensively to obtain the mechanical loss value of materials for gravitational wave detectors [149] [152] [153] [154] [155].

2.2.3 Reducing the thermal energy

The GEO600, Virgo and LIGO detectors currently run at room temperature. A final way to reduce the thermal noise is to simply reduce the temperature, T, and hence thermal energy, within the final-stage suspension.

However, it is important to note that fused silica suspensions are not amenable to cryogenic operations for two primary reasons.

Firstly, fused silica exhibits a wide mechanical loss peak around 40 K, significantly increasing thermal displacement noise [156] [157].

Secondly, fused silica has poor thermal conductivity. Fused silica's room temperature thermal conductivity is already very low compared to that of metals, which are typically on the order of 100s of $\mathrm{W\,m^{-1}\,K}$ [158]. Worse still, this thermal conductivity decreases with decreasing temperature, making extraction of any deposited laser power in the cryogenic test mass mirrors potentially challenging [159].

The KAGRA detector, and its Cryogenic Laser Interferometer Observatory (CLIO) prototype, were the first experiments to look to exploit the reduction in temperature via cryogenic operation [119] [160]. This is also the proposal for the cryogenic suspensions of 3G detectors such as the Einstein Telescope (ET) [126], LIGO Voyager [122] and stage 2 of the Cosmic Explorer (CE) [123] detector.

Sapphire has been proposed as an alternative material for cryogenic detector operation in response to the need to move away from fused silica for 3G detectors [119]. KAGRA uses sapphire as the suspension elements and mirror material instead of fused silica [119].

In order to gain the benefits from sapphire, KAGRA must run at operating temperatures <20 K. As will be discussed in section 2.5.3, cooling large test masses to cryogenic temperatures is not trivial, and the lower the operational temperature of the detector, the greater the complexity of cooling required [161].

Silicon has also been proposed as an alternative material for cryogenic detector operation [122] [123] [126] [136]. Silicon encompasses many advantageous properties that are attractive to cryogenic suspensions. Most importantly, are its high thermal conductivity, low or zero coefficient of thermal expansion and low mechanical loss compared to fused silica. These properties and their advantages are explored in detail in the next section. Compared to sapphire, silicon is considered to have superior optical properties in the near-infrared region [161] and potentially benefits from higher average strength bonds compared to sapphire [162]. Like sapphire, silicon potentially has an issue in the availability of large pieces for manufacturing a heavy test mass [126] [127]. While sapphire is not covered in the work of this thesis, this chapter displays plots of relevant properties for comparative purposes.

2.3 Silicon as a detector material

This section covers silicon as a material for 3G detectors. It will cover the fundamental properties of silicon, highlighting the main advantages of silicon over other materials.

2.3.1 Elemental silicon production

Silicon, in its most abundant form, makes up around 30% of the Earth's crust as silicate igneous rock [163]. When purified into elemental silicon it takes the form of a shiny, dark grey solid semi-metal as shown in figure 2.3.

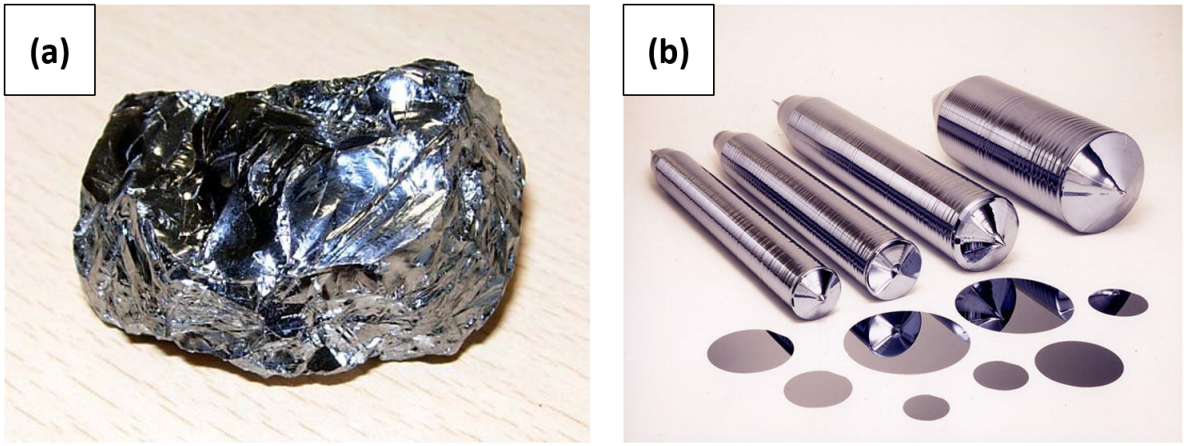

Figure 2.3: (a) A lump of elemental silicon [163]. (b) Purified single crystalline silicon ingots with corresponding wafers sliced from their cross-section [164].

This purified silicon is then typically manufactured into large single crystal boules, or ingots, using two common techniques. The first, known as the Czochralski (CZ) method, melts the silicon in a SiO_2 crucible, where a seed crystal is then contacted with the molten silicon and slowly drawn upwards and rotated in order to extract a large single-crystal boule [165]. This method allows for the deliberate addition of impurities in order to dope the silicon to change its electrical properties but can also add unwanted impurities such as oxygen, given off from the SiO_2 crucible [166]. This method has been further modified in what is known as magnetic field applied Czochralski (mCZ), whereby a magnetic field is applied to the silicon melt which generates a corresponding electric field and Lorentz force, which confines oxygen impurities from the crucible to the outer edges of the melt [167]. This leaves a purer central silicon section with lower oxygen impurities within it [168]. Fundamentally, all crucible processes where the silicon melt contacts the crucible will still impart impurities, deliberate or unwanted, into the melt, particularly at larger diameters where it becomes more difficult to confine them.

The second method produces a very pure silicon compared to any CZ method, known as float-zone (FZ) silicon, where a silicon ingot is melted at the bottom by an inductive heating coil

[169]. This molten zone is then contacted to a seed crystal below and the inductive coil is slowly moved up the ingot while it rotates. This molten silicon then solidifies into a large single crystal with the lack of any direct contact to the ingot, significantly reducing any impurities [169]. FZ silicon is currently limited to ingot diameters of around 200 mm [170].

These boules and ingots can then be mechanically diced and polished into single crystalline silicon wafers. While boules up to 300 mm in diameter are now common, the semiconductor industry is slowly transitioning to wafer diameters up to 450 mm although these are still not the industry standard at the time of writing [171]. Furthermore, increasing the wafer diameter beyond this is is highly unlikely due to both technical difficulty and economic inefficiency [172]. This is an important consideration for the gravitational wave field as it limits the dimensions of any silicon suspension elements and mirrors that may be produced by traditional Czochralski or FZ silicon production methods.

2.3.2 Single crystalline silicon

Silicon wafers are the most common form of single crystalline silicon and are used heavily in the electronics industry. These wafers can be modified by dopants and produced with different crystal structure orientations to create wafers with varying thermal, electrical and mechanical properties.

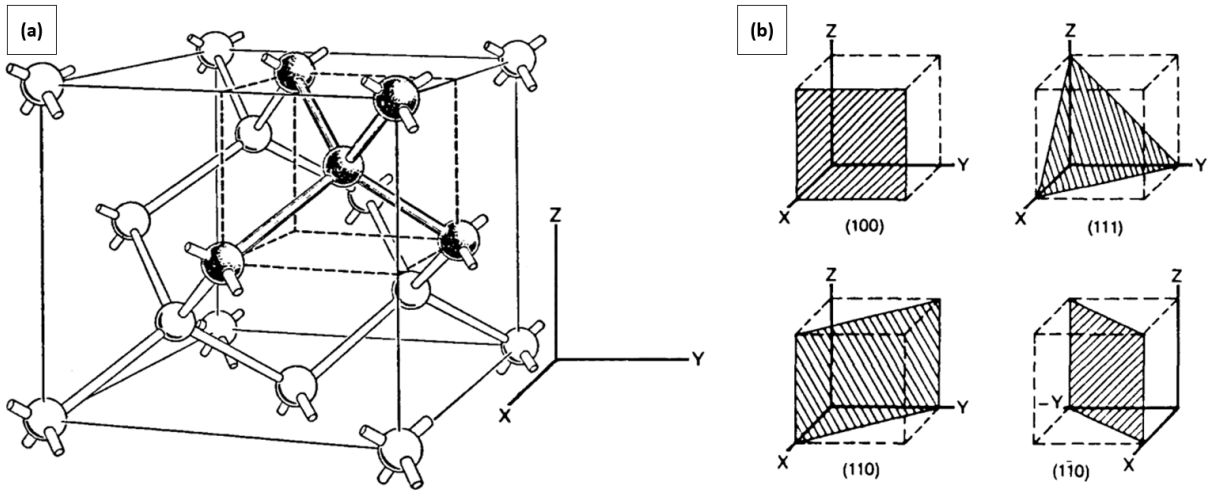

Figure 2.4: (a) Unit cell showing atomic structure of single crystalline silicon with black highlighting that the silicon atom has 4 bonds connecting it to its nearest neighbours. (b) Illustration of 3 primary crystallographic axes of single crystalline silicon. Both images taken from [173] and image (b) has been modified.

As shown in figure 2.4(a), single crystalline silicon has a face-centered diamond-cubic crystal structure. It has 3 primary crystal planes, namely (100), (110) and (111), illustrated in figure 2.4(b).

Axis	Young's modulus (GPa)
E_x, <110> axis	169
E_y, <110> axis	169
E_z, <100> axis	130

Table 2.1: Young's modulus of silicon defined for xyz co-ordinate system of a standard (100) silicon wafer [174].

It is important to understand the crystal axis orientations of any silicon samples as the mechanical and thermal properties differ along each axis [174]. Of particular note, the Young's modulus of silicon changes with crystal axis, as shown in table 2.1.

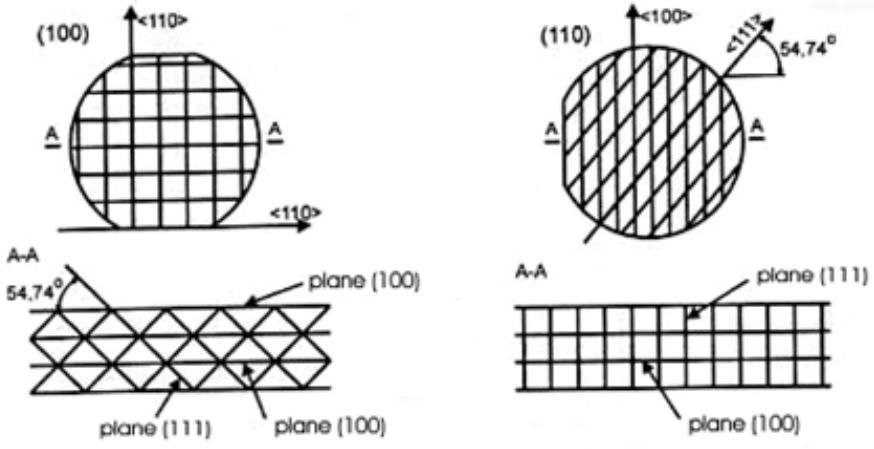

Figure 2.5: Relation of crystal orientations in a (100) and (110) silicon wafer [175].

Typically, silicon wafers are produced in the (100) and (110) silicon surface planes and can be identified by wafer "flats" or notches manufactured as part of the production process. Figure 2.5 show how one can identify silicon wafer axes orientations.

2.3.3 Mechanical loss mechanisms of detector materials

In reality, materials are not ideal, suffering from defects such as dislocations, mechanical imperfections, impurities and dangling bonds [149]. These defects contribute to the overall mechanical loss of a material.

It is important to note here that loss must be considered within the context of energy distribution within the suspension system. If one has a particularly lossy material, but with little energy residing within the lossy region, then the thermal noise contribution may be small. Conversely, if a lot of energy is found within a very low loss region, this may increase the thermal noise contribution to be equal to or higher than a high loss material with little energy within it.

For materials in gravitational detectors, there are a number of distinct loss mechanisms detailed below.

2.3.3.1 Bulk loss

Bulk loss is an intrinsic property of a material and is defined as the mechanical dissipation arising from internal friction within the bulk material itself, separate from other loss mechanisms. This internal friction can arise from impurities and mechanical imperfections in the crystal structure [154]. Doping of the silicon also has a potential effect on increasing the bulk loss at low temperatures [154].

As can be seen in figure 2.6, bulk loss in silicon generally decreases with decreasing temperature over the proposed 3G detector operating temperatures of 18 K and 123 K. Loss peaks in figure 2.6 arise from impurity losses that are avoidable with purer silicon [126]. Specifically, the Si-O-Si peak corresponds to a vibration of oxygen and silicon atoms, arising from oxygen impurities as part of the CZ growth process [155], which is discussed in more detail in section 2.5.1.

Furthermore, crystal orientation is shown to have an effect on bulk loss at temperatures below 40 K with (100) silicon showing lower loss compared to (111) silicon [155].

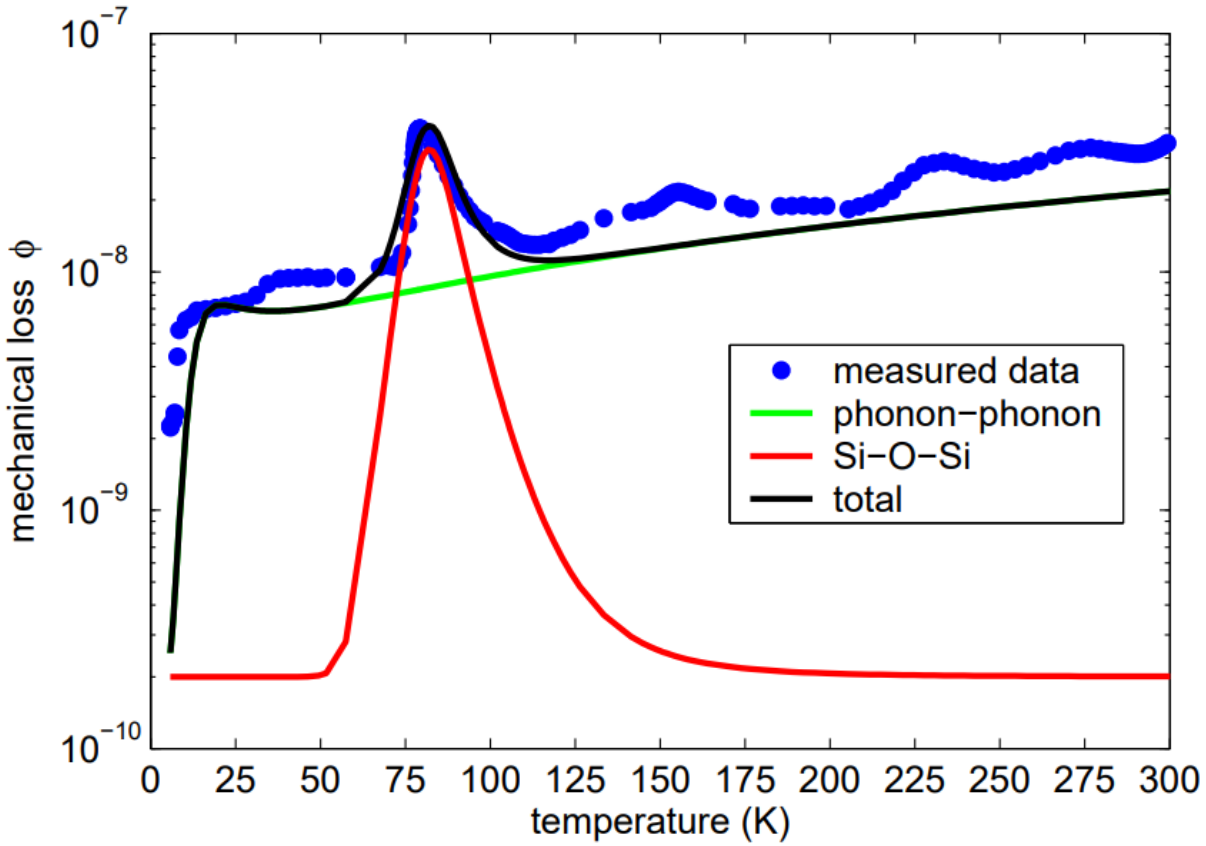

Figure 2.6: Measured bulk loss of (100) silicon [155].

2.3.3.2 Thermoelastic loss

Thermoelastic loss occurs when temperature variations occur in an object. These variations give rise to heat flow within the object.

If one considers a stationary suspension ribbon or fibre, the aim for a gravitational wave detector, local temperature variations will cause heating and cooling throughout the suspension element. This will cause localised expansion and contraction, as determined by the material's thermal expansion coefficient, α, of the ribbon or fibre, in turn mechanically deflecting the ribbon or fibre, causing the suspended test mass to deflect accordingly. This is the source of thermoelastic noise in the detector.

The thermoelastic loss for a tensioned ribbon or fibre can be written as [176]:

$$\phi_{thermoelastic}(\omega) = \frac{ET}{\rho C}\left(\alpha - \sigma_0\frac{\beta}{E}\right)^2\left(\frac{\omega\tau}{1+(\omega\tau)^2}\right) \tag{2.5}$$

where E is the Young's modulus, T is the temperature, ρ is the density, C is the specific heat capacity, α is the coefficient of thermal expansion, σ_0 is the tensile stress in the fibre, ω is the angular frequency and $\beta = \frac{1}{E}\frac{dE}{dT}$ is the thermal elastic coefficient. τ is the characteristic time for the heat to cross the ribbon or fibre which for ribbon geometries is:

$$\tau = \frac{1}{\pi^2}\frac{\rho C y^2}{\kappa} \tag{2.6}$$

where y is the thickness of the ribbon and κ is the thermal conductivity.

For reference, the thermoelastic noise for an infinite sized test mass can be written as [177]:

$$\bar{x}_{thermoelastic}(\omega) = \sqrt{\frac{8\alpha^2\kappa k_B T^2(1+\sigma)^2}{\sqrt{2\pi}\rho^2 C^2 r_0^2\omega^2}} \tag{2.7}$$

where all terms as defined in equation 2.5 are the same, k_B is the Boltzmann constant, σ is the Poisson's ratio and r_0 is the radius at which the incident laser beam intensity has fallen to $\frac{1}{e}$ of its maximum value. A correction factor for a finite sized test mass was derived in [178].

Thermoelastic noise is both frequency and temperature dependent.

2.3.3.3 Surface loss

All materials have finite edges, leading to a region considered as the surface of the material. Surface loss arises from a surface layer, of thickness h, which is lossier than the bulk of the material with a dissipation factor, ϕ_s.

These higher losses can come from a variety of microscopic sources. Local lattice distortions, friction caused by shallow damage and cracks, materials adsorbed onto the surface and dangling chemical bonds are all surface loss contributors although, in general, these mechanisms are not fully understood [153] [179] [180] [181]. One also notes that smoother surfaces

will have lower surface loss, primarily due to the removal of aforementioned features that will minimise mechanical friction between surface features, as shown in figure 2.7 [153]. This is a relevant consideration for the work conducted in chapter 3.

It is worth also noting, silicon grows a native oxide layer when exposed to air [182]. The layer of native oxide typically found on a silicon surface is on the order of nm [182] and is not expected to contribute significantly to the surface loss at detector operating temperatures due to the small thickness. This has been experimentally verified by the inability to measure the loss contribution from native oxides on silicon surfaces. [179] [183].

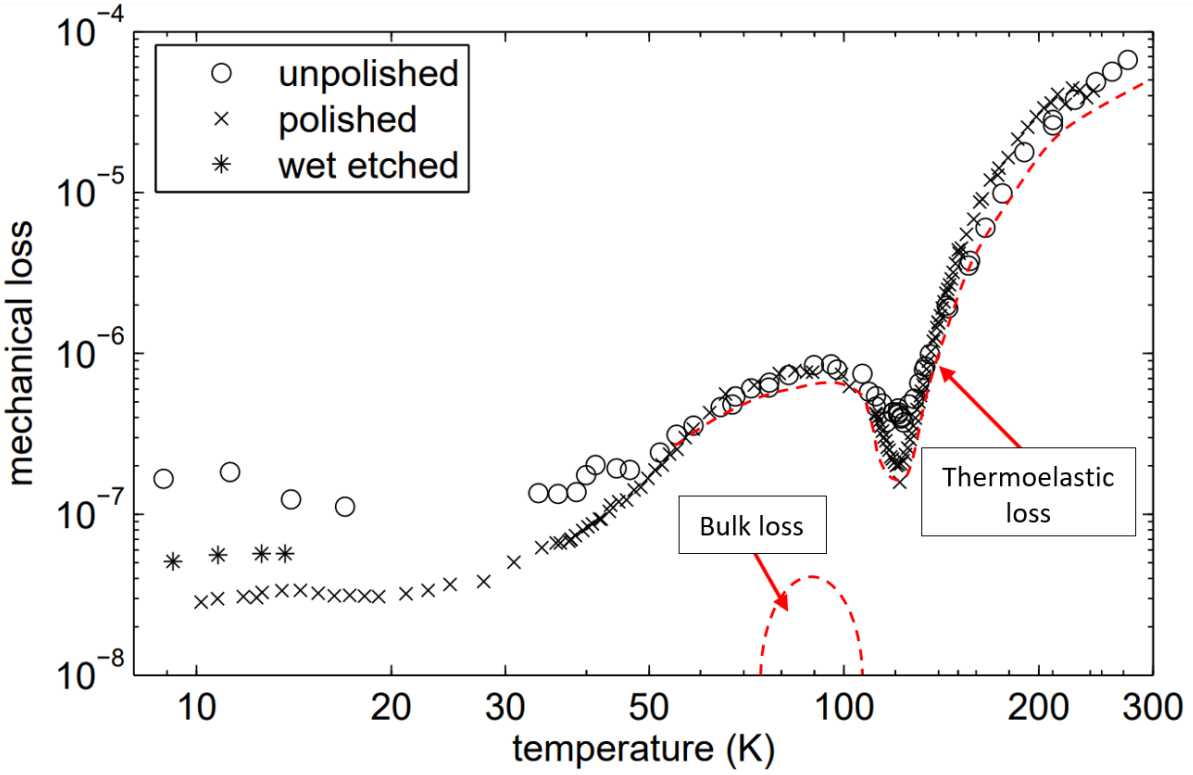

Figure 2.7: Influence of surface roughness on mechanical loss of silicon sample [153]. This image has been modified to give an indicative idea of the relative magnitude of each loss mechanism for these measured samples.

Surface loss can be a significant source of mechanical loss, particularly in geometries with large surface-to-volume ratios such as suspension fibres or ribbons [153].

Surface loss for ribbon geometries can be calculated from the following equation [184]:

$$\phi_{surface\,ribbon} = \frac{3+\frac{x}{y}}{1+\frac{x}{y}} h\phi_s \frac{2(x+y)}{xy} \tag{2.8}$$

where x is the width of the ribbon, y, is the thickness of the ribbon, ϕ_s is the mechanical loss of the material surface and h is the depth over which surface loss mechanisms are believed to occur. Due to a difficulty in measuring the actual surface depth, h, and the corresponding ϕ_s, it is difficult to decouple the two terms so it is common to use the $h\phi_s$ value for thermal noise

calculations. $h\phi_s$ for silicon is taken as 5×10^{-13} [136] [153]. One notes that surface loss is frequency independent.

Figure 2.7 has been modified to highlight the difficulty in directly measuring surface loss. By plotting bulk and thermoelastic loss mechanisms for these samples in red, this gives an indicative idea on the relative magnitude of these different loss mechanisms for a silicon sample measurement. One notes that thermoelastic loss dominates above ≈50 K [153].

2.3.3.4 Bonding loss

The jointing of detector suspension elements is another significant contributor to suspension thermal noise. Traditional mechanical jointing is unsuitable; epoxies have very high relative mechanical loss and mechanical bolting will suffer from flexing and friction that will significantly increase the mechanical loss [185]. Furthermore, jointing techniques have the additional requirements of being UHV compatible and will need to also be cryogenically compatible for 3G detectors.

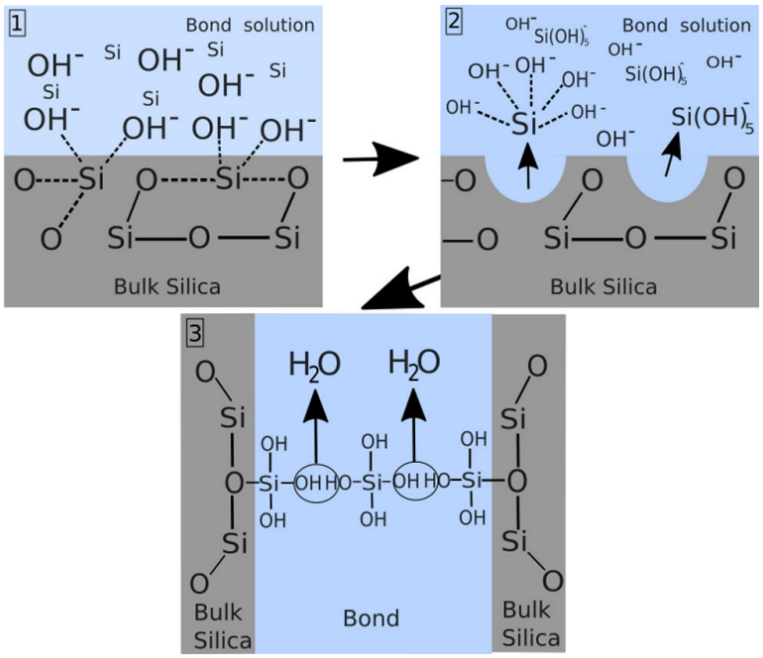

Figure 2.8: Three step process of HC bonding, (1) hydration of silicon dioxide forming weak bonds with silicon atoms on surface, (2) polymerisation where silicon atoms release from bulk to form $Si(OH)_5^-$, (3) dehydration where siloxane chains occur with water migration out of the bond region [186].

For these reasons, a new jointing technique was developed for GEO600 and utilised in the aLIGO and AdV detectors called hydroxy-catalysis (HC) bonding [87] [187] [188] [189]. This technique was used to attach the ears to the fused silica test masses [189]. HC bonding is a precise optical jointing technique that creates chemical bonds between oxide materials using an aqueous hydroxide solution [186].

The process requires oxide surfaces which then undergo three main chemical processes; etching by hydration, polymerisation and dehydration, as shown in figure 2.8.

Fused silica is an oxide material and so can be bonded directly. Silicon has a native oxide layer, typically 1 nm thick [182]. Work conducted by Beveridge noted this native oxide is not thick enough to successfully bond silicon substrates, instead a thermal oxide layer is grown on silicon surfaces to increase this to an order of 100s nm thick [190].

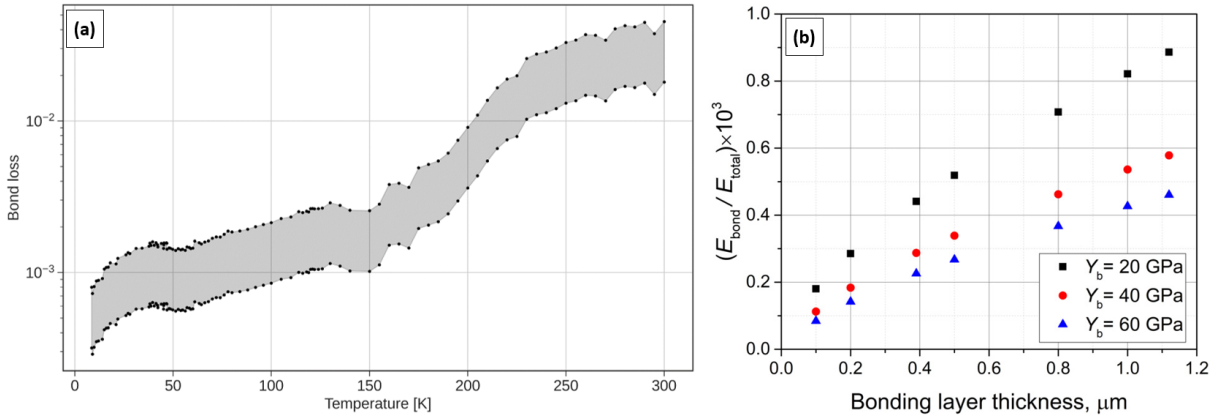

Figure 2.9: (a) Bond loss value range of a sample in which the bond thickness is known to be between 390 nm and 1120 nm over 8–300 K [110]. (b) Ratio of bond energy to total energy against bond thickness as calculated using FEA, image modified from [111].

Silicon-silicon HC bonds benefit from being extremely thin (10-100s nm) [190], strong (30 ± 17 MPa) [186] and with relatively low mechanical loss, as shown in figure 2.9(a) [110].

Figure 2.9(a) shows the possible range of bond loss values measured from a bond thickness that is known to be between 390 nm and 1120 nm [110].

The thicker the bond, the more energy is contained within the bond region, as shown in figure 2.9(b) [111]. It is possible to remove energy from the bond by thinning bonds, however thinning bonds too far will decrease the ultimate bond strength, as shown in figure 2.10 [191].

This issue was tackled in LIGO by careful design of the geometry of the bonded ears in order to keep energy away from the high loss HC bond region between the ears and the test mass. Currently, silicon is not able to be shaped in this way for either the ear, or the ribbons or fibres, and so this method is unlikely to be applicable for current silicon detector designs.

Bond loss values measured so far (figure 2.9(a)) imply bond loss to be one of the largest contributing factors to suspension thermal noise in silicon gravitational wave detectors due to a combination of high loss in comparison to other loss mechanisms and little ability to design silicon geometries to keep energy away from this high loss region. Work conducted in chapter 5 discusses modelling these bonds through finite element analysis (FEA), their contribution to the overall thermal noise and some solutions for mitigation in a silicon detector.

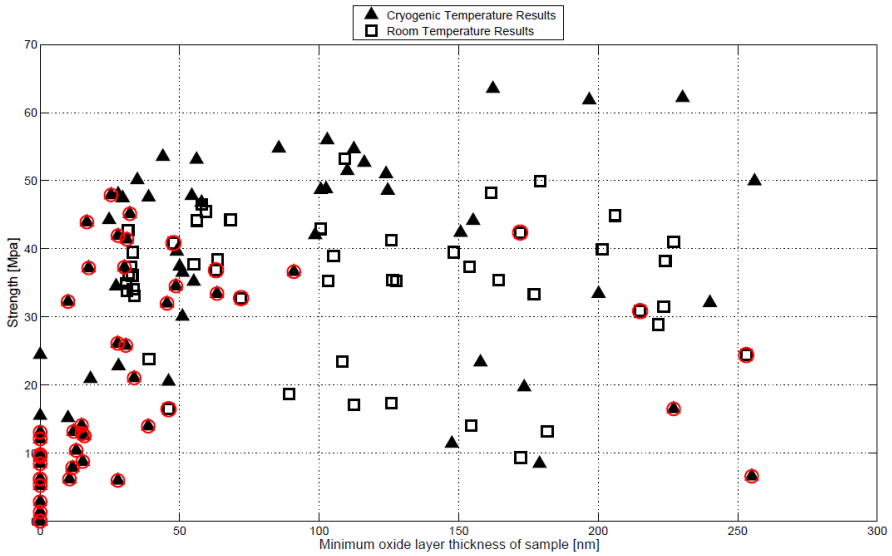

Figure 2.10: Cryogenic and room temperature measurements of silicon-silicon HC bond strengths against oxide layer thickness. From this data a lower limit of 50 nm is required for consistent bond strength [191]. Red circles indicate where fractures occurred across the bond interface.

2.3.3.5 Mechanical loss of silicon

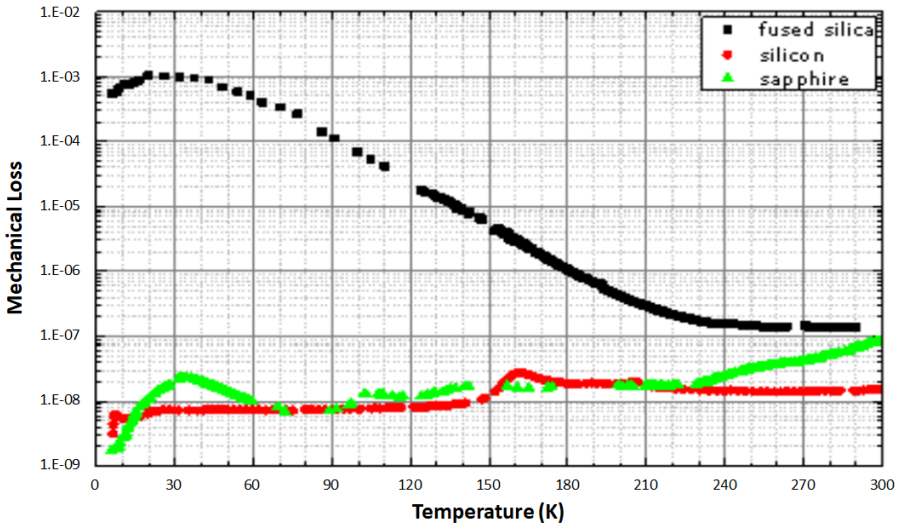

Figure 2.11: Mechanical loss of fused silica, silicon and sapphire from room temperature to 4 K [192].

The bulk mechanical loss of silicon shows promising results compared to fused silica and sapphire at cryogenic temperatures, as shown in figure 2.11. At room temperature, the mechanical loss is dependent on the thermoelastic dissipation in the sample, however at lower temperatures other loss mechanisms, such as surface or bulk loss will dominate [136] [193].

One also notes silicon's mechanical loss was found to vary with crystal axis orientation at

low temperatures [155]. This is due to the fact the thermal noise (and loss) of a bulk substrate is dependent on the Young's modulus which, from table 2.1, varies with crystal orientation.

2.3.4 Coefficient of thermal expansion of silicon

As shown in figure 2.12(a), silicon passes through a zero coefficient of thermal expansion, α, around 18 K and 123 K [194]. A comparison to other relevant gravitational wave materials is provided in figure 2.12(b).

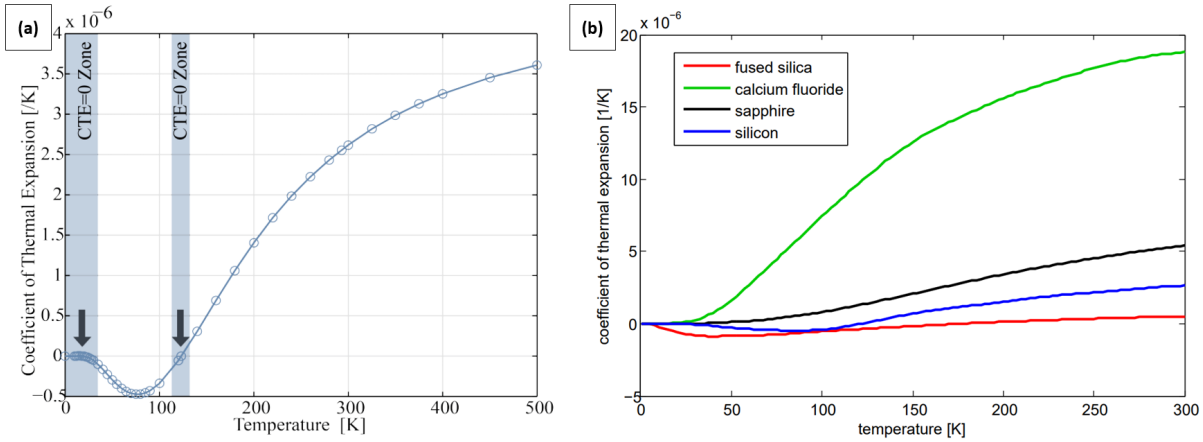

Figure 2.12: (a) Coefficient of thermal expansion of undoped silicon [195]. (b) Comparison of the linear coefficient of thermal expansion of silicon, sapphire, fused silica and calcium fluoride [196].

From equation 2.5, one can effectively minimise or nullify thermoelastic dissipation by reducing the thermally dependent terms, α, β and T. For suspension fibres at room temperature, since α, β and T are fixed values, one can design the fibre geometry so that σ_0 nullifies this term. This is how LIGO fused silica fibres were designed [109], effectively nulling the thermoelastic loss contributions from the fibres.

For silicon, one can not simply nullify α due to the other terms temperature dependence. However, since the value of $\frac{\beta}{T}$ closely matches α, with the correct selection of σ_0, one can minimise thermoelastic dissipation near these temperature nulls as demonstrated in figure 2.13(a). Furthermore, thermoelastic loss from silicon ribbons is not a dominant noise source, compared to surface and bulk loss, for the silicon ribbon geometry shown in figure 2.13(a) [136].

However, one also has to consider thermoelastic loss from the test mass mirror itself because thermoelastic displacement noise is directly proportional to α as shown in equation 2.7 [177]. As shown in figure 2.13(b), this *is* a dominant noise source for silicon test mass mirrors when operating away from the 18 K and 123 K nulls compared to the intrinsic material noise arising from bulk and surface loss. This is the main motivation for running silicon gravitational wave detectors at one of these two operating temperatures.

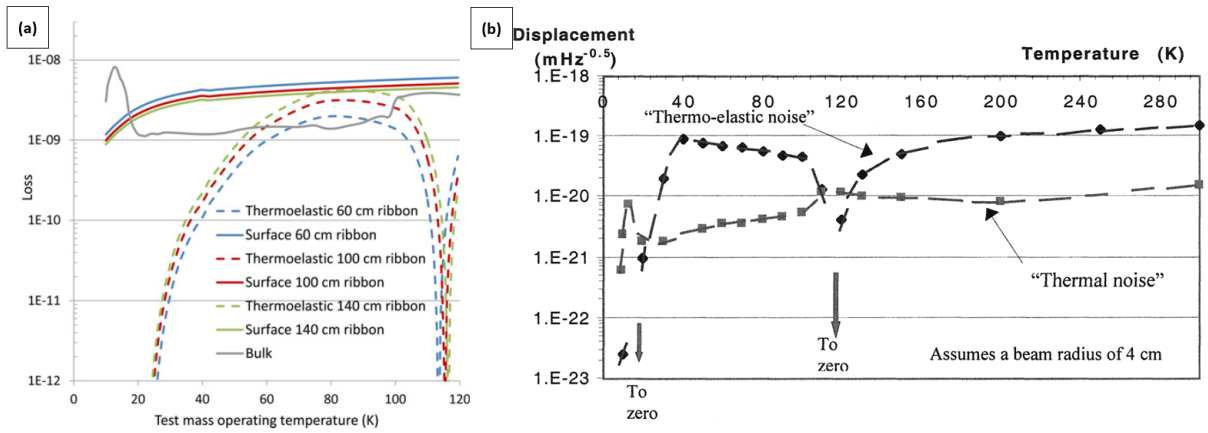

Figure 2.13: (a) Surface, bulk and thermoelastic loss components for a 5:1 aspect ratio silicon ribbon, calculated at 10 Hz [136]. (b) Calculated temperature dependence of thermo-elastic displacement noise and thermal displacement noise at 10 Hz in a single mirror silicon substrate [197]. Note (a) shows loss and (b) shows displacement noise.

Furthermore, operating the detector at these operating temperatures will help to minimise thermal deformation of the test mass mirror, known as thermal lensing [198].

2.3.5 Thermal conductivity of silicon

For 3G cryogenic detectors, heat will be deposited by the interferometer laser into the cold test mass mirror. This arises from a variety of sources but can broadly be split into radiation from nearby objects and absorption of heat into the test mass mirror via point defects and free carrier absorption arising from impurities in the mirror substrate or coating material [122].

For higher temperature operating regimes, e.g. $\approx$123 K, some of this heat can be radiated away from the test mass through careful design [122] [199]. However for lower temperature regimes, e.g. $\approx$18 K, the primary mechanism for heat transfer will be through thermal conduction along the suspension fibres or ribbons [200]. As such, high thermal conductivity of the suspension elements is a requirement for cryogenic detectors.

Heat transfer in crystalline silicon is dominated by phonon transport, even with a large amount of impurities or dopant charge carriers [201]. The thermal conductivity of silicon is limited by 4 main mechanisms; anharmonic phonon-phonon scattering (also known as the Umklapp process), boundary/geometrical scattering, phonon scattering from point defects/impurities and free electron-phonon scattering [202].

Boundary or geometrical scattering is of particular interest for suspension design since the fibre geometry may reduce the suspension elements expected thermal conductivity. Boundary scattering arises when the mean free path of phonons is comparable to the dimension of the sample, and so this has the ability to scatter the phonons [203] [204]. Very thin fibre or ribbon geometries may give rise to boundary scattering, impeding the heat flow through the suspension elements. This is discussed further in chapter 4.

These mechanisms can be reduced by lowering the temperature of the material, ensuring a smooth surface finish and removing defects and impurities within the silicon respectively.

Silicon is known to have a very high thermal conductivity compared to fused silica with a peak around the 25 K as shown in figure 2.14. This peak can significantly change dependent on the aforementioned purity of the silicon due to the mechanisms highlighted, as shown in figure 2.14.

While this high thermal conductivity is advantageous for extracting laser-deposited heat in the test mass mirror, one has to balance the geometry of the suspension elements with the heat extraction. Thickening the suspension elements to tackle this problem will lead to lower dilutions, discussed in section 2.4.5 below, negatively affecting the resultant thermal noise of the suspension [136].

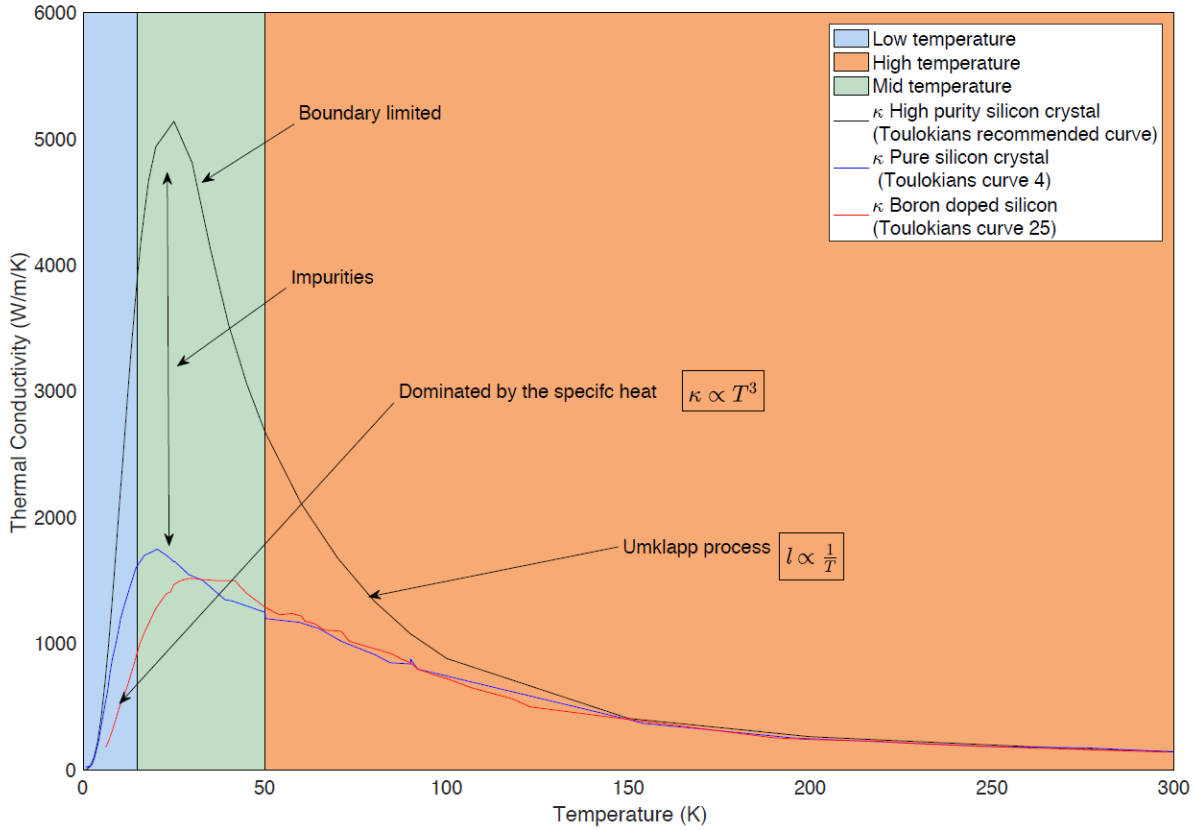

Figure 2.14: Summary of processes affecting the thermal conductivity of single crystalline silicon at different temperature ranges and dopant levels [205]. Image obtained from [202].

Furthermore, similar to the thermal expansion coefficient advantage, high thermal conductivity of the test mass material will help to minimise thermal lensing due to minimisation of thermal gradients [198].

2.4 Resonant modes of the suspension system

Mechanical and thermal excitation of the detector will excite certain suspension modes both in the suspension elements and the test mass mirror itself. Some of these excited modes couple directly, some indirectly, to detector displacement noise.

For the purposes of this thesis, the co-ordinate referencing system is consistent throughout where Z is the primary pendulum direction, parallel to the interferometer beam incident on the test mass, X is horizontally perpendicular to this and Y represents the vertical. However, this varies from the standard LIGO suspension co-ordinate system [206]. This change was for convenience of FEA modelling discussed in chapter 5 and is merely a change of nomenclature.

2.4.1 Pendulum modes

The primary resonant mode of the suspension is the pendulum mode. Mechanically, this mode primarily reduces seismic noise within the detector through the multi-stage pendulum suspension chain. However, this mode directly couples into the position of the test mass mirror surface and changes the arm length in the detector, inducing noise.

For a circular fibre geometry, the mode frequency is the same in both primary and secondary axes of the detector due to the uniform directional stiffness perpendicular to the fibres longitudinal axis. However, for a ribbon geometry of different width and thickness, this creates two pendulum modes, defined as pendulum X and pendulum Z modes in figure 2.15.

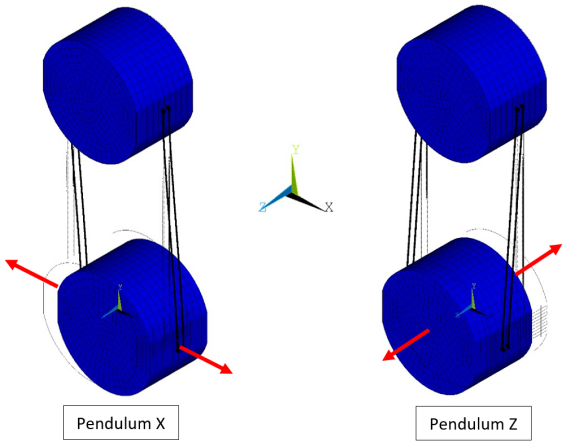

Figure 2.15: Pendulum modes in the X and Z directions of a suspension with a ribbon geometry.

One can calculate the frequency of a pendulum with a ribbon geometry as follows [149]:

$$f_p = \frac{1}{2\pi}\sqrt{\frac{k_{(x/z)}}{m}} \tag{2.9}$$

where m is the point mass and $k_{(x/z)}$ is the spring constant for the ribbon in X and Z directions, as defined by [149]:

$$k_{(x/z)} = \frac{T}{L_{eff}}\left(1+\sqrt{\frac{EI_{(x/z)}}{TL_{eff}^2}}\right) \tag{2.10}$$

where T is the tension in the ribbon, L_{eff} is the effective length of the pendulum, E is the Young's modulus and $I_{(x/z)}$ is the second area moment of inertia in the X and Z directions defined as [149]:

$$I = \frac{wt^3}{12} \tag{2.11}$$

where t is the thickness in the direction of bending and w is the width perpendicular to t.

By increasing the length of and thinning the suspension ribbons, one can push the frequency of the pendulum modes down, out of the detection band of interest as well as reducing the baseline noise for frequencies above the pendulum frequency. This effect can be clearly observed in figure 2.2.

2.4.2 Violin modes

Violin modes are transverse motion of the fibre or ribbon, analogous to the transverse modes of a string vibrating when fixed at both ends. Violin modes for a ribbon geometry will give rise to different fundamental violin modes and higher order harmonics. Fundamental violin modes in the X and Z directions are shown in figure 2.16.

One can calculate the frequency of a violin mode with a ribbon geometry as follows [149]:

$$f_n = \frac{n}{2L}\sqrt{\frac{T}{\mu}}\left(1+\frac{2}{L}\sqrt{\frac{EI_{(x/z)}}{T}}+\frac{EI_{(x/z)}}{2T}\left(\frac{n\pi}{L}\right)^2\right) \tag{2.12}$$

where, n is the harmonic of interest, L is the length of the ribbon, T is the tension in the ribbon, μ is the mass per unit length of the ribbon, E is the Young's modulus of the ribbon material and $I_{(x/z)}$ is the second area moment of inertia as defined in equation 2.11.

While these modes typically enter the detection band, due to the loss low nature of the suspension elements used, much of the energy is focused into a narrow peak around each resonance. These noise peaks can then be removed from the data using filtering.

The thermal noise for violin modes can be calculated by modifying equation 2.1 to give [136] [207]:

$$x_{violin\,mode\,\nu}(\omega) = \sqrt{\frac{4k_BT}{\omega}\frac{2m_{fibre}}{\pi^2m^2\nu^2}\left(\frac{\omega_{violin}^2\phi_{total\,violin}(\omega)}{\omega_{violin}^4\phi_{total\,violin}(\omega)^2+(\omega_{violin}^2-\omega^2)^2}\right)} \tag{2.13}$$

where all terms as defined in equation 2.1 are the same, m_{fibre} is the mass of each fibre, ω_{violin} is the respective violin mode frequency, ν is the respective violin mode number and $\phi_{total\,violin}$

is the total loss for the violin mode, calculated using [208] [209]:

$$\phi_{total\,violin}(\omega) = 2\phi_{total\,pendulum}(\omega). \tag{2.14}$$

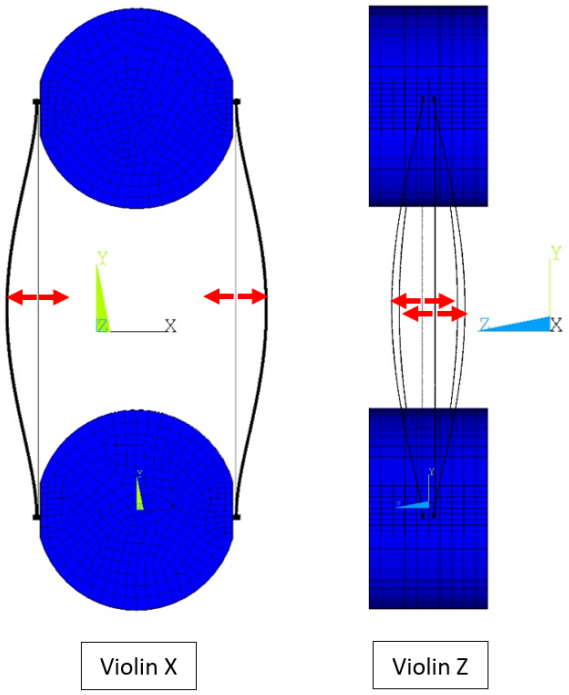

Figure 2.16: Violin modes in the X and Z directions of a suspension with a ribbon geometry.

Violin modes couple into detector displacement noise by producing a small recoil in the test mass due to their motion.

2.4.3 Bounce mode

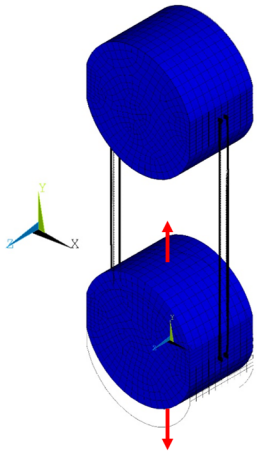

Figure 2.17: Bounce mode of the suspension geometry.

The vertical bounce mode, shown in figure 2.17, can directly couple into the horizontal motion of the test mass. This occurs due to the length of the interferometer arms. Due to the local gravitational acceleration and Earth's curvature there is a small non-vertical coupling contribution arising from the non-parallel hanging of the Fabry-Pérot cavity mirrors, as shown in figure 2.18.

The horizontal cross-coupling is estimated to be approximately 0.1 % for aLIGO and GEO600 suspensions [64] [65].

To calculate the thermal displacement noise for the bounce mode, one can then modify equation 2.1 to give:

$$x_{vertical}(\omega) = 0.001\sqrt{\frac{4k_BT}{\omega m}\left(\frac{\omega_{vertical}^2\phi_{total\,vertical}(\omega)}{\omega_{vertical}^4\phi_{total\,vertical}^2(\omega)+(\omega_{vertical}^2-\omega^2)^2}\right)} \tag{2.15}$$

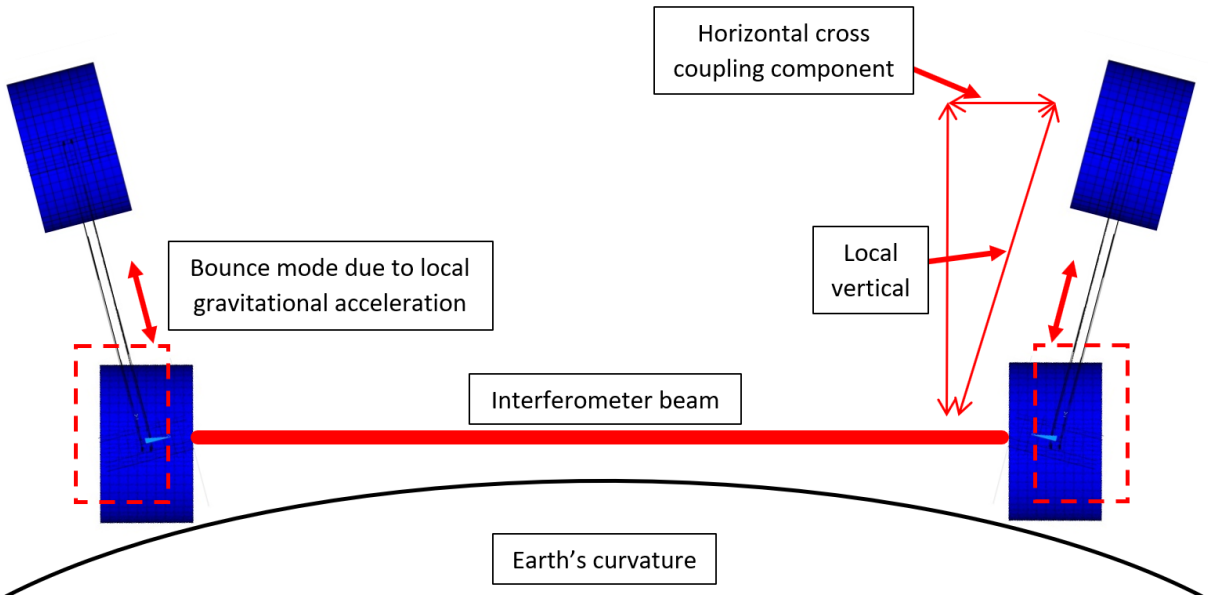

Figure 2.18: Schematic of horizontal cross coupling from vertical bounce mode.

2.4.4 Pitch and torsional modes

Other modes exist that can affect displacement noise in the detector. Pitch and torsional modes, shown in figure 2.19, can couple directly into detector noise by misalignment of the laser beam incident on the test mass mirror surface. This can be minimised by keeping the incident beam as close to the centre of mass as possible.

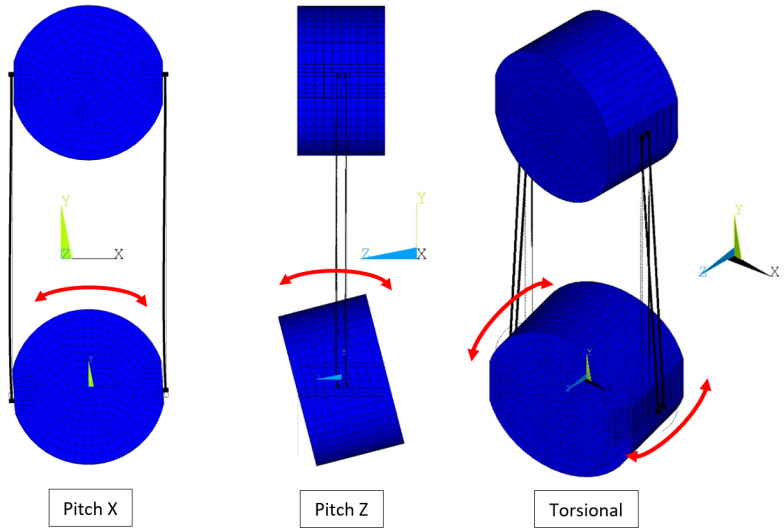

Figure 2.19: Pitch X, pitch Z and torsional modes of the suspension geometry.

2.4.5 Dissipation dilution

The inherent physics of a pendulum system allows for an advantageous effect called dissipation dilution.

If one considers a simple pendulum of mass, m, suspended by a wire of length, L, from an infinitely rigid ideal clamp, when the pendulum is displaced, x from its equilibrium position the pendulum experiences two forces. A gravitational restoring force tries to return the pendulum to its equilibrium rest position while there is an elastic force arising from the stiffness, k_{wire}, of the tensioned wire itself.

As the pendulum is displaced, potential energy is stored in the gravitational field due to the displacement of the pendulum mass, however potential energy is also stored in the elastic force within the tensioned wire.

The potential energy stored in the wire is:

$$E_{wire\,potential} = \frac{1}{2} k_{wire} x^2 \tag{2.16}$$

Dissipation of energy will occur in the bending of the wire from internal friction, resulting in dissipation of some fraction δ of the potential energy stored with every bending cycle. The energy lost per cycle is:

$$E_{lost\,per\,cycle} = \frac{1}{2} \delta k_{wire} x^2 \tag{2.17}$$

Therefore, rearranging equation 2.4, the mechanical loss of the wire is:

$$\phi_{wire}(\omega) = \frac{\frac{1}{2}\delta k_{wire} x^2}{2\pi(\frac{1}{2} k_{wire} x^2)} = \frac{\delta}{2\pi} \tag{2.18}$$

The energy stored in this displaced system is the summation of the potential energy stored in the wire and the energy stored in the gravitational field, giving:

$$E_{potential} = E_{wire\,potential} + E_{gravitational\,potential} = \frac{1}{2}k_{wire}x^2 + \frac{1}{2}k_{gravity}x^2 \tag{2.19}$$

It is vital to recognise here that there is no loss mechanism in the gravitational field. Therefore the mechanical loss, $\phi_{pendulum}(\omega)$, for this whole system can be written as [144]:

$$\phi_{pendulum}(\omega) = \frac{\frac{1}{2}\delta k_{wire}x^2}{2\pi\left(\frac{1}{2}k_{wire}x^2 + \frac{1}{2}k_{gravity}x^2\right)} = \phi_{wire}(\omega)\frac{k_{wire}}{k_{wire}+k_{gravity}} \tag{2.20}$$

where $k_{gravity}$ is the lossless gravitational spring constant and k_{wire} is the elastic spring constant. It is clear to see that $k_{wire} < k_{wire} + k_{gravity}$, therefore $\phi_{pendulum}(\omega) < \phi_{wire}(\omega)$.

Since gravitational wave suspension wires are typically very thin and long, supporting large test masses, this means $k_{wire} \ll k_{gravity}$, so one can simplify equation 2.20 to:

$$\phi_{pendulum}(\omega) \approx \phi_{wire}(\omega)\frac{k_{wire}}{k_{gravity}} \tag{2.21}$$

The ratio of gravitational to wire elastic spring constants is known as the dissipation dilution factor, D, and can be expressed as [144]:

$$D \approx \frac{k_{gravity}}{k_{wire}} = \frac{\frac{mg}{L}}{\frac{\sqrt{TEI_{(x/z)}}}{2L^2}} = \frac{2mgL}{\sqrt{TEI_{(x/z)}}} \tag{2.22}$$

where g is the gravitational acceleration, T is the tension in *one* wire of the suspension, E is the Young's modulus of the wire material and $I_{(x/z)}$ is the second area moment of inertia as defined in equation 2.11.

For a pendulum with more than one suspension wire, such is the case for gravitational wave detectors, two more terms must be introduced to give:

$$D = \frac{2mgL}{\beta n\sqrt{TEI_{(x/z)}}} \tag{2.23}$$

where n is the number of wires and β takes the value of 1 or 2.

The use of n identical wires results in bending, and hence dissipation, in all wires, therefore the dilution must be multiplied by the number of wires.

The β term is dependent on the suspension wire bending, taking the value of 1 when the suspension wire bends only at the top and 2 where it bends at the top and bottom, as illustrated in figure 2.20. This will be entirely dependent on the suspension mode considered as each mode will have its own dilution.

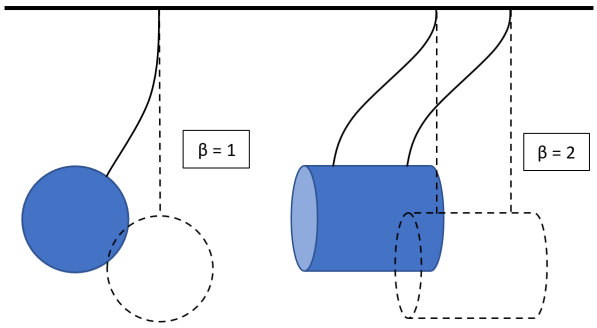

Figure 2.20: Illustration of suspension wire bending with corresponding β values to be used.

Dissipation dilution shows, on top of seismic isolation, the clear advantages for pendulum suspensions, particularly when care is taken to select the correct fibre or ribbon geometries, material properties and total mirror mass, in order to increase dilution. With increasing dilution, one benefits in a reduction of the resultant thermal noise for the corresponding suspension mode since less thermal energy is present off resonance.

2.5 Challenges for silicon detectors

There are a number of significant challenges that will need to be addressed in order to realise 3G cryogenic silicon gravitational wave detectors that will be discussed in this section.

2.5.1 Silicon manufacturing and jointing

As discussed in section 2.4.1, one requires the flexibility to design suspension ear and fibre or ribbon geometries in order to push suspension resonant peaks out of detection band frequencies.

Shaping of silicon on the scales required for an ET-like suspension remains a significant challenge for the field. Samples obtained for this thesis work, discussed in chapter 3, rely on cutting samples from silicon wafers using lasers but, as discussed in section 2.3.1, ultimate wafer sizes are likely to be limited to 450 mm in diameter, far short of the proposed 1.5 m length ET suspensions.

For long suspension fibres, CZ, mCZ and FZ methods of growing silicon may not be feasible due to impurity levels and locations within the boules or ingots. Furthermore, it is not known whether these long boules or ingots can be cut vertically along their length in order to produce long, thin fibre geometries. If one used mechanical dicing, it may be extremely difficult to cut such a long delicate geometry from the boule in the first place, and the surface finish would likely be far too rough to meet strength and surface requirements.

As such, other techniques may provide possible options. One such technique is the μ-pulling technique, involving the downward pulling of a crystal fibre though a micronozzle at the bottom of a crucible heated using a radio frequency (RF) generator [194] [210]. Currently this method

requires better control of diameter regulation and crystal orientation alignment of the produced fibres. Furthermore, this technique uses a crucible meaning impurities are likely to be found within the finished silicon product.

Another technique, developed for low optical loss in infra-red telecommunications, involves laser-induced recrystallisation to produce single crystalline silicon fibres. This method starts by extruding a silicon core clad in fused silica, which can be created to any arbitrary length. This glass fibre, with a polycrystalline silicon core up to 100 µm thick, can then be heated by a CO_2 laser that is absorbed strongly by the silica cladding [211]. This locally melts the polycrystalline silicon interior and allows it to recrystallise into a single crystalline silicon fibre. This has been demonstrated for both circular and square cross-sectional geometry [212]. One could then theoretically remove this silica cladding by dissolving the assembly in hydrofluoric (HF) acid, leaving behind a silicon crystal.

A final technique being developed at Leibniz-Institut für Kristallzüchtung (IKZ) involves melting silicon granules via inductive heating with a silicon "self-crucible" stabilised around the melt. A silicon ingot is then grown from the melt. This "self-crucible" method means no contamination from the crucible however this approach is currently limited by the available purity of the silicon granules which contain contaminants in the form of carbon, although work is in progress to increase the purity of the granules [213] [214]. Furthermore, IKZ is working on developing and improving experimental methods for silicon fibre production [215].

These production techniques are currently immature and have yet to be extensively demonstrated. The work in this thesis does not explicitly cover the production techniques for silicon ribbons or fibres, rather the work in chapter 3 investigates the concepts of post-production strengthening of silicon, via surface treatments. These techniques provide promising abilities to produce the silicon ribbons or fibres required for 3G generation detectors.

Precise shaping of silicon fibres or ribbons have yet to be demonstrated. As in aLIGO, suspension element geometries will likely have to be precisely shaped to push the bending energy far away from high loss bonded regions in the suspension, such as the HC bonded ears. Ear geometries akin to aLIGO have yet to be demonstrated in silicon, especially at the required surface quality and exact geometry required to minimise thermal noise.

Research is also required in order to obtain large silicon test masses, on the scale of 200 kg required for ET. Currently no such silicon exists in high purity form on these scales. However, work is ongoing, including investigations of directional solidification techniques to produce quasi-monocrystalline silicon. This silicon is available in sizes larger than that required for 3G gravitational wave detectors and shows initial promising results for mechanical loss [216].

Finally, as discussed in section 2.3.3.4, if one is to pursue HC bonding in 3G silicon detectors, this requires a layer of thermally grown oxide to allow chemical bonding to take place. Thermal oxides on silicon typically grow at temperatures >900 °C [217]. Oxidising large test masses of this size locally may not be possible and oxidising the entire test mass is unlikely to

be practically feasible or desirable. Preliminary work on direct bonding of silicon is under way at the University of Glasgow to investigate high strength, low loss bonds [218].

2.5.2 Welding silicon

One particular challenge is how to install and repair suspensions. As will be shown in chapter 4, proposed silicon suspensions, developed for this thesis, require ribbons to be HC bonded to silicon ears which in turn will likely be HC bonded to the test mass mirror. HC bonds can not be de-bonded once cured [219].

In aLIGO, the quadruple suspension system is built by welding fused silica fibres onto the ears of the test masses. Welding has the added advantage where broken fused silica fibres can be removed and a new fibre re-welded in place [148]. This property is highly desirable in order to reduce detector downtime should a suspension fibre break. This also means the entire ear or test mass mirror and remaining fibres do not need to be replaced, reducing costs substantially.

Preliminary investigation of silicon welding shows particular difficulties when trying to emulate such a repair. Work was conducted by the author and Dr Alan Cumming to try and weld two pieces of silicon together using the CO_2 laser apparatus used to create fused silica fibres [220], as shown in figure 2.21(a).

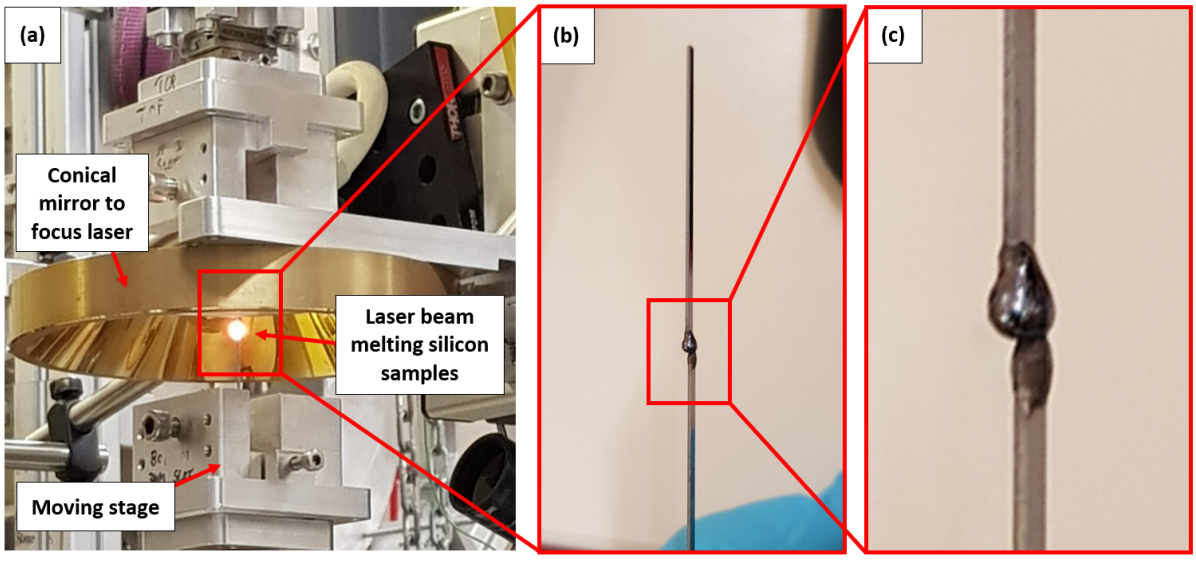

Figure 2.21: (a) CO_2 laser pulling apparatus, originally designed for pulling fused silica fibres. (b) Result of melting and joining two silicon ribbons together. (c) Zoom on (b) of the weld region showing a "balling" effect on the ends of the silicon when molten.

While it was seemingly possible to attach two pieces of silicon together, the joint was weak and easily broken. When the silicon became molten, it was noted the silicon transitioned from solid to liquid almost instantaneously at a certain laser power. This behaviour is consistent with silicon's emissivity profile where the emissivity dips at the point of phase transition [211] [221].

The silicon then "balled", due to surface tension of the liquid silicon, and behaved almost like water with a very low viscosity; which contrasts strongly with the very viscous nature of fused silica when molten, which is key to control when welding. The liquid balls on either end of the silicon pieces were slowly pressed together while molten. The balls showed no change in surface tension behaviour when in contact and the weld was allowed to cool. The resultant weld is shown in figures 2.21(b) and (c). One can clearly see this ball shape has been retained as the weld has cooled, leaving a poor joint between the pieces.

2.5.3 Cryogenics

Cryogenic detectors using silicon raise a number of mechanical and thermal challenges to be overcome. Operating detectors in cryogenic regimes is not trivial. The infrastructure required to cool detectors down can be a source of noise itself, and careful design is required in order to operate with efficient cooling while not imparting vibrational noise into the suspension system [222].

All materials, sensors and actuators must be cryogenically compatible, and considerations must be made for differing material coefficients of thermal expansion and radiative properties [223]. Chapter 4 will discuss these challenges and offer some practical solutions for silicon in more detail.

Thin conductive suspension elements for ET-LF will provide cooling challenges, particularly with the lower proposed operating temperature of 18 K. This temperature regime will necessitate cooling via conduction only through the suspension elements [200]. Conductive heat transfer within the suspension system needs to be carefully designed to minimise any potential heat flow inhibitors through the highly conductive silicon suspension elements. This will require high thermal conductivity bonds in order to maximise conductance of the suspension system, though preliminary work has shown that detector cool-down times are achievable on realistic timescales [202]. Proposals to reduce cooling times further involve the application of a helium contact gas into the vacuum system that can then be pumped out before laser light injection [127]. All of this is still to be experimentally proven for silicon in practice with such large masses suspended from thin conductive suspension elements.

Figure 2.22 shows work conducted by Cumming *et al*, demonstrating the ribbon and fibre strength in an ET-LF type detector is currently the limiting factor; rather than thermal extraction [136]. However, one must also consider the fibre thickness required to extract the deposited laser heat from the test mass should strength be improved to a level to be no longer limiting.

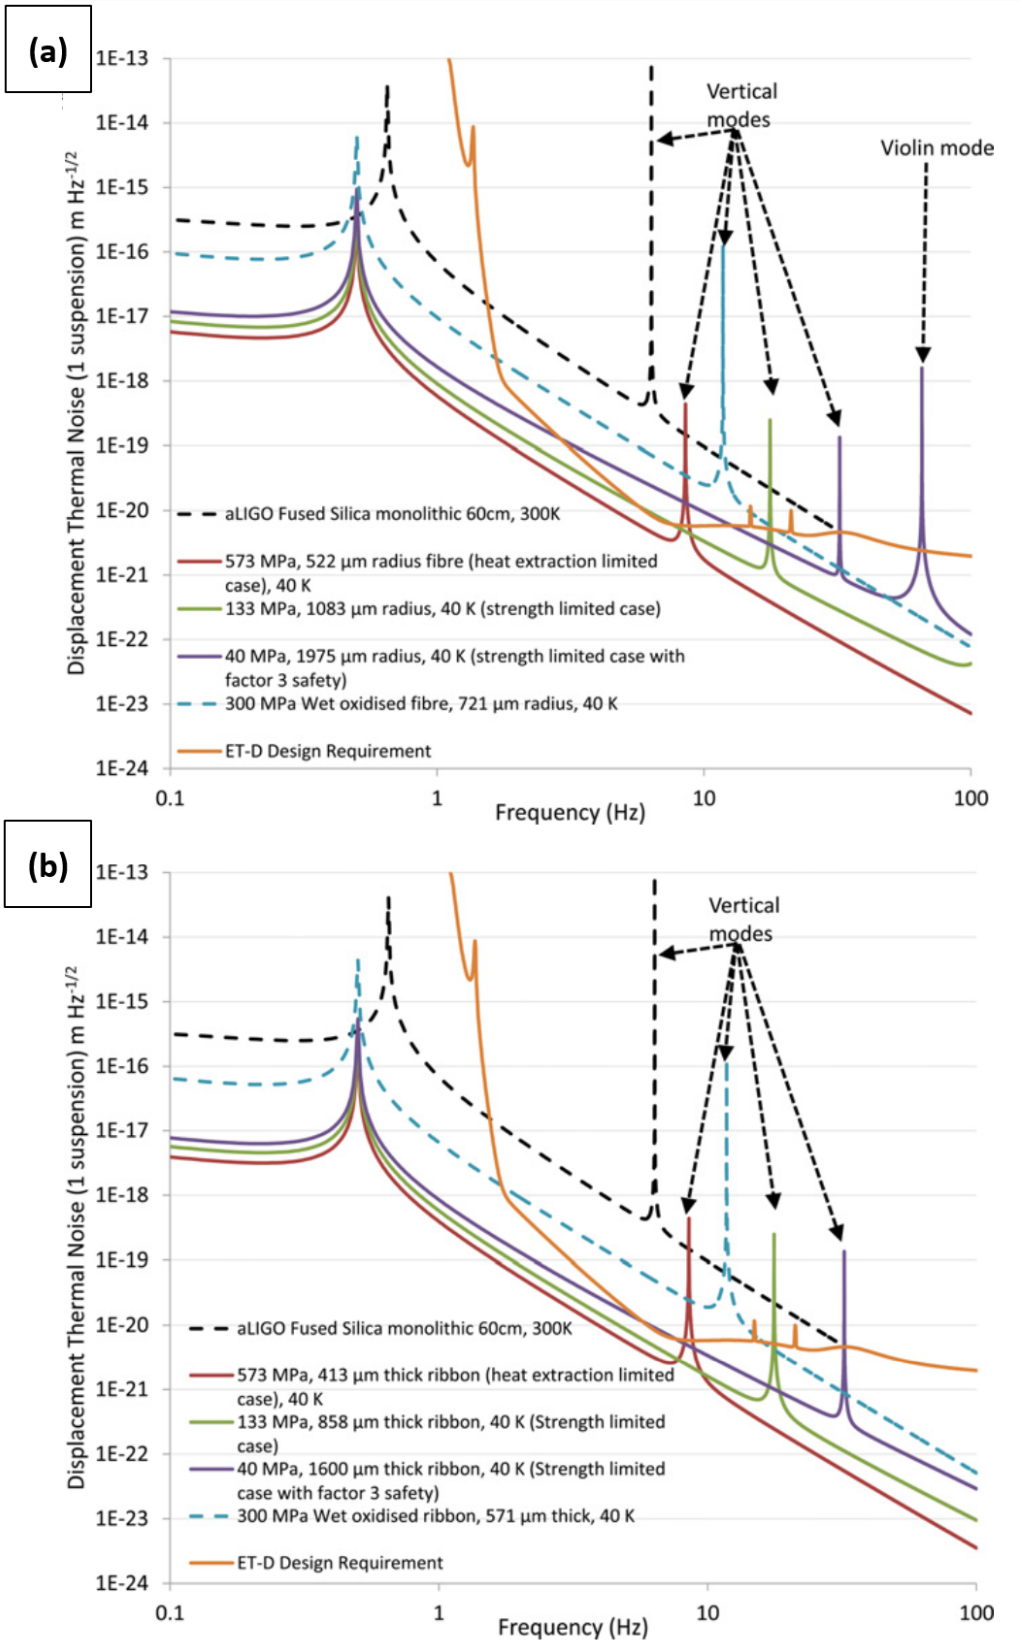

Figure 2.22: Displacement thermal noise of a single 1 m long silicon mirror suspension for (a) circular cross section fibres, (b) rectangular cross section ribbons [136].

2.5.4 Silicon strength

If one is to significantly increase the mass of the test mass, as discussed in section 2.2.1, the primary consideration must be for the suspension fibres or ribbons to be able to mechanically support the mass.

Silicon ribbons have been observed to be significantly weaker than fused silica fibres, with reported average strengths of 133 MPa for mechanically polished samples [136] [224]. By contrast, the current aLIGO suspensions run at ≈768 MPa [225]. A safety factor is obviously important considering people will need to work in and around the suspension during installation and commissioning of the detector. One also should consider the cost and time lost should a suspension fibre fail. aLIGO currently runs with a safety factor of around 6 [226], made possible due to ultimate tensile strengths for fused silica of ≈4.4 GPa [227]. Brittle materials such as silicon generally require a slightly higher safety factor since they do not have a well-defined yield point [228]. Silicon ribbons are currently a factor of 33 times weaker than fused silica which poses a significant engineering challenge.

It is important to note, one can not simply increase the cross-sectional areas of the silicon fibres or ribbons to mitigate their low breaking stress. From equations 2.5 and 2.23, this would directly impact on the thermal noise, as thermoelastic loss would increase and dissipation dilution would decrease for increasing fibre or ribbon thickness. Ultimately, increasing fibre or ribbon thickness to accommodate for stress will have a negative impact on the suspension thermal noise. From figure 2.2, decreased ribbon tension would also bring the bounce mode frequencies further into the detection band as well as bringing the violin mode frequencies down. The work conducted in chapter 3 is aimed at increasing the ultimate tensile strength of silicon by way of surface treatments.

Furthermore, silicon's mechanical properties pose further engineering challenges. If one conservatively uses a safety factor of 6 with current reported silicon strengths, for an ET-LF design, this would require a stress in each fibre or ribbon of 22 MPa. For the proposed suspension length of 1.5 m and assuming a Young's Modulus of 169 GPa along the suspension length, each fibre or ribbon would stretch by around 0.2 mm [127]. For comparison, aLIGO fibres stretch by 6.5 mm when loaded [229]. Since the suspension extension for silicon is very small, there is a much more significant engineering challenge to ensure they can not become overstressed during the assembly process. This stretch shows that even small additional movements of the mass could put significant extra stress on the suspensions elements. Moreover, in a 4 fibre or ribbon suspension, keeping all the fibre or ribbon lengths the same, such that the tensions and stresses in fibres are the same, will also be much more challenging. Typically machining tolerances of tooling used to align and attach fibres are 25–50 μm, this being significant as compared to the silicon fibre or ribbon stretch [230]. aLIGO addresses this issue by de-stressing as part of the laser welding process which equalises stress throughout the fused silica suspension elements [231]. As discussed in section 2.5.2, laser welding is not currently possible for silicon. Engineering

challenges are further explored in chapter 4.

2.6 Conclusion

This chapter discusses the benefits and remaining challenges of using silicon as a material for 3G detectors.

By reducing the temperature through cryogenic cooling, one can directly improve suspension thermal noise by removing thermal energy from the suspension system. Cryogenic cooling still presents challenges for practically cooling suspension structures and test masses.

Reduction in temperature then also exploits the multiple advantageous properties of silicon; the coefficient of thermal expansion nulls, the high thermal conductivity and the low mechanical loss. Jointing and manufacturing technologies will need to be developed in order to minimise their loss contributions and allow suspension resonant frequencies to be moved to place them outside the gravitational wave detection frequency bands. Silicon manufacturing technologies will also allow for lengthening of the suspension chain, which will increase dissipation dilution, further reducing loss contribution and hence improving suspension thermal noise.

Furthermore, increasing the mass of the test mass will also increase dissipation dilution, lowering thermal noise further and helping to push suspension resonant frequencies out of the detection band. However, it is vital that the strength of silicon is increased in order to support the proposed large-scale test masses for ET-LF with low-noise suspension fibres or ribbons. Surface treatments for increasing the ultimate tensile strength of silicon are discussed in the following chapter.

Chapter 3

Silicon surface treatments for strength optimisation

3.1 Introduction

As discussed in chapter 2, 3G gravitational wave detectors will operate with higher mirror masses in both room temperature fused silica and cryogenic silicon regimes [122] [123] [127]. Silica fibres will be scaled up in cross-section and in suspended stress, this being possible as the current working load is well below the ultimate tensile strength of fused silica [125]. However, silicon fibres have been observed to be much weaker than silica [136] [224], therefore pursuing methods to increase this strength are important for 3G suspension development.

For ET, particularly ET-LF, the proposed cryogenic silicon test mass is 211 kg [126] [127]. With a current proposed fibre radius of 2.2 mm [127], this leads to a stress of 34 MPa in each fibre.

Fundamentally, if one operates the detector at lower stresses in each fibre, this means larger cross-sectional fibres which then leads to an increase the thermoelastic noise and a reduction in the dilution, increasing the overall detector noise. This also brings violin frequencies further down, or the vertical bounce and pendulum frequency up, into the detection band.

It is clear to see, increasing the ultimate tensile strength of silicon will ultimately allow for lower noise suspensions.

3.2 Silicon surface damage

The mechanical strength of single crystalline silicon can be affected by the damage on the surface layer [232] [233] as is common for all crystalline materials. Since the surface to volume ratio for a ribbon or fibre is larger than for a bulk material or mirror substrate [144], it is generally believed to be one of the main factors in reducing silicon's strength [136]. Point defects (interstitial and vacancy defects), linear defects (dislocations in the atomic lattice) and bulk defects

(atomic impurities) are commonly attributed to a further decrease in mechanical strength.

Some sources quote theoretical values of the ideal ultimate tensile strength of single crystalline silicon ranging from 16 GPa [234], calculated by means of modelling silicon nanowires from first principles, up to 23 GPa [235] using two-atom unit cell calculations. However, experimental work has rarely demonstrated these strength values, since defect and edge damage free crystals do not exist in a practical setting. Silicon nanowires have produced estimated fracture strengths of up to 20 GPa, with experimental results showing strengths of 16 GPa [236]. While it is tempting to use these values as the ultimate benchmark; they are likely far from attainable strengths in the context of macroscopic silicon fibres or ribbons.

Experimental results have shown, that by reducing the length or diameter of a silicon nanowire, one can actually increase its fracture strength due to what is known as the Da Vinci size effect [237]. The thinning to a nanowire scale results in a significant reduction in the quantity of material, reducing the statistical likelihood of a flaw, defect or crack being present.

Bulk strength measurements of silicon wafers have yielded experimental values of up to 8.8 GPa [238]. This is regularly quoted as benchmark strength for bulk, macroscopic crystalline silicon. However, great care should be taken when interpreting this work, as the method of strength testing was deliberately designed to eliminate edge effects of the wafer and is therefore not analogous to fibre usage in suspensions where exposed fibre edges are prominent and unavoidable.

Quantitative data on subsurface silicon damage is not easily found as most literature tends to look at qualitative models of subsurface silicon damage [239] [240] [241]. The few quantitative sources found suggest the subsurface damage extended to 10s of μm into the surface of the silicon although these examples related to fine polishing or grinding of silicon wafers designed to smooth surfaces at microscopic scales [242] [243]. Since silicon ribbons will need to be cut on the macroscopic scale, work damage from dicing saws or cutting lasers is of more practical interest. From literature investigation of the cutting of macroscopic silicon via these methods, Schwartz and Robbins states, "*each etching period had been chosen to remove 4 to 6 mils from the specimen, it may be assumed that the work damage had been removed after the first etch*" [244]. Taking the upper limit, this suggests a removal of at least 150 μm should remove all work damage on silicon samples. Further evidence would suggest, subsurface damage on the scale of 100–200 μm seems a realistic target [238].

Furthermore, the surface quality will also have an effect on the surface loss mechanism, contributing to the overall loss of any gravitational wave detector design, as discussed in chapter 2.

3.2.1 Fundamentals of stress, strain and surface energy

Before investigating how surface quality impacts on the structural and mechanical properties of a material one must first discuss the fundamental quantities used to characterise these materials.

3.2.1.1 Strain energy

In a one dimensional situation, a linear application of force, F, resulting in a small deflection, dx, of an atom with a spring constant, k, from equilibrium position causes a change in its potential energy, dW. The energy can be determined from Hooke's law:

$$\begin{aligned} dW &= F\,dx \\ F &= kx \\ W &= \int kx\,dx = \frac{1}{2}kx^2 \end{aligned} \tag{3.1}$$

where this potential energy, W, is termed strain energy.

If the stress between two atoms is increased until the bond is broken, the strain energy becomes available as bond potential energy, neglecting dissipative losses due to heat, sound, etc. The resulting two separate atoms have the potential to form bonds with other atoms. These now separate atoms can be considered to be a "surface". For a solid consisting of many atoms, the surface atoms have a higher energy state compared to those in the interior.

3.2.1.2 Surface energy

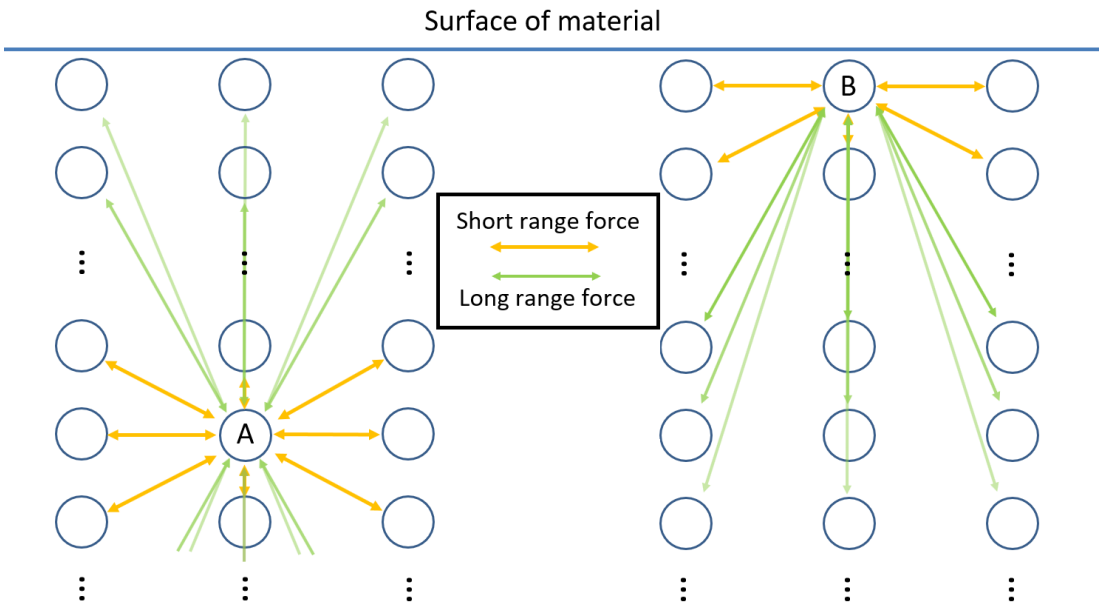

Figure 3.1: Illustration of forces acting upon atoms of the interior, A, and at the surface, B, of a solid.

Consider an atom A within a solid, as illustrated in figure 3.1. Long-range van der Waals forces, typically acting over 0.3–0.6 nm, and short-range Coloumb repulsive forces, typically acting over <0.03 nm, act equally in all directions on this particular atom and the atom takes an equilibrium position within the material [245]. Now consider atom B on the surface of the solid.

Such an atom is attracted by attractive forces extending over many atomic dimensions [246], in the direction into the material volume only. The creates a net attractive forces into the interior of the solid. The corresponding short-range repulsive force extends only to within the order of an atomic diameter. This repulsive force acts to halt the inward movement of atom B, however due to the net attractive force this creates a new equilibrium position for the surface atom. Thus atoms on the surface move inwards until the repulsive short-range forces of the nearby atoms balances with the long-range attractive forces from adjacent atoms *and* the next few rows of atoms beneath them.

These surface atoms therefore have greater potential energy and act like a thin tensile skin, which is shrink-wrapped onto the body of the material. In liquids, this effect manifests itself as the familiar phenomenon of surface tension. Surfaces of solids also have surface potential energy but the effects of surface tension are not readily observable because solids are not so easily deformed as liquids. The surface energy of a material, stored as an increase in compressive strain, represents the potential that a surface has for making chemical bonds with other like atoms.

3.2.2 Griffith crack criterion

This section investigates the theory surrounding surface damage and its effect on the ultimate tensile strength of silicon rather than internal defects such as dislocations and impurities.

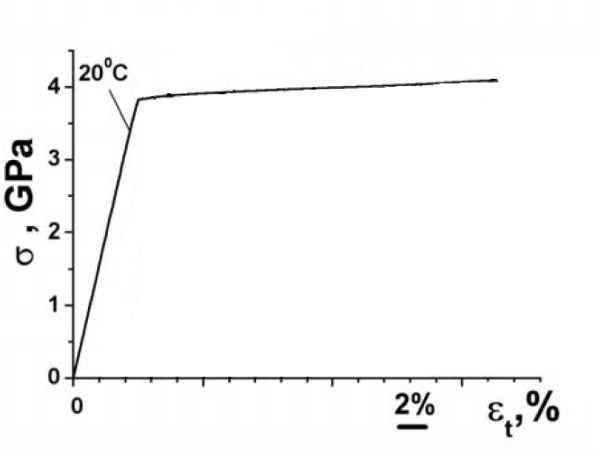

Figure 3.2: Stress–strain curve for single crystalline silicon at 20 °C. Image modified from [247].

Silicon is an ideal brittle material with a very short elastic zone and virtually no plastic zone at room temperature that leads to catastrophic failure in the material under loading as illustrated in figure 3.2 [248] [249].

Surface damage and its effect on tensile strength is well described by the Griffith crack criterion which is applicable to the fracture dynamics of a brittle material [250].

3.2.2.1 Stress concentrations

The presence of sharp corners, notches or cracks in a material can concentrate the applied stress around these points. Work conducted by Inglis in 1913 [251], showed that local stresses around a corner or hole in a stressed plate could be magnified to many times higher than the average applied stress. It was shown that the degree of stress magnified was related to the radius of curvature of the hole or corner. The smaller the radius of curvature, the higher the stress concentration.

Inglis discovered that the stress concentration factor, κ, for an elliptical hole was given by:

$$\kappa = 1 + 2\sqrt{\frac{c}{\rho}} \tag{3.2}$$

where c is the hole radius and ρ is the radius of curvature of the tip of the hole. This equation shows that the stress concentration factor is independent of the absolute size or length of the hole, instead depending only on the ratio of the size to the hole's radius of curvature.

3.2.2.2 Energy balance criterion

Griffith noted that larger cracks appeared to propagate more easily than smaller ones, which seemingly contradicted Inglis' claim that the stress concentration factor was independent of the absolute size or length of the stress concentrator [250].

This helped lead Griffith to the addition of a second condition in order to reconcile the apparent anomaly by considering the minimum potential energy [246]:

- The bonds at the crack tip must be stressed to the point of failure. The stress at the crack tip is a function of the stress concentration factor, which depends on the ratio of the hole radius to its radius of curvature
- For an increment of crack extension, the amount of strain energy released must be greater than or equal to that required for the surface energy of the two new crack faces.

where the second condition can be express mathematically as:

$$\frac{dU_s}{dc} \geq \frac{dU_\gamma}{dc} \tag{3.3}$$

where dU_s is the strain energy, dU_γ is the surface energy and dc is the crack length increment. This equation shows that for a crack to extend, the rate of strain energy released per unit of crack length must be at least equal to or greater than the rate of surface energy required to create two new crack face surfaces.

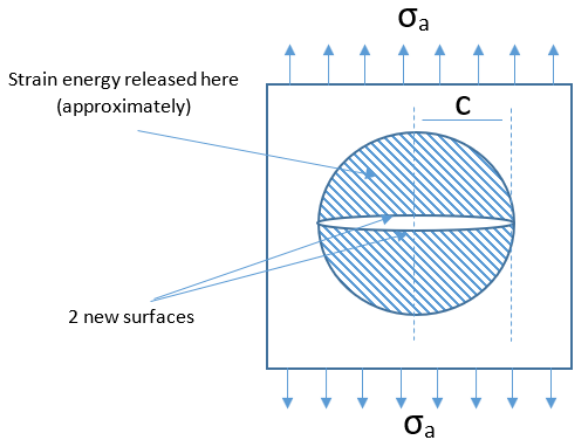

Figure 3.3: Illustration of Griffith's application of Inglis' stress calculations to a narrow elliptical crack in an infinite plate of unit width under a uniform stress, σ_a.

Griffith applied Inglis' stress calculations for a very narrow elliptical crack to show the strain energy released by introducing a double-ended crack of length $2c$ in an infinite plate of unit width under a uniform stress σ_a, as illustrated in figure 3.3, giving [250]:

$$U_s = \frac{\pi \sigma_a^2 c^2}{E} \tag{3.4}$$

where E is the Young's modulus.

The total surface energy for two surfaces of unit width and length $2c$ is given as:

$$U_\gamma = 4\gamma c \tag{3.5}$$

where γ is the fracture surface energy of the solid. This is usually larger than the free surface energy of the material since fracture processes may involve energy dissipative mechanisms such as microcracking, phase transformations and plastic deformation [246] The factor 4 arises from the fact there are two crack surfaces each of length $2c$.

Thus, by differentiating equations 3.4 and 3.5 with respect to c, in order to take the form of second condition equation 3.3:

$$\frac{\pi \sigma_a^2 2c}{E} \geq 4\gamma. \tag{3.6}$$

Simplifying this, one ends up at the Griffith energy balance criterion for crack growth:

$$\frac{\pi \sigma_a^2 c}{E} \geq 2\gamma. \tag{3.7}$$

Equation 3.7 highlights that, below a critical crack length, the crack will not extend spontaneously. Beyond this point, more energy becomes available by the released strain energy than is required by the newly created crack surfaces which leads to unstable crack growth and ultimate fracture of the material.

It is important to note the energy balance criterion is an indicator of whether crack growth is possible, as a necessary, but not always sufficient, condition for material fracture. For example, if the crack tip is blunted or rounded then the crack may not extend because of an insufficient stress concentration. Fracture *only* occurs when the stress at the crack tip is sufficient to break the bonds there.

For a given stress, there is a minimum "safe" crack length that is not self-propagating. A crack will not extend if its length is less than the critical crack length, c_c [246]:

$$c_c = \frac{2\gamma E}{\pi \sigma_a^2} \tag{3.8}$$

Rearranging equation 3.8 one can find the theoretical failure stress, σ_f, of a material if one knows the level of surface damage by way of a crack length, c_c:

$$\sigma_f = \sqrt{\frac{2\gamma E}{\pi c_c}} \tag{3.9}$$

Further work continued to develop the Griffith crack criterion through Orowan's investigation of plastic crack tip zones [252] and Irwin's modification to include ductile materials [253] but due to the aforementioned properties of silicon, namely its lack of plastic zone and brittleness, this work will not be explored in any more depth.

3.3 Previous work

Tensile strength testing work has previously been conducted to support the design of fused silica fibres and develop and enhance their tensile strength [77] [125] [225] [227] [254] [255] [256]. Further work was completed on surface treatments and tensile strength testing of silicon ribbons for future 3G cryogenic instruments [136] [224].

Previous silicon work [136], shown in figure 3.4, is the precursor to the work reported in this chapter. Nine sample sets were strength tested with six sample sets undergoing surface treatment to improve the surface quality and ultimate tensile strength. This work investigated samples cut from (110) and (100) silicon wafers, whereas the work in this chapter will only look at samples cut from (100) silicon wafers due to availability.

Control sample sets one and two were mechanically polished silicon samples cut from the (110) and (100) wafers respectively. Sample set 3 was obtained through etching the sample from a silicon wafer with no further surface treatment conducted to provide a comparison with mechanically polished samples. The remaining 6 sample sets came from three surface treatments applied to both crystal orientations.

Sample sets 4 and 5 had the sample faces coated with silicon nitride Si_3N_4. Sample sets 6 and 7 were wet oxidised. Sample sets 8 and 9 were wet oxidised on the edges with the sample faces coated with Si_3N_4.

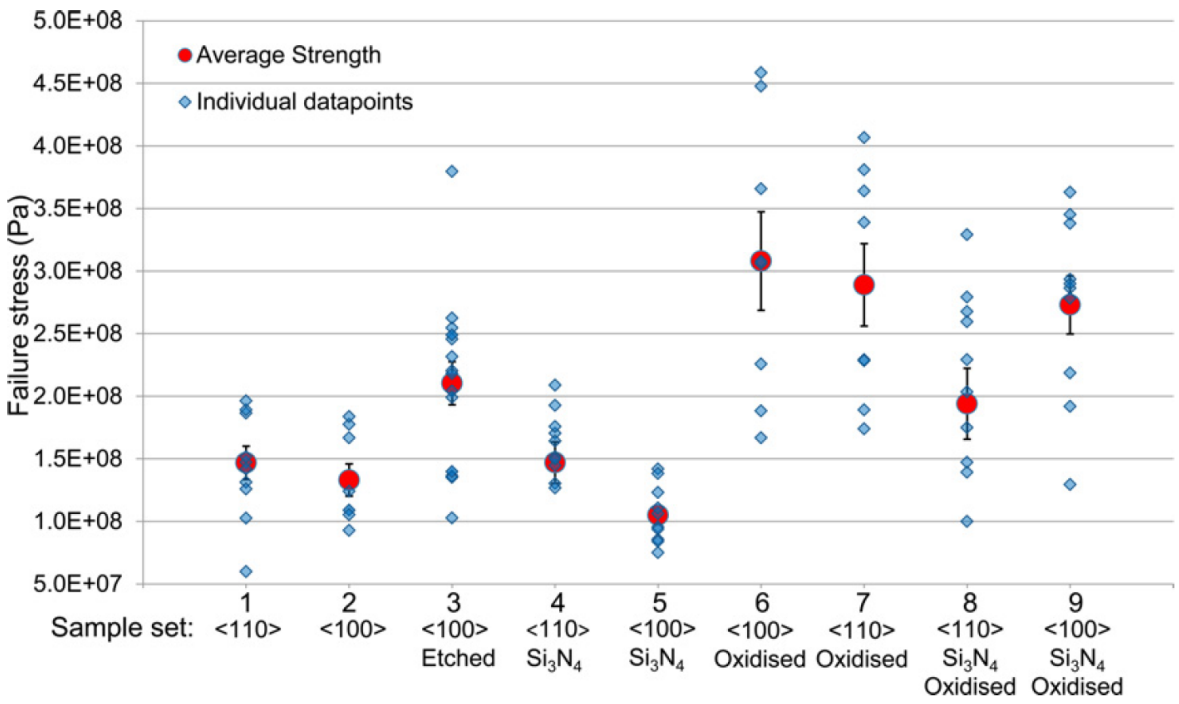

Figure 3.4: Strength test results for nine different types of sample, showing individual data points and average values with their standard error [136].

As can be seen from the results in figure 3.4, some key points were deduced. Firstly, looking at sample sets 1 and 2, it appears that crystal orientation of the silicon sample is not a dominant influence on tensile strength for the crystal orientations investigated. Secondly, the chemically etched samples were stronger, on average, by 50 % compared to the mechanically polished control sample sets one and two, providing strong evidence that the surface quality of the edges correlates to the tensile strength. Scanning electron microscope (SEM) imaging of surface damage of the mechanically polished control samples, combined with the results of the Si_3N_4 coated faces, suggested the surface quality was the limiting factor on strength.

Finally, of particular note, was the strength for the sample sets 6, 7, 8 and 9, all of which were oxidised along the edges of the samples. It is suspected that by creating an outer barrier of oxide, one can effectively "patch up" surface damage cracks, reducing these stress concentrators and increasing the ultimate tensile strength. While it is tempting to pursue this technique due to the substantial increases in tensile strength, the authors note the fundamental issue with coating silicon ribbons, in what is effectively a layer of fused silica, is the high mechanical loss arising from a wide mechanical dissipation peak at cryogenic temperatures [157].

This work concluded that the concept of surface treatment of silicon can ultimately increase its tensile strength. However, this work also notes that any surface treatment applied must not add excess loss which would increase thermal noise.

3.4 Silicon samples for surface treatment and strength testing

The work in this chapter aims to build on this previous work to increase the ultimate tensile strength of silicon by focussing on the removal of surface damage from silicon samples. Three surface treatment methods were investigated; chemomechanical polishing, argon ion etching and wet chemical etching. These treatments were applied to silicon ribbons where a control was used to investigate the effects of the surface treatments.

Silicon ribbons were laser cut from wafers by Laser Micromachining Ltd (LML), Wales [257], due to cost and previous experimental work with their samples. Examples of the silicon ribbons are shown in figure 3.5(c). LML's industrial pulsed laser cutting technology is proprietary so the exact cutting process is unknown but it is generally assumed to involve a high power, short pulse laser, suspected to be around 1 µm [258]. This is fired perpendicular to the surface of the sample in order to locally pierce the material and then repeated either linearly or at varying angles in order to cut the shape required. This technology is limited by wafer thickness, up to around 1 mm thick. Generally, simpler geometries with lower accuracies are obtained with increasing wafer thickness [259].

50 samples were left over from previous work. These were cut by LML in 2014 using an "older technique" of laser cutting. 561 new samples, using a newer technique, with new lasers and processes, were cut in 2018 and 2019. Details of the changes are proprietary, but there was a noticeable difference in edge quality of the silicon as will be shown in section 3.4.2.4.

3.4.1 Substrate silicon wafer and sample details

All wafers used were (100) single crystalline silicon double-side polished and either 525 µm or 775 µm thick. The wafer properties are recorded in table 3.1.

Wafer thickness (µm)	Wafer diameter (mm)	Dopant	Resistivity (Ωcm)
525	100	Phosphorous (n-type)	0.1-0.5
775	300	Boron (p-type)	1-100

Table 3.1: Wafer properties for laser cut samples.

All 525 µm thick samples were cut to align as best as possible with the <110> crystal axes on the (100) wafer. LML's rotational error, quoted as 0.02° [259] is much smaller than the wafer manufacturing error of 1° which was therefore taken as the accuracy of alignment. Approximately 75 % of the 208 775 µm thick samples were cut to align as best as possible with the <110> crystal axes on the (100) wafer leaving approximately 25 % cut in a random, off-axis orientation. This was to maximise the usage of wafer space on wafers used for blade spring designs manufactured for a separate study, shown in figure 3.5(a).

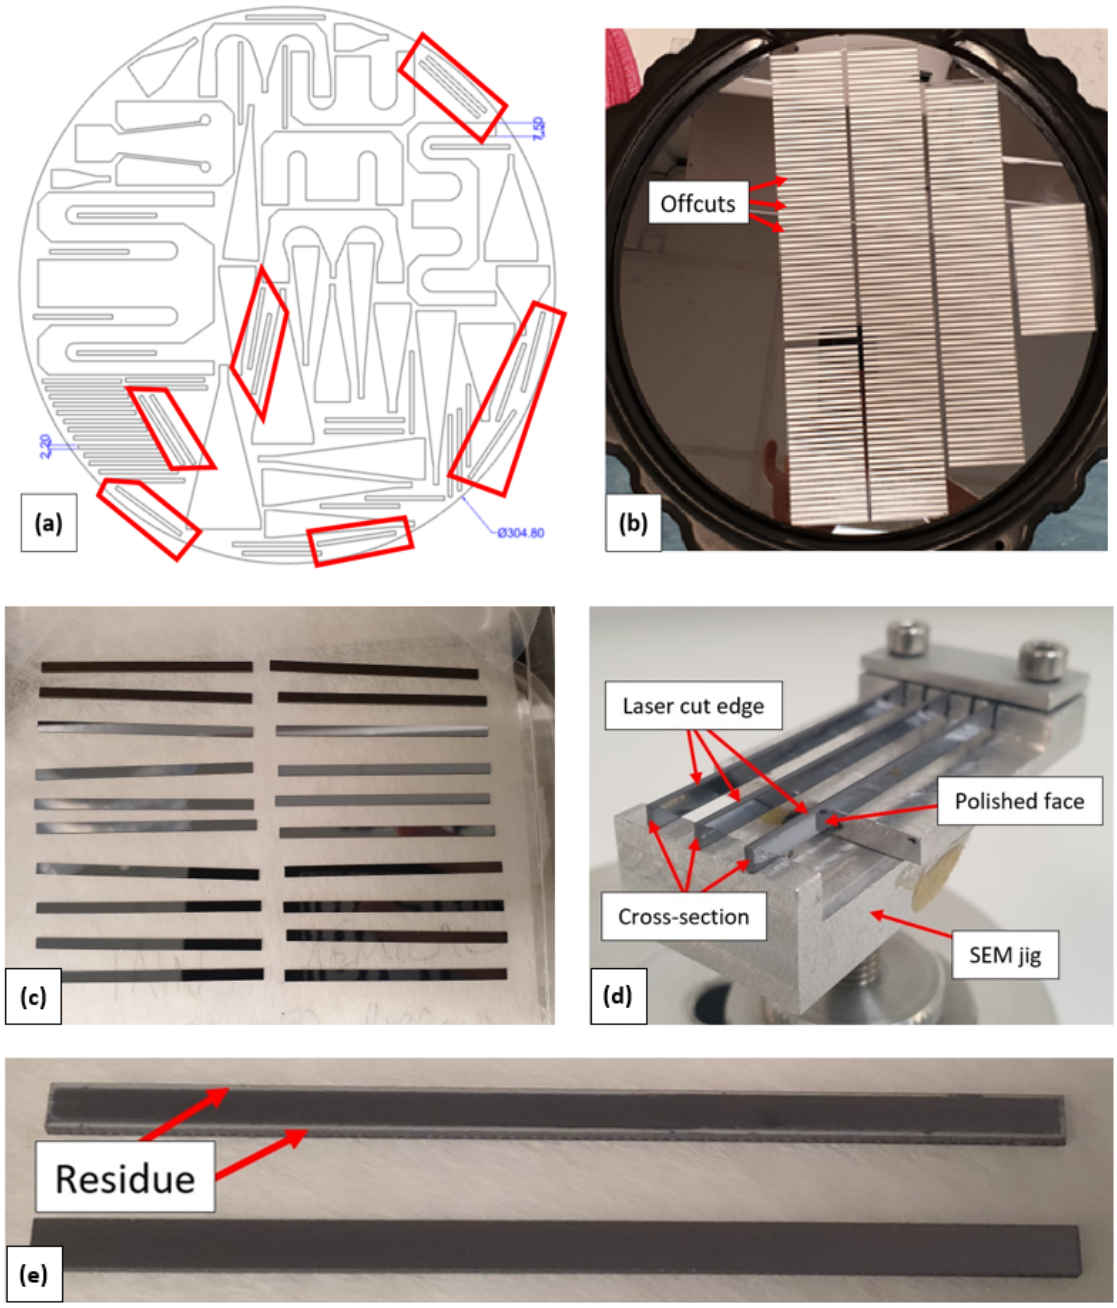

Figure 3.5: (a) Example of randomly orientated, off-axis cut ribbon samples amongst blade spring designs, highlighted in red. (b) Crystal axes orientated cut samples. (c) Typical example of a batch of 20 silicon ribbons as cut and delivered from LML. (d) Photograph of 3 775 µm thick samples, cut using the new laser technique, labelling the edges, faces and cross-sections of the ribbon geometry held in a jig. This jig is also used for imaging the samples in a SEM, in this case in an edge-on configuration. (e) An example of a silicon ribbon (top) as received from LML showing white residue around the laser cut edges. The sample below shows another silicon ribbon (bottom) after the residue has been cleaned.

Going forward, it should be assumed that all samples are crystal axis orientated; unless otherwise stated. An example of the negative from which ribbons have been cut along the

crystal axes from a 775 µm thick, 300 mm diameter wafer is shown in figure 3.5(b).

All samples were rectangular in shape, with dimensions of 45 mm by 2.2 mm in order to repeat the dimensions used in previous work [136]. LML documentation states cutting tolerances are ± 50 µm [260].

When the samples were received, some exhibited white dust or residue left over from the laser cutting process as shown in figure 3.5(e). All samples were wiped with acetone using Anticon Gold wipes [261], followed by a subsequent wipe with isopropyl alcohol (IPA) to remove this residue and any other surface contaminants before inspection.

Sample sets of between 17 and 20 were allocated to each surface treatment, giving suitable statistics for analysis, consistent with previous work [136].

3.4.2 Laser cut silicon surface and edge quality

In total, 474 silicon ribbon samples were characterised. Some additional testing was conducted on offcuts. These offcuts can be seen in figure 3.5(b) still attached to the silicon wafer.

The primary surface laser cutting damage exists around the edges of the silicon ribbons. The damaged edges along the length of the ribbons are of paramount interest due to their larger surface area and imperfect finish. The damaged cross-section edges, shown in figure 3.7, are not of interest due to the tensile strength testing techniques employed in section 3.4.2.6. The untreated, laser cut samples discussed within this section were used as a control group. Surface characterisation of these samples relied primarily on SEM imaging, followed by optical surface roughness measurements, photography, dimensional profiling and visual inspection.

For each sample set, between 3 and 5 samples were taken as a representative set for further imaging to measure their surface morphology using the SEM and surface roughness measurements. All samples underwent photography, dimensional profiling and visual inspection. For surface treatment methods, sample sets were further characterised pre/post-surface treatment and/or before strength testing.

3.4.2.1 Visual inspection

Under visual inspection, chips were occasionally observed on face-edge corners of the samples, and were clearly observed on sample faces due to their high reflectivity, an example of which is shown in figure 3.6. This damage was noted to allow attempted correlation with the failure point when later strength tested.

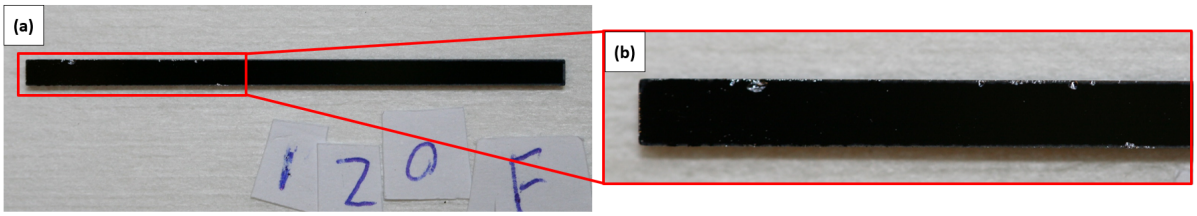

Figure 3.6: (a) Photograph of the front face of sample 120, 775 µm cut using the new laser technique. (b) Magnification of (a) on the face-edge vertex chipping on sample 120.

3.4.2.2 Photographic imaging

Figure 3.5(c) shows the 45 mm × 2.2 mm silicon ribbon faces which were reflective in nature on both sides due to them being cut from the double-side polished wafers. The laser cut edges of the sample were dull and non-reflective where the laser ablation has pierced and ablated the silicon material.

It was noted that the old laser cutting technique produced a non-rectangular cross-sectional ribbon which was latterly seen to improve with the newer laser cutting technique. This can be seen in figure 3.7 and was further verified using SEM imaging discussed below. This cross-sectional shape may lead to an over-estimation of cross-sectional area and therefore an under-estimation of the sample tensile strength. Pixel analysis was conducted on the cross-section of the old laser cut control samples. This yielded a 2 % difference maximum between the widest and narrowest sections.

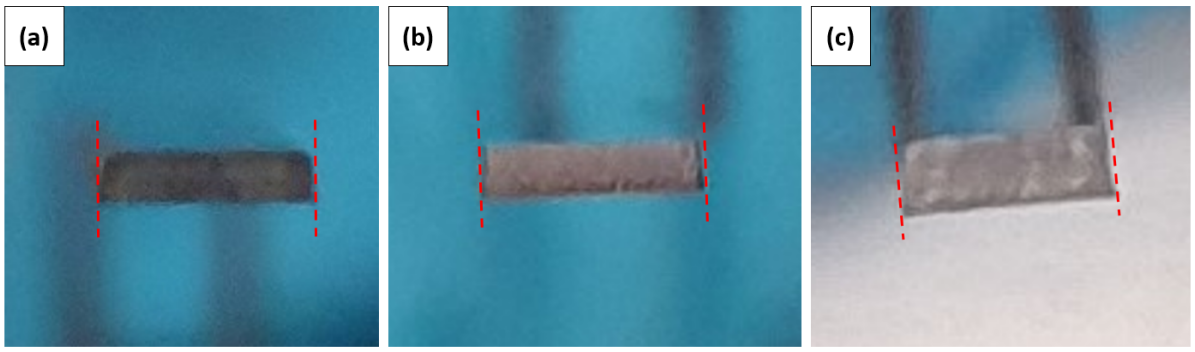

Figure 3.7: Photographs looking along the 2.2 mm×525 µm or 775 µm cross-section of the ribbons. (a) Non-rectangular cross-section of a 525 µm sample cut using the old laser technique. (b) A more rectangular cross-section of a 525 µm sample cut using the new laser technique. (c) A more rectangular cross-section of a 775 µm sample cut using the new laser technique.

3.4.2.3 Surface roughness

For each sample set, between 3 and 5 samples were imaged using a Wyko NT1100 Optical surface profiler [262]. This profiler uses white-light interferometry to image the surface roughness of optically reflective samples. These samples were imaged on the front and back polished faces to quantify their surface roughness since these surfaces were reflective enough for analysis. The imaged areas were approximately 2 mm $\times$ 2 mm. All sample faces showed an arithmetic average surface roughness, R_a, of < 5 nm.

Due to the dull, non-reflective damaged edges of the sample, it was not possible to profile the surface of the sample edges.

3.4.2.4 Scanning electron microscope imaging

For each batch, between 3 and 5 samples were imaged using a Hitachi TM-1000 SEM [263].

A jig was designed to hold the silicon ribbon samples in both face-on and edge-on orientations within the SEM chamber, as shown in figure 3.5(d). To avoid damaging or chipping the ribbons it was important to ensure the silicon faces and edges were not exposed to contaminants or contacted with any hard materials. As will be discussed in section 3.4.2.6, it is acceptable to hold the samples 8 mm from either end, therefore the jig was designed to mechanically clamp these ribbons within this region.

Comparing the old and new laser cutting technique

Figure 3.8 shows typical examples representing the new and old laser cutting technique. As can be seen, on one face, there is little difference between the samples when comparing the two techniques. Typically one face exhibited more damage, in the form of micro-chipping around the edges than the opposing face. It is suspected that the face which has greater chipping is the exit side of the laser and that these chips exist due to ablation on the back side of the sample when pierced by the cutting laser. This effect did not appear to change between the 525 µm and 775 µm thick samples. Small cracks and darkened damaged areas can be seen along the top and bottom edges of the face in figure 3.8(b).

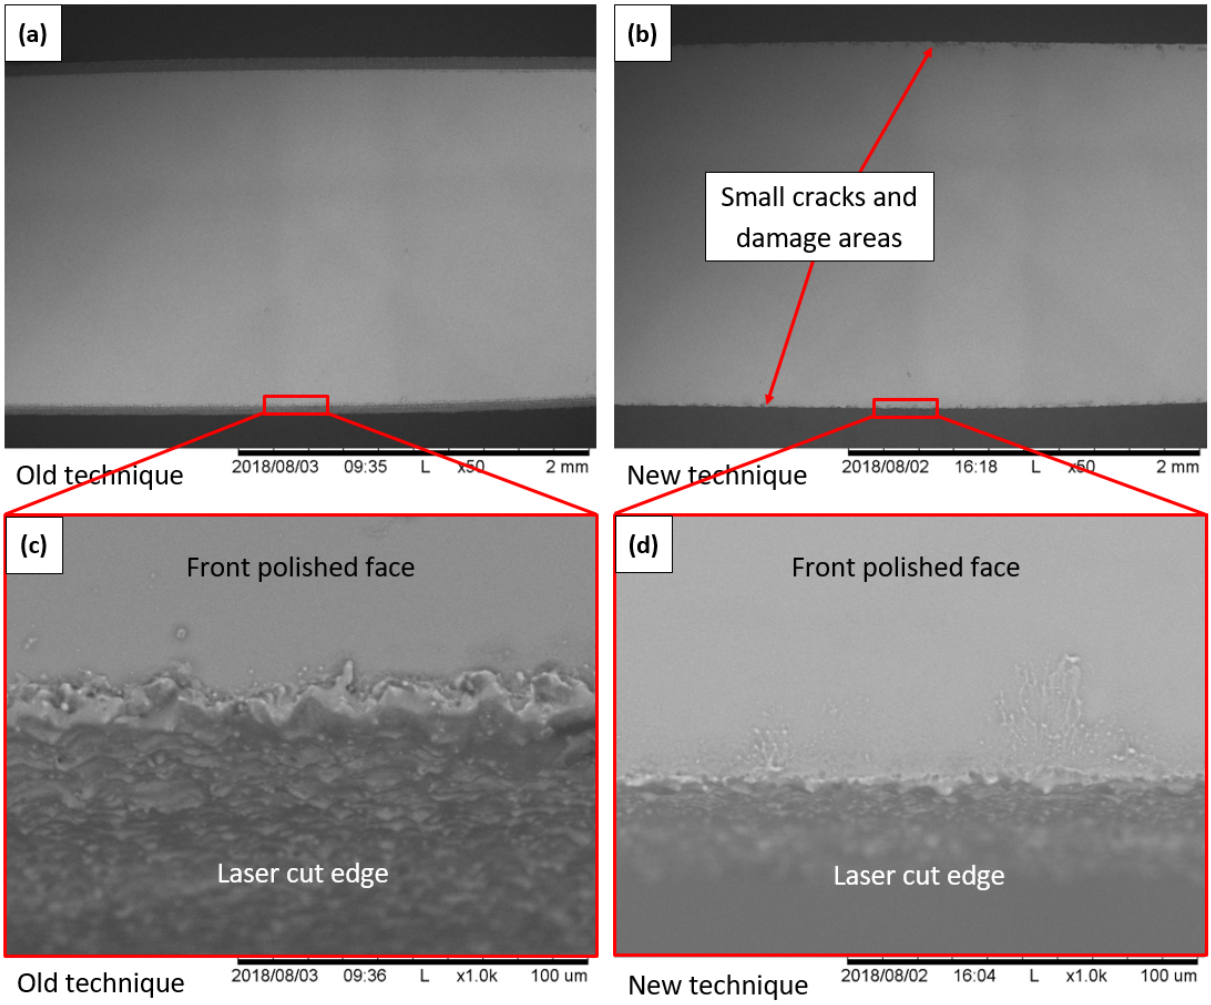

Figure 3.8: SEM magnified images of two 525 µm thick samples, face on. (a) Front face of sample 26 cut using the old laser cutting technique. (b) Front face of sample 1026 cut using the new laser cutting technique. (a) and (b) are magnified 50×. (c) Vertex of the front polished face and bottom laser cut edge of sample 26 cut using the old laser cutting technique. (d) Vertex of the front polished face and bottom laser cut edge of sample 1026 cut using the new laser cutting technique. (c) and (d) are magnified 1000×.

Zooming in on this vertex it was evident that the new laser cutting technique produced a much sharper and cleaner cut into the silicon compared to the old technique as shown in figure 3.8(c) and figure 3.8(d). The shallower cut angle, perpendicular to the polished face, using the old technique is clearly seen in figure 3.8(c) as one is able to look along the edge of the sample.

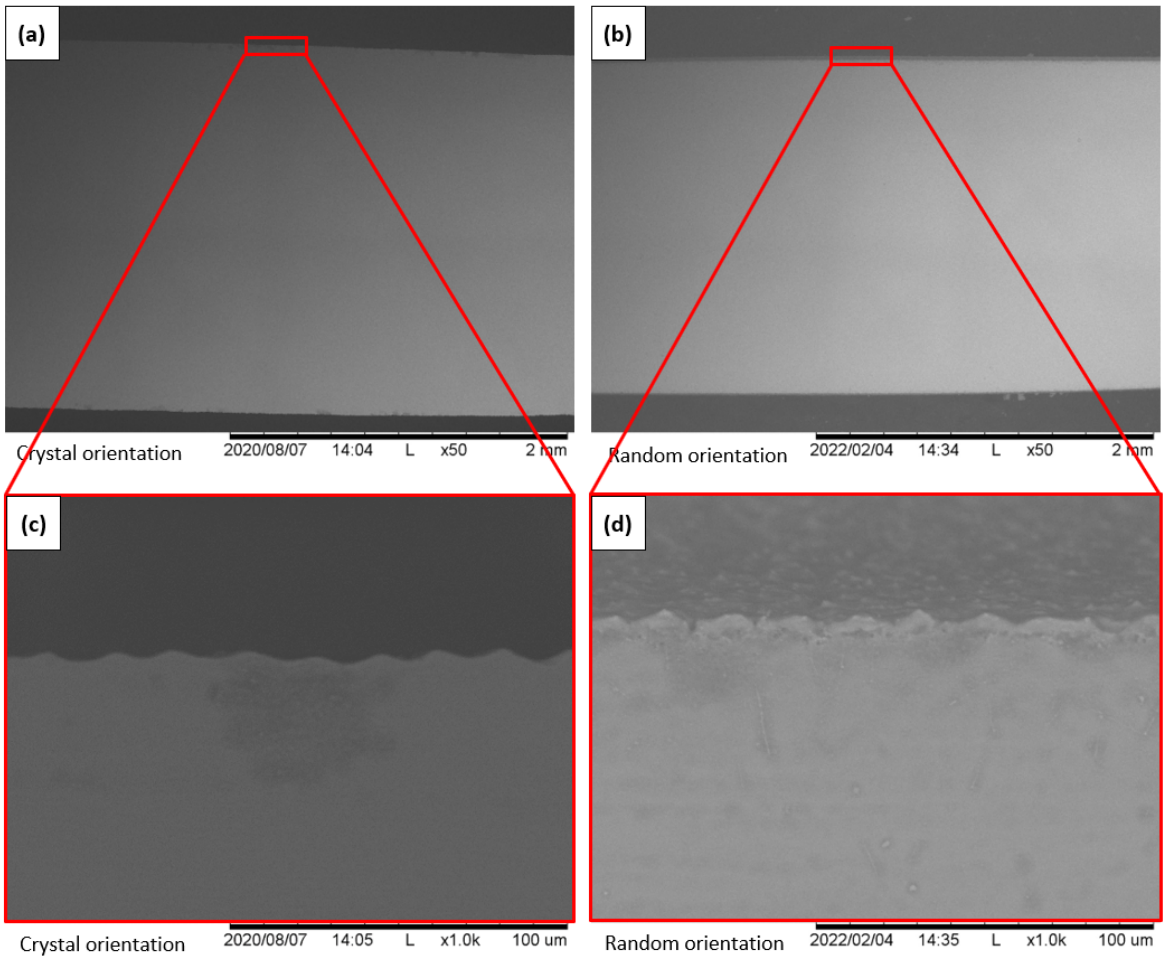

Figure 3.9: SEM magnified images of four 775 µm thick samples, face on. (a) Front polished face of sample 43c cut along the crystal axis. (b) Front polished face of sample 1 cut in a random orientation. (a) and (b) are magnified 50× (c) Front polished face of sample 43c cut along the crystal axis, focused on the top edge of (a). (d) Front polished face of sample 1 cut in a random orientation, focused on the top edge of image (b). (c) and (d) are magnified 1000×.

The features of the new laser cutting technique samples were similarly observed in the 775 µm both crystal-aligned and random orientated silicon ribbon samples as shown in figure 3.9.

Figure 3.9 also shows a tooth-like, repeating pattern along the cut edge was visible, consistent with laser pulsed cutting [264]. For 525 µm thick samples, this feature was not particularly clear on either the new or old laser cutting technique samples, although some repeating linear features are seen as shown in figure 3.10, particularly when zoomed out.

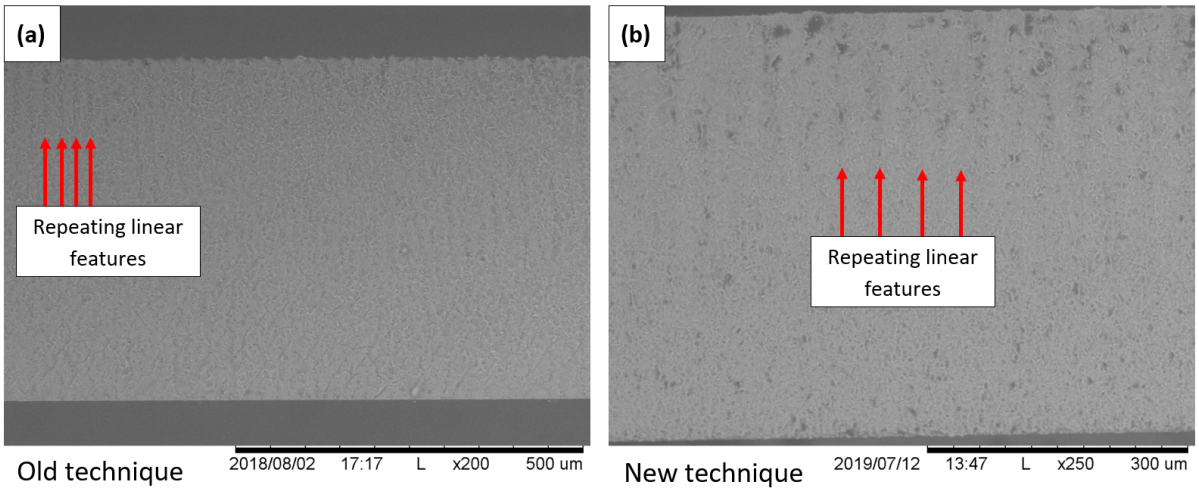

Figure 3.10: SEM magnified images of two 525 µm thick samples, edge on. (a) Bottom cut edge of sample 44 cut using the old laser cutting technique, magnified 200×. (b) Top cut edge of sample 6122 cut using the new laser cutting technique, magnified 250×.

For 775 µm thick samples, this feature was clearer, particularly on the crystal axis cut samples, as shown in figure 3.11.

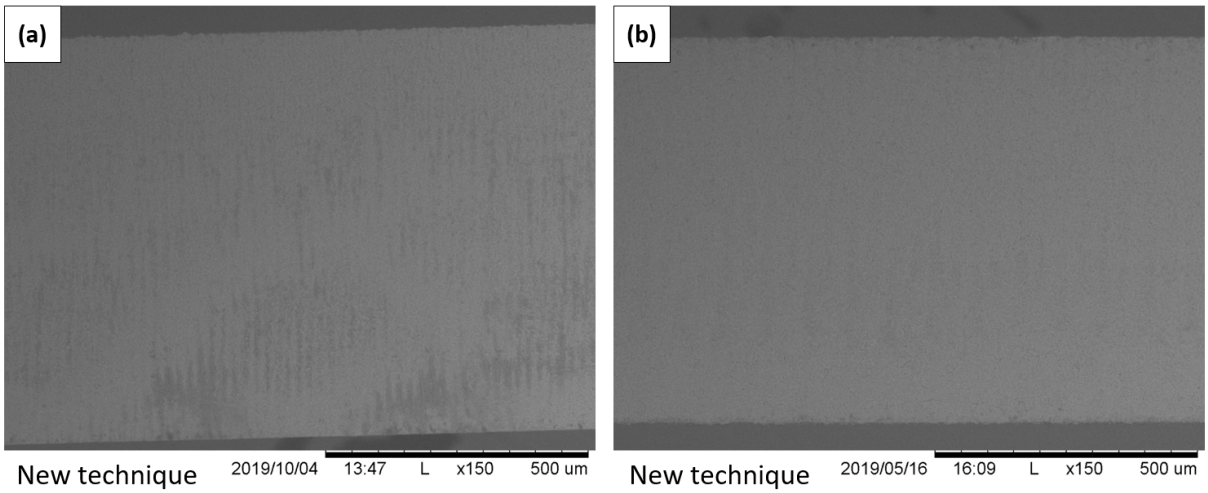

Figure 3.11: SEM magnified images of two 775 µm thick samples, edge on. (a) Top cut edge of sample 160c cut along the crystal axis. (b) Top cut edge of sample 115 cut in a random orientation. (a) and (b) are magnified 150×.

3.4.2.5 Silicon ribbon dimensional profiling

It is important to know the cross-sectional area of the sample in order to calculate the corresponding stress at breaking load, and therefore two cross-section measurement methods were used.

The first method used RS PRO 150 mm digital callipers [265] ensuring only the very ends of the sample were contacted to avoid damaging the tensile testing area. The quoted accuracy

of these callipers is $\pm 20\,\mu m$ [265]. Random error is a potentially large source of uncertainty when using callipers. It is particularly difficult to measure sample widths ensuring the callipers are perfectly at 90°. This is exacerbated by the non-rectangular cross-section of the samples evidenced in section 3.4.2.2. From repetition tests, the random measurement error was $\pm 130\,\mu m$ when measuring the width of the samples. The random measurement error when measuring the thickness of the samples was $\pm 22\,\mu m$.

The second method involved using two fibre profiling set-ups. The first set-up, referred to as the "LIGO fibre profiler", has been previously used to measure the profile of LIGO silica fibres and can provide much greater accuracy than digital callipers [144] [266] [267].

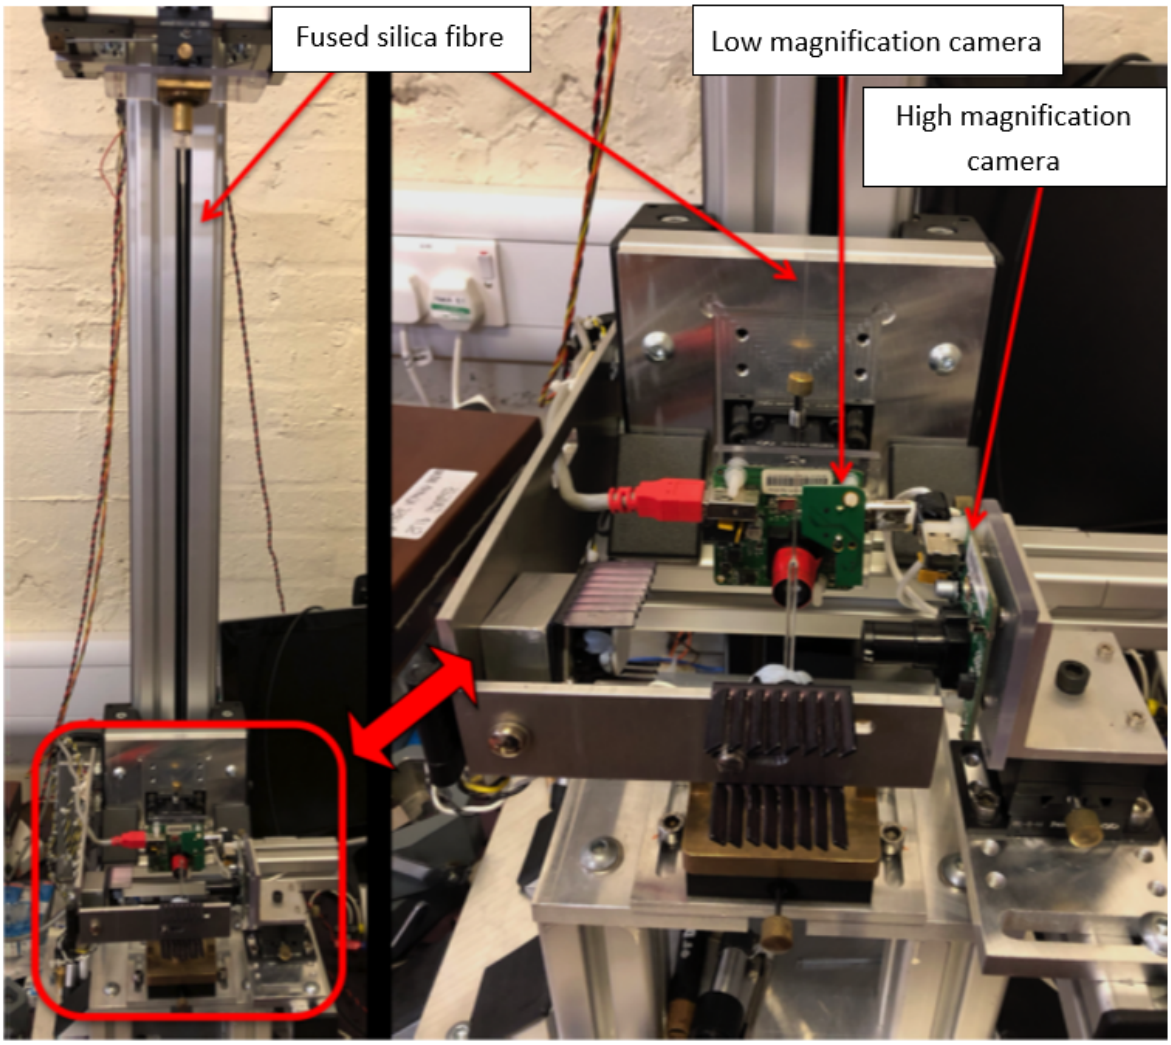

Figure 3.12: Image of LIGO fibre profiler. Image modified from [268].

The profiler uses two cameras and backing lights positioned opposite to characterise a samples dimensions through shadow sensing. An example of a fused silica fibre undergoing fibre profiling can be found in figure 3.12. The low magnification camera was used to measure the sample width and had an error of 0.25 %. The high magnification camera was used to measure

the sample thickness and had an error of 1 % [144].

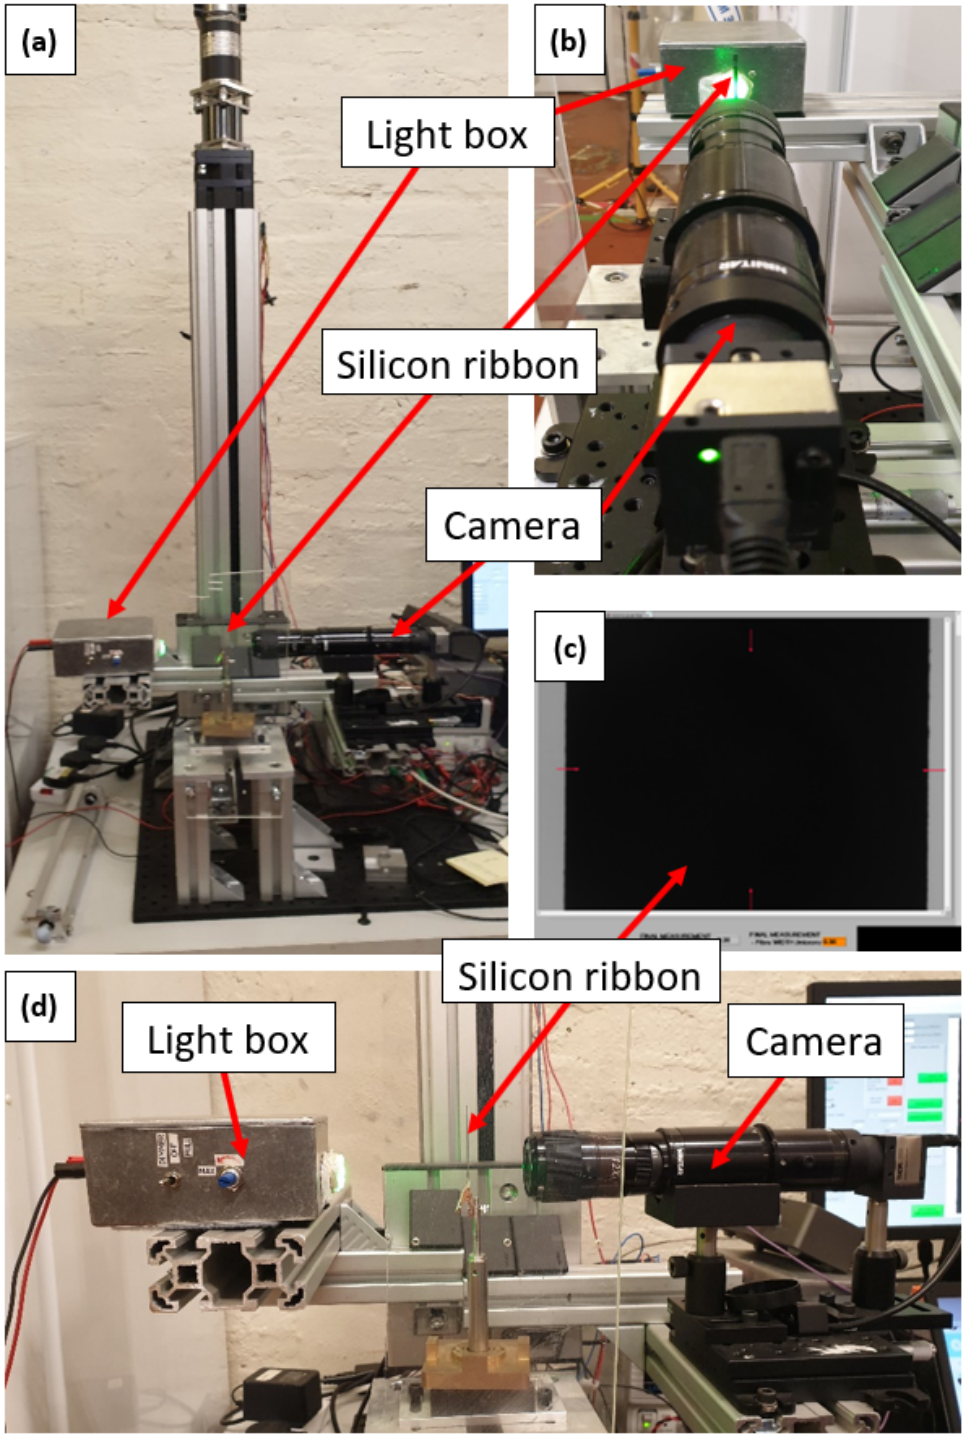

Figure 3.13: (a), (b) and (d) Photographs showing the thin fibre profiler from various angles. (c) Typical image from the profiler, showing a silicon sample (black section) being profiled.

As shown in figure 3.13, a second conceptually similar "thin fibre profiler" was developed to allow for more accurate measurement of silica fibres [227] [268]. This profiler only utilises one camera with much greater magnification, focus and control ability. The errors for this pro-

filer were determined in line with section 3.6.1.7 in [144]. The pixel error on this device was quantified by 10 repeated measurements of the same point on the sample to gauge the spread of measurements. Camera instability can contribute to pixel error. For focussing error, 10 repeated measurements were taken, this time re-focusing the camera in between each measurement, since determining when the sample is in focus is done by eye using the observed camera image, therefore experiencing unavoidable user subjectivity.

	Ribbon thickness				Ribbon width	
	775 µm		525 µm		2.2 mm	
	Pixel error	Focusing error	Pixel error	Focusing error	Pixel error	Focusing error
	791.14	790.72	513.04	513.1	2232.68	2228.65
	791.58	791.28	513.55	513.61	2232.67	2233.06
	791.79	791.36	513.47	513.62	2231.09	2232.92
	791.86	790.85	513.11	513.12	2231.12	2228.91
	791.43	791.08	513.07	513.57	2231.04	2231.07
	788.74	791.23	513.68	513.21	2230.91	2230.95
	791.53	791.19	513.54	513.16	2230.92	2230.72
	793.29	791.36	513.52	513.11	2231.04	2228.86
	791.23	790.55	513.11	513.12	2231.12	2228.58
	791.09	791.17	513.04	513.05	2230.94	2228.88
Average	791.37	791.08	513.31	513.27	2231.35	2230.26
Standard Deviation	1.12	0.28	0.26	0.23	0.70	1.74
Error (%)	**0.14**	0.04	0.05	0.05	0.03	**0.08**

Table 3.2: Data from experiment determining camera shake/pixel error and focussing error of thin fibre profiler for each ribbon dimension. The percentage error in bold is the error determined for each dimension.

Table 3.2 shows the resultant percentage errors. The larger of the pixel error or focusing error was taken as the error for that dimension. Due to the larger pixel error on the 775 µm thickness this will be taken as the conservative upper estimate on the error for the thickness for both samples since camera instability is possible for every measurement. The focusing error was largest on the width measurement.

3.4.2.6 Tensile strength testing set-up

With the samples sufficiently characterised, testing on a series of control samples for their ultimate tensile strength took place.

A sample was loaded into the tensile strength testing mechanism. A low-geared, high torque mechanism then pulls the sample apart through its longitudinal axis at a constant rate. The testing mechanism has a load cell able to measure the peak force at the point of failure.

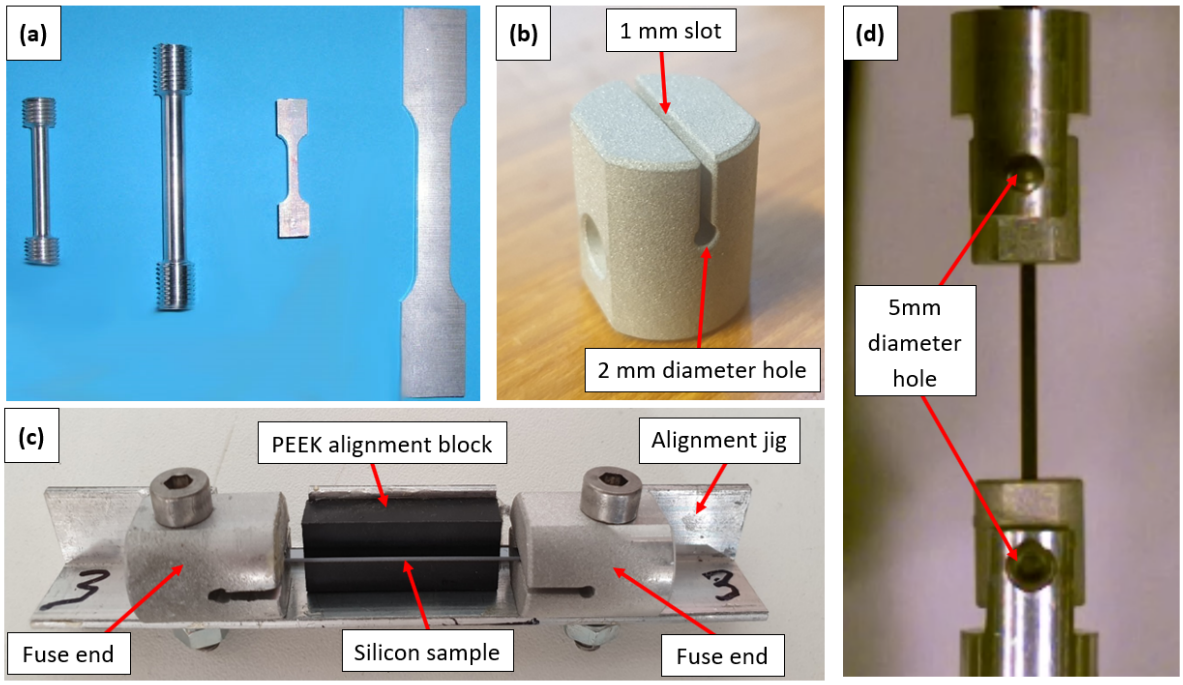

Figure 3.14: (a) Example of typical metal specimens designed for tensile strength testing [269]. (b) Fuse end used to hold silicon ribbon samples for tensile strength testing. (c) Silicon ribbon curing inside two fuse ends placed into the alignment jig. (d) Cured silicon ribbon fuse end set-up ready for tensile strength testing.

The standard methodology for testing material is to use a shaped dumbbell form in order to allow for easy gripping, and to constrain the failure area to the exposed, well-defined thinner cross-sectional area as shown in figure 3.14(a).

Due to the brittle properties of silicon at room temperature, it is extremely difficult to hold samples by mechanical clamping onto their polished surfaces. This is a similar problem faced in previous work conducted to strength test silica fibres that also had to be secured without damaging the stock material at the clamping point [144] [268]. A similar solution to this problem was implemented for the silicon samples by use of the "fuse end" concept previously employed in [144] [268] [270]. This is a 6082-T6 aluminium holster that can hold silicon with epoxy without damaging it, as shown in figure 3.14(b). Since the samples were affixed by adhesive, the inner slot of the fuse end was sandblasted to roughen the surface, with a 2 mm $\times$ 2 mm diameter hole also drilled at the slot base in order to aid the gripping of the adhesive. Two part Araldite 2012 adhesive was used, as in previous work, as it is known to have sufficient properties for strength testing [144] [271].

To attach a sample, the epoxy was applied into both fuse end slots using a mixing nozzle. The silicon sample was inserted into both fuse end slots. It is also important to align the sample so that tension is applied axially along the sample. As shown in figure 3.14(c), this was achieved by the use of a fuse end alignment jig.

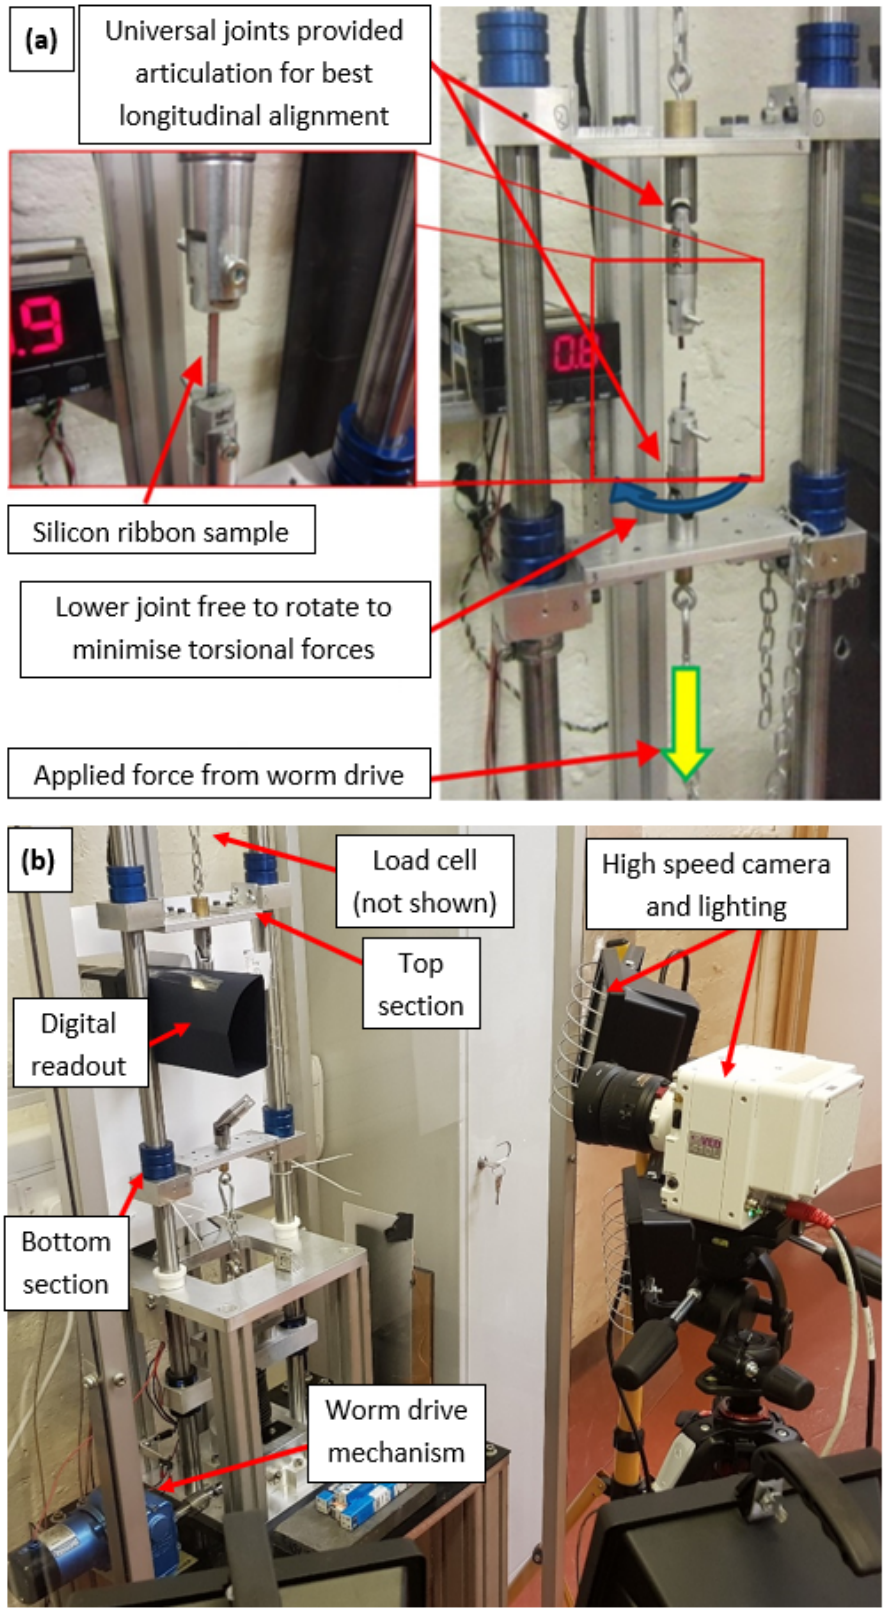

Figure 3.15: (a) Articulation enabled by the use of universal joints as well as the affixing bolts that hold the sample in place, image taken and adapted from [136]. (b) Digital readout alongside the top and bottom sections, including high speed camera and lighting. In this image a sample is hung from the top section but the bottom section has not yet been connected.

Immediately after epoxy application the sample assembly was inserted into the alignment jig and gently pressed against the alignment block using PEEK tipped tweezers. This block was also made from PEEK to minimise any damage to the sample edge. The 90° shape of the alignment jig helped align the sample and fuse ends in all three axes. Finally the fuse ends were secured with two retaining bolts while the epoxy was left to cure for at least 24 hours at room temperature. A completed sample ready for testing is shown is shown in figure 3.14(d).

Figure 3.15 shows the tensile strength testing set-up [272]. The top section of the tester is fixed by a chain to an Omega DP25B-S-DC10 load cell with a quoted error of $\pm 0.03\,\%$ of reading, ± 1 count [273]. The bottom section, consisting of a sliding plate on linear rails, is held only by the tensile strength sample. Both top and bottom sections are attached to the sample using universal joints, that allow for articulation to further minimise any longitudinal misalignment that may cause premature fracture due to off-axis forces. The top and bottom sections are free to slide up and down on the rails. The sample is affixed into the set-up by placing a bolt through the holders and fuse ends on the top and bottom sections of the tester. Once the sample is placed in the tester, load is applied by an electric motor with a worm drive and reduction gearbox. A high speed camera with sufficiently powerful lighting was also used to record breaking specimens [274]. The camera was set to a resolution of 512×384 pixel with a frame rate of 23000 frame per second and an exposure time of 18 µs. The aim was to conduct image analysis on the samples in order to correlate tensile stress failure points to catalogued defects. A ruler was attached to one of the rails to aid in this analysis.

3.5 Tensile strength testing control samples

Control name	Thickness (µm)	Laser technique	Crystal orientation	No. of samples
Old	525	Old	Along axis	17
New	525	New	Along axis	20
Crystal	775	New	Along axis	20
Random	775	New	Mix of off-axis/along axis	20

Table 3.3: Control group names and parameters.

A full set of control sample groups were established for each sample parameter, as shown in table 3.3.

3.5.1 Results and error analysis

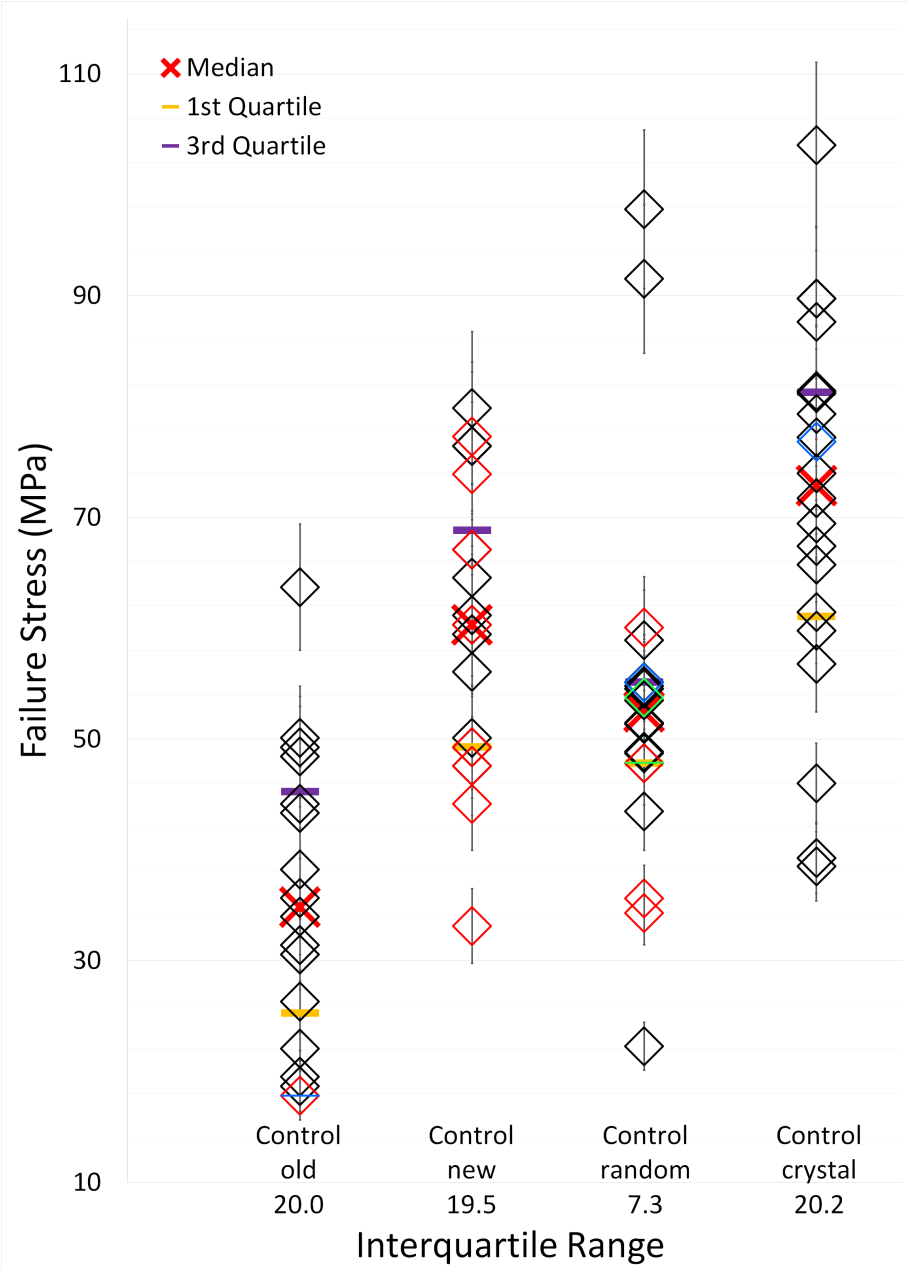

Figure 3.16: Results of tensile strength testing on four control sample groups. ◇ indicates sample failed at fuse end interface. ◇ indicates a sample was noticeably misaligned in the fuse end assembly. ◇ indicates the sample pulled free from a fuse end due to epoxy failure so the sample was re-epoxied into a new fuse end and re-tested. Where more than one anomaly exists on a sample, the respective colours are applied to the data marker.

The tensile strength testing results of the four control sample groups are shown in figure 3.16.

For all control set dimensions, it was decided to check a random subset of 5 samples from the batch of 20. LML quotes a laser cutting tolerance of $\pm 50\,\mu m$. Since the digital callipers have an error of $\pm 130\,\mu m$, it was deemed if the subset was within $2.20\,mm \pm 130\,\mu m$ that this width

and the corresponding error be applied to all samples. For thickness of the samples, the wafer manufacturing tolerance was ±25 µm, so the upper error of manufacturing tolerance, ±25 µm, was applied to both wafer thicknesses.

It was noted under no loading that the digital readout would occasionally flicker between readings by 0.1 kg, so the load cell error was taken to be ±0.1 kg as a conservative estimate over the manufacturer stated error given in section 3.4.2.6.

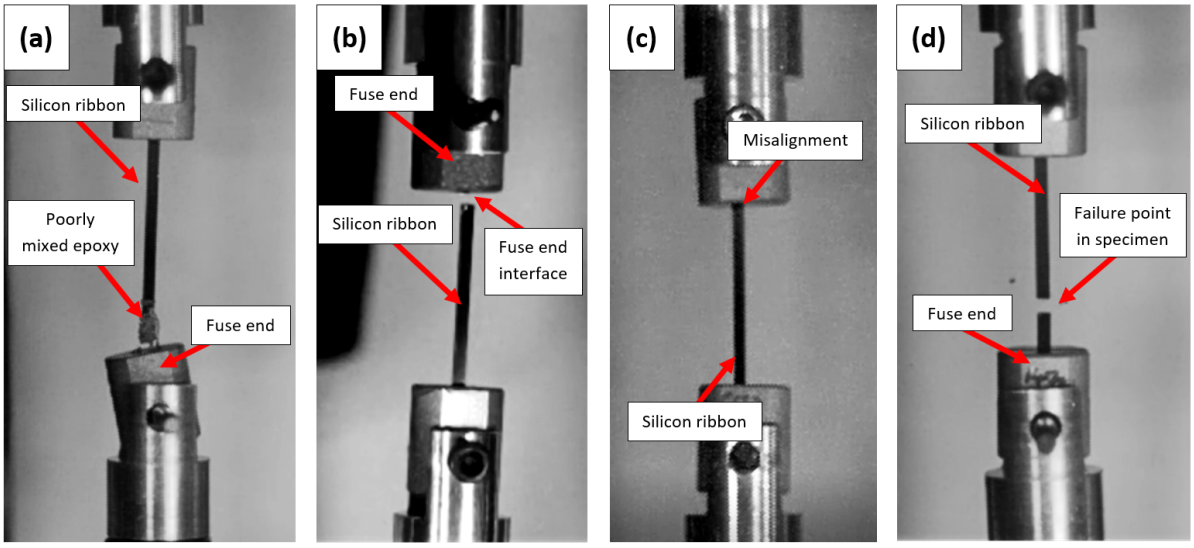

Figure 3.17: (a) Example of epoxy failure. (b) Example of a failure along the fuse end interface. (c) Example of sample misalignment. (d) Example of an ideal tensile failure. These images are frames taken from the high speed camera footage.

During testing, some samples broke at lower strength than expected or did not break at all. Occasionally the Araldite epoxy would not set properly. This has been observed before and can be attributed to poor or incomplete mixing of the two-part epoxy with curing failure only becoming apparent during tensile strength testing. When observing the load cell digital readout, the load value would rise, stabilise, then fall as the sample slipped out of the epoxy as shown in figure 3.17(a). These samples were recovered and where the end of the sample had slipped out, a new fuse end was epoxied onto it using the alignment jig.

It was noted some samples occasionally failed across the interface where the sample was epoxied into the fuse end, as shown in figure 3.17(b). It is worth noting when this happens as the failure mechanism may be different compared to when the sample fails purely in the silicon.

Finally, some samples were found to be misaligned when removed from the alignment jig, either due to initial misplacement or movement during the epoxy setting process, as shown in figure 3.17(c).

In figure 3.16 and all strength testing figures, samples failing across the fuse end interfaced are marked with a red diamond. Where a sample was noticeably misaligned, this is marked by a blue diamond. Finally, for a sample has slipped out of a fuse end but was re-epoxied into a new

fuse end, a green diamond has been added.

3.5.2 Discussion

The median strength of the control samples produced using the new laser cutting technique was 60.3 MPa, 73 % stronger than the samples cut using the old technique. Due to the proprietary processes, it is difficult to tell what improvements have been made in the new cutting technique. It was observed that the angle of cutting into the silicon creates a more rectangular cross-section which will reduce vertex angle and may decrease the effects of a vertex stress concentrator. This is a possible explanation for the increased strength which was observed. The error in width measurement, discussion in section 3.4.2.2, arising from the non rectangular shape of samples was found to be negligible when compared to the combined strength testing error.

Another possible explanation is that the new process led to smoother surfaces, which have smaller crack sizes and tend to yield higher failure stress. However, as discussed in section 3.4.2.3, it was not possible to measure surface roughness.

	Control old	Control new	Control random	Control crystal
Median strength (MPa)	34.8	60.3	52.5	72.8
Critical crack length, c_c (µm)	168.8	56.2	74.2	38.6

Table 3.4: Results of median strength for each control set and corresponding critical crack length.

Referring to equation 3.8, one can calculate a crude estimate of the critical crack size on a sample from the median values of the control sample sets, as shown in table 3.4. This is a useful test to see if critical crack sizes are on scales consistent with surface features observed during characterisation. The silicon surface energy, γ was taken as $1.9\,\mathrm{J\,m^{-2}}$ [275] across the <110> plane through which most samples broke. This value was also used for random orientation samples since it was not possible to obtain the surface energy of crack surfaces through a random, off-axis path in the silicon samples. It is likely that the surface energy will be higher for these *control random* sample sets when breaking across non-crystallographic planes, so this critical crack length value, c_c, is likely an underestimate [276]. A value of 169 GPa was used for the Young's Modulus along the <110> axis [174].

The results in table 3.4 appear to be realistic estimates of surface damage size on sample edges. This can be further interrogated by reversing the applied method, taking surface damage observed in a sample and calculating the corresponding failure stress, σ_f.

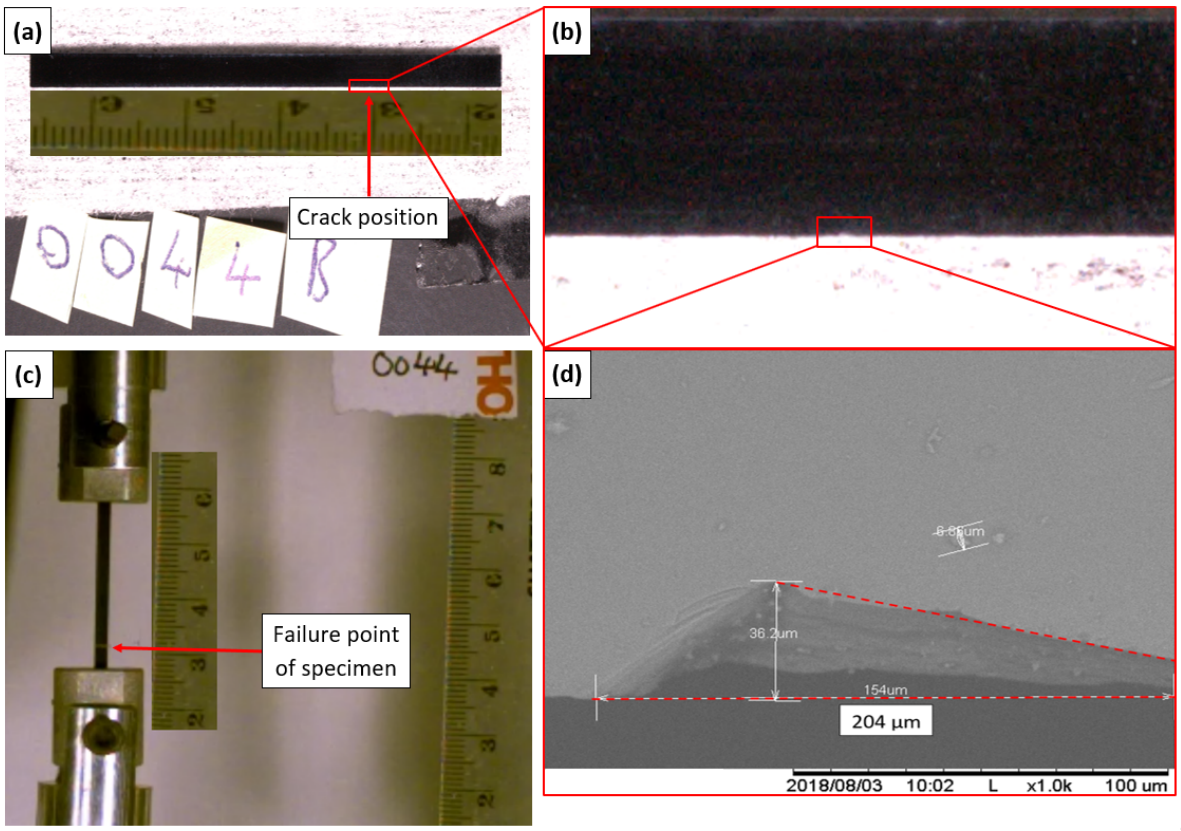

Figure 3.18: Example of a damaged area on a tensile silicon ribbon sample. (a) Photograph of the back face of sample 44. (b) Magnification of the bottom edge of (a). (c) Point of failure during tensile testing of sample 44. (d) SEM measurement and extrapolation of the damage dimensions on the back face of sample 44.

An example of surface damage on a sample is shown in figure 3.18. This crack was photographed, as shown in figures 3.18(a) and (b). Due to the very narrow field of view and mobile stage limits of the SEM, damage had to be in the central portion of the sample in order to be imaged. By chance, sample 44, from the *control old* sample set, had a damaged area on the bottom vertex of the sample's back face near the edge of this region. This was measured on the SEM top the edge of the imaging region and predicted damage geometry determined from extrapolation as shown in figure 3.18(d). The was also able to be correlated with imagery from the high speed camera, shown in figure 3.18(c) indicating this crack to be the likely cause of failure of the sample.

Crack length (µm)	Predicted failure stress, σ_f (MPa)	Measured failure stress, σ_m (MPa)	% difference
36.2	75.1	31.4	82 %
204.0	31.6		1 %

Table 3.5: Comparison of sample 44 predicted and measured failure stress from crack lengths shown in figure 3.18(d).

This damaged geometry is complicated compared to idealistic cracks in theory, so a simplistic calculation using the crack depth into the surface, 36.2 µm, as well as the predicted length of the crack along the vertex, 204.0 µm, were used as indicative values to calculate the predicted failure stress of the sample. The results are shown in table 3.5. Looking at Griffith's theory in section 3.2.2, it is more valid to use 36.2 µm, as this represents the crack distance perpendicular to the damaged surface. This yields a failure stress 82 % higher than the measured stress and can not be reconciled even with the ±3.2 MPa error applied. It is likely coincidence that the 204 µm result is so close to the measured failure stress since Griffith's theory does not use the length of the crack perpendicular to the crack direction. It is unlikely that the sample failed elsewhere, at a different crack or damage area, since it was possible to correlate the damaged region to the high speed camera footage sample failure point. However, one can see that the *magnitude* of stresses observed in figure 3.16 and in this critical crack calculation are consistent with theory.

One expects the *control new* and *control crystal* samples to yield similar median values due to the fact the only difference between the samples is the sample thickness. The *control new* and *control crystal* median values were within 15 % of each other which suggests this is indeed the case.

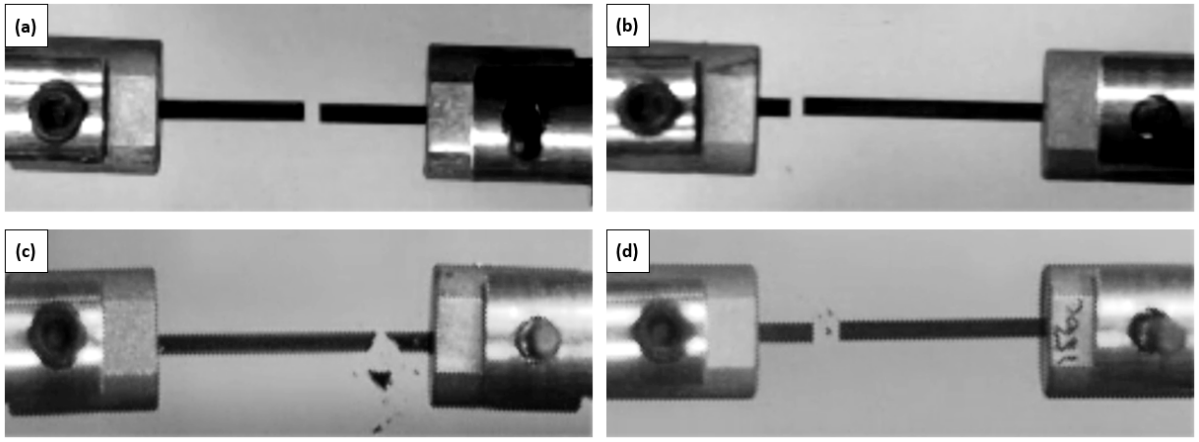

Figure 3.19: Freeze frame from high speed camera moments after fracture of samples. (a) Control old sample 3 breaking along the <110> plane. (b) Control new sample 1025 breaking along the <110> plane. (c) Control random sample 105 breaking off-axis. (d) Control crystal sample 156c breaking along the <110> plane.

The *control random* samples had some significant outliers, which may be explained by the random crystal orientation of the samples. All samples in the *control old*, *control new* and *control crystal* groups, sixty in total, fractured along the <110> plane, with examples shown in figure 3.19(a), (b) and (d). While previous work concluded that crystal orientation of the silicon sample is not a dominant influence on strength [136], this referred to the comparison of samples cut along the crystal axes of (100) and (110) wafers, rather than samples that are randomly off-axis. It is known silicon cleaves most easily along its crystal axes [276]. For the *control random* set, samples broke in both the <110> crystal axis plane and off-axis, the latter is shown

in figure 3.5(c). The strongest sample, 114, registered at 82 % stronger than the *control random* median value of 52.5 MPa. However, the second outlier and next strongest sample broke along the crystal axis. For the other six samples that broke off-axis, four broke at a stress above and two broke below the median value. The results are inconclusive for correlating off-axis samples to strength and this should be repeated with a sample set consisting entirely of silicon ribbons cut off-axis.

Failure of samples along the fuse end interface did not show any particular trend in the data. As such, no correlation was found between samples that broke at the fuse end interface and their strength.

For the few samples that were re-epoxied there is no evidence to suggest that this process weakened the samples. Similarly for samples that were bonded squint for the control group, there appeared to be no correlation between this misalignment and the corresponding stress, although the sample size for this was small.

It is interesting to note the interquartile range (IQR) is broadly similar for the *control old*, *control new* and *control crystal* axes which are all cut along the crystal axis of the wafer. The *control random* set has a much smaller IQR which can not be readily explained by crystal axes orientation; as discussed above, random and crystal orientation samples within this set did not appear to follow any particular strength pattern or trend.

3.6 Chemomechanical polishing

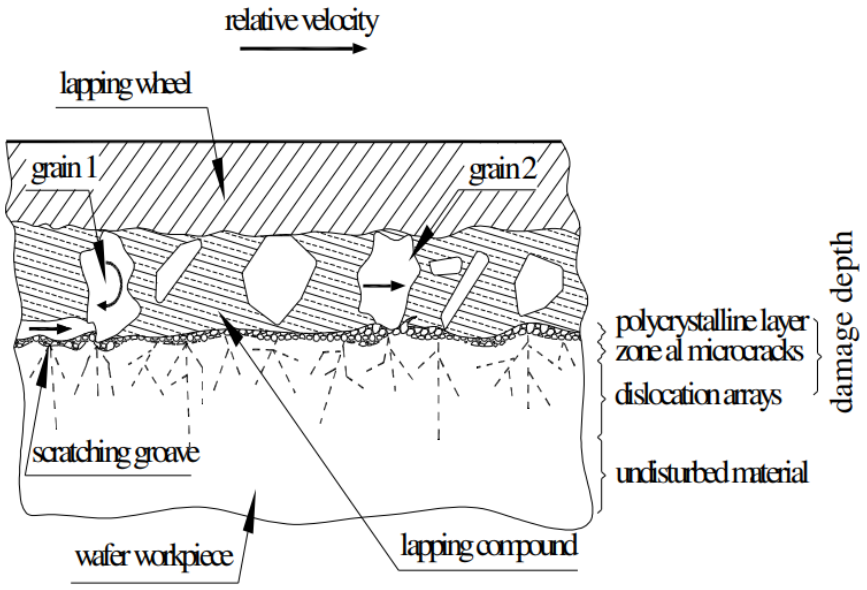

Figure 3.20: Schematic of lapping slurry particles abrading the surface of a silicon substrate. Image obtained from [277].

One method for improving the surface of silicon by reducing surface roughness is chemomechanical polishing (CMP). This is a well-known process for polishing single-crystalline silicon wafers to provide sufficient surface quality for the purpose of producing integrated circuits on mass industrial scales [278] [279].

All samples were laser cut from wafers that were chemomechanically polished on both faces. This produces the near featureless polished faces of the silicon samples as measured on the surface profiler in section 3.4.2.3. The same type of treatment was therefore applied to the laser cut edges of a batch of samples.

Typically, CMP occurs in two steps, lapping followed by polishing. Lapping involves preparing the silicon surface by removing the surface damaged layer with abrasive particles in the form of a slurry, working at macroscopic scales to remove roughness via mechanical abrasion [277], as shown in figure 3.20.

Polishing involves applying a colloidal solution to the lapped surface with mechanical motion to increase the substrate surface layer's reactivity towards water. The colloidal solution typically contains nanoparticles of silica, approximately 40 nm in size [280], covered with reactive OH or O^- groups. These groups act as attraction points leading to temporary Si-O-Si bonds, chemically removing the surface material molecule by molecule, as shown in figure 3.21 [281]. This allows typical surface roughness, R_q, of single figure Å to be achieved, far below any surface roughness caused if mechanical abrasion from the suspension particle size was directly involved [282]. This polishing process has the advantage of both removing the surface damage induced by lapping, which should increase the ultimate tensile strength, combined with finishing or smoothing any remaining stress concentrators on the molecular level which may also aid in keeping surface loss low [283].

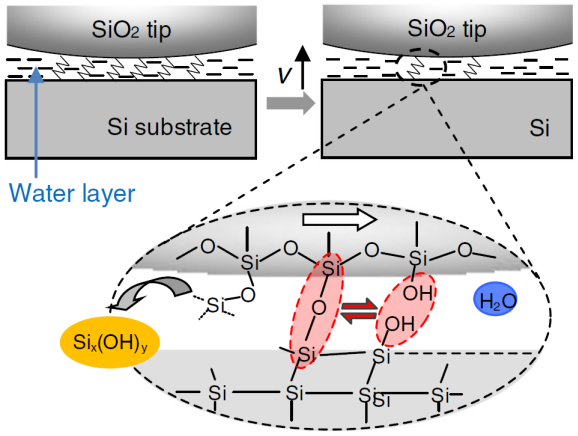

Figure 3.21: Schematic of polishing process showing the formation of molecular silanol, Si-O-Si and OH groups between the silicon substrate and silica nanoparticle, allowing for the chemical, non-abrasive removal of surface silicon atoms. Image obtained from [284].

3.6.1 Logitech PM5 polishing machine

A Logitech PM5 lapping and polishing machine [285], shown in figure 3.22 was used for this surface treatment method. To ensure the best polish, the PM5 was fully overhauled and re-set up from scratch to ensure it was working optimally. This included conditioning of the lapping and polishing plates. Part of conditioning the system involves measuring how convex or concave the lapping plates are with a flatness monitor. A conditioning mass, which can be moved across the plate width, is then used to shape the plate itself to the flatness desired. The cast iron smooth lapping plate was measured to be concave by 2.5 µm and so the plate was reconditioned until it was flat.

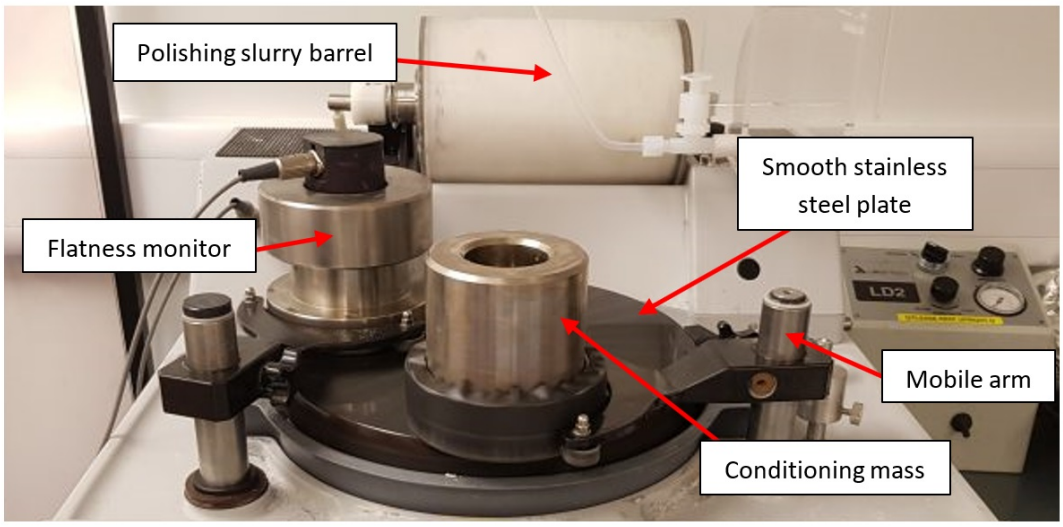

Figure 3.22: Set-up of the Logitech PM5 polishing machine showing conditioning of the smooth stainless steel lapping plate with 9 µm aluminium oxide slurry.

Conditioning plates before lapping or polishing is always necessary as it also prepares the surface by impregnating it with the lapping or polishing slurry as well as removing contaminants. The slurry barrel at the top is rotated on rollers in order to provide a constant, fresh supply of lapping solution onto the lapping plate.

To lap or polish samples, the Logitech PM5 is supplied with a PP5 jig, shown in figure 3.23. Samples can be waxed onto a sample holder, on the underside of the PP5 jig, shown in the inset of figure 3.23. When the jig is then orientated upright, gravity, combined with a central pressure spring located axially within the PP5 jig, applies pressure to the samples where they contact the lapping or polishing plates on the bed of the machine. The PP5 jig is supported by a metal ring around its base, which is also exposed to the lapping or polishing process. Channels can be seen in the metal ring which allow for the inflow of lapping slurry or polishing fluid.

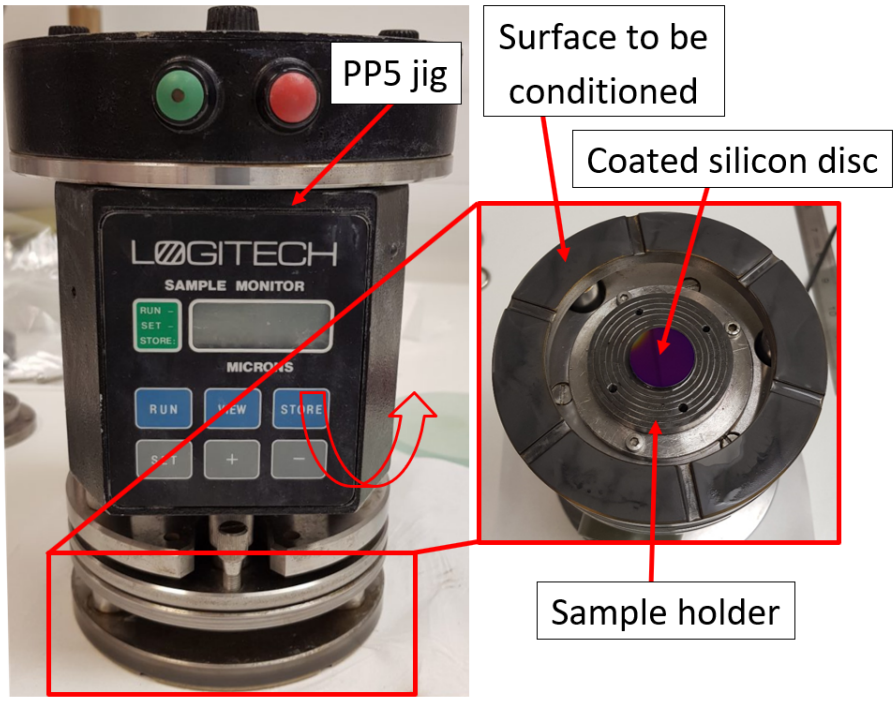

Figure 3.23: PP5 jig supplied with Logitech PM5 polishing machine. Inset image shows the metal ring, with 6 channels, that contacts the lapping or polishing plate and the sample holder bolted onto the jig. In this image, a coated silicon disc is attached onto the sample holder.

Conditioning of the PP5 jig itself is as important as conditioning the lapping or polishing plates. Before any lapping or polishing was conducted, the system would always be run with the appropriate lapping or polishing plate and slurry as shown in figure 3.24. The sample holder would be affixed to the PP5 jig but no samples were attached in order to condition the metal ring shown in figure 3.23 and expose the sample holder area to the slurry and remove any contaminants.

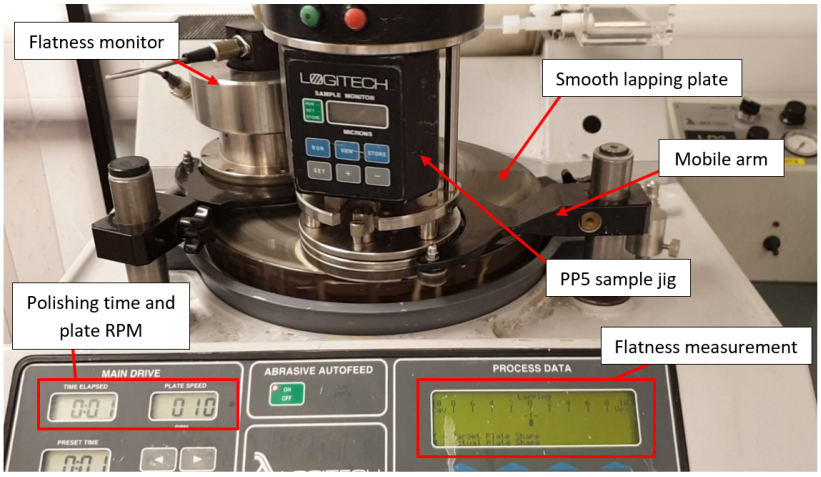

Figure 3.24: Still image from a video showing the PP5 jig and cast iron smooth plate undergoing conditioning with 9 µm aluminium oxide lapping slurry. No silicon samples are affixed to the sample holder at this stage. This is the standard lapping configuration.

In order to keep the plate flat, a mobile arm held the PP5 jig in one position and occasionally adjusted the position across the plate during conditioning to ensure even wear of the samples and plate as well as stopping the plate from becoming too concave or convex. The mobile arm also has rollers to allow the PP5 jig to freely rotate around its vertical axis.

It is vital the equipment is thoroughly cleaned before lapping and especially before polishing as any remaining lapping material could inadvertently scratch a polished surface.

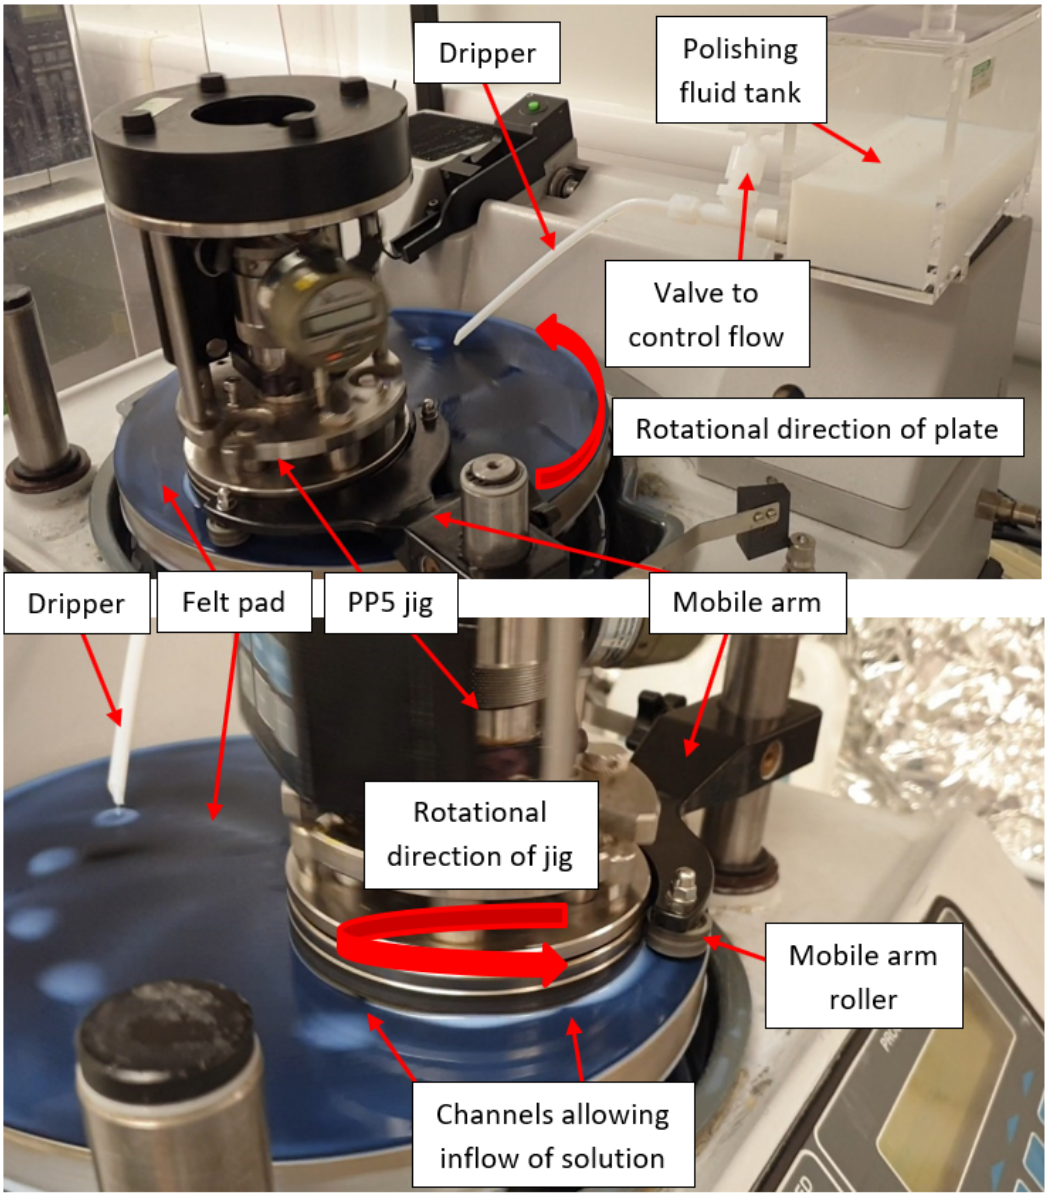

Figure 3.25: Still images taken from a video of the Logitech PM5 set-up in a polishing configuration using the felt pad. Note the flatness monitor is removed and the polishing solution is applied via a box and dripper system mounted on the top right of the machine. The direction of rotation of the PP5 jig and plate is also highlighted.

For polishing with colloidal silica, a flat aluminium plate is used with a felt polishing pad affixed to the plate via adhesive, as shown in figure 3.25. Conditioning is required for polishing

for the purposes of pad impregnation and contaminant removal, however no flatness monitoring is conducted. This is due to the fact polishing does not mechanically abrade the polishing pad or deform the underlying plate, so the plate always remains flat. Also any non-smooth (i.e. non-lapped) surfaces, such as the flatness monitor's diamond impregnated support ring, would tear and significantly damage the felt pad. Unlike lapping, the moveable arm constantly sweeps the PP5 jig backwards and forwards across the polishing plate in order to supply fresh colloidal solution and remove used fluids. As shown in 3.25, instead of a slurry barrel providing a constant feed of fresh solution, a plastic box filled with polishing fluid is used, complete with a dripper tube and valve to control flow.

3.6.2 Chemomechanical polishing surface treatment

For all CMP work conducted, the lapping material was a 9 µm aluminium oxide powder [286]. This was made up in line with manufacturer recommendations, 15 % by volume of powder, by adding 150 ml of powder to 1.5 L of deionised (DI) water to create the lapping slurry [287].

For polishing, the colloidal solution used was Syton SF1 polishing fluid [288]. A measurement of the pH of the solution was conducted by fellow PhD student Ms Victoria Graham, which resulted in a pH of 9.4.

3.6.2.1 Chemomechanical polishing treatment #1

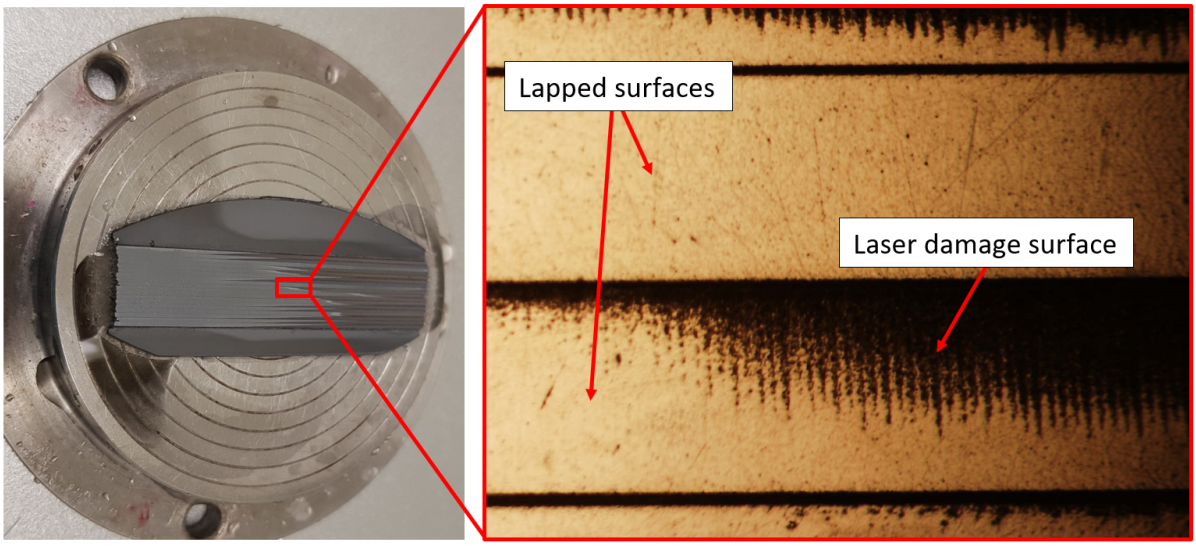

Figure 3.26: Offcut samples after 2.5 hours of lapping at 50 RPM with 9 µm aluminium oxide slurry. Inset shows Olympus optical microscopy [289] of lapping interface where one can see the beginning of the removal of the damaged surface.

For the first CMP surface treatment experiment, the samples were to be lapped by 150 µm on both edges, consistent with the target removal amount stated in section 3.2. The samples were

then to be polished to have a similar surface roughness of the polished sample faces, $\approx 5\,\text{nm}$, R_a.

Preliminary investigation conducted on offcuts yielded a significant underestimate of lapping rate. This was thought to be due to initial misalignment of the sample holder on the PP5 jig with resulting polish shown in figure 3.26. This experiment did however allow for the inspection of the lapping boundary, where one can see the removal of the damaged surface to the left hand side of the image shown in the inset of figure 3.26.

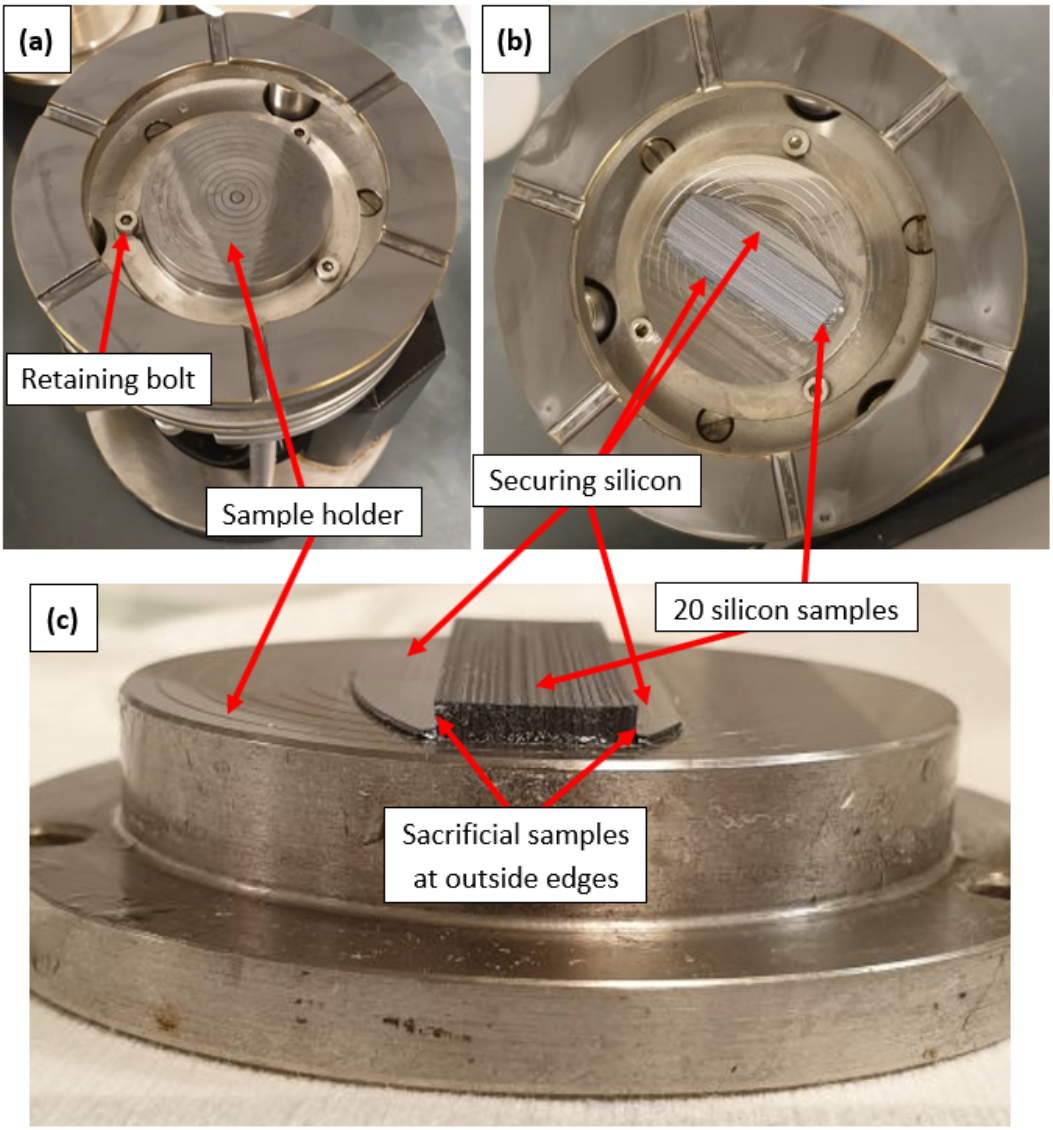

Figure 3.27: (a) Sample holder undergoing lapping to achieve flatness. (b) Samples mounted onto the sample holder which is bolted onto the PP5 jig. (c) Samples waxed onto the sample holder.

After further investigation, it was discovered the sample holder itself was found to be misaligned, with the top and bottom surfaces not parallel. Using retaining bolts, the sample holder was aligned to be flush with the PP5 jig metal ring using a wetting technique. By placing a

glass slide over the metal ring, one can adjust the alignment of the sample holder by wetting it and contacting it with the glass slide. When one retracts the sample holder into the body of the jig, the wetting pattern indicates areas of misalignment. Once best aligned by this technique, the PP5 jig, complete with sample holder was conditioned and lapped for 10 minutes and then imaged as shown in figure 3.27(a). It was noted this time was insufficient to suitably flatten the sample jig, so the system was further lapped and conditioned for total time of 20 minutes at 50 revolutions per minute (RPM).

20 samples from the new laser cut technique were treated, the same as the *control new* samples. Before undergoing treatment, all samples were subject to a pre-treatment clean consisting of a high quality (>99.9 % purity) acetone flush followed by a wipe. The sample was then rinsed in DI water and wiped again. Finally, a (>99.9 % purity) methanol flush followed by a wipe rendered the sample clean. This pre-treatment cleaning method was developed and verified via a number of tests conducted on a DSA30B drop shape analyser [290]. Results showed an increase in hydrophilicity of the silicon surface and hence cleanliness.

The samples were then secured onto the sample holder using paraffin wax, melted by placing the assembly on a hotplate, exposing the top laser cut edge of each sample to the first lapping procedure. Two larger offcuts of silicon were used as securing pieces for the sample assembly, shown in figures 3.27(b) and (c). From initial experiments, it was noted the outermost samples tended to suffer significant chipping, likely due to lack of sufficient slurry replacement causing silicon chips from lapping to travel around one full revolution and damage the edges of the samples. To mitigate this, two sacrificial samples of the same dimensions were placed on the outermost assembly to protect the 20 samples, shown in figure 3.27(b) and (c). Once the samples were fully cooled (typically overnight), surplus wax was removed by liberally applying acetone to the sample holder and rubbing a wipe back and forward over the exposed edge of the samples.

The sample holder was attached to the PP5 jig which was then placed onto the polishing machine, set up in the lapping configuration shown in figure 3.24. The silicon samples were lapped with 9 µm aluminium oxide slurry for 2 hours at 50 RPM. Lapping was then stopped in order to estimate material removal rate. The sample holder was removed and DI water thorough rinsed over the samples and sample holder to remove any lapping slurry. By using callipers to measure the difference in height between the thickness of the sample holder and the sample edges, it was discovered approximately 1100 µm has been removed from one side, approximately half the entire sample width. This equated to an approximate lapping rate of $9\,\mu m\,min^{-1}$. It should be noted the calliper error is significantly smaller than discussed in section 3.4.2.5, and was taken and verified as the quoted error of $\pm 20\,\mu m$ due to sample geometry and measurement repeatability.

As shown in figure 3.28(a), a smooth, uniform grey surface was observed across the entire lapped surface. Chipping was noted on the sacrificial sample outer vertices, highlighted in the inset of figure 3.28(a), showing their effectiveness at protecting the 20 surface treated samples.

While this sample had been over-lapped by 7 times the intended amount, there was no reason to discard this sample set. The next step was to polish the samples.

With no sample holder affixed, the PP5 jig was placed onto the polishing machine, set up in the polishing configuration shown in figure 3.25. The PP5 jig was conditioned for 10 minutes at 30 RPM.

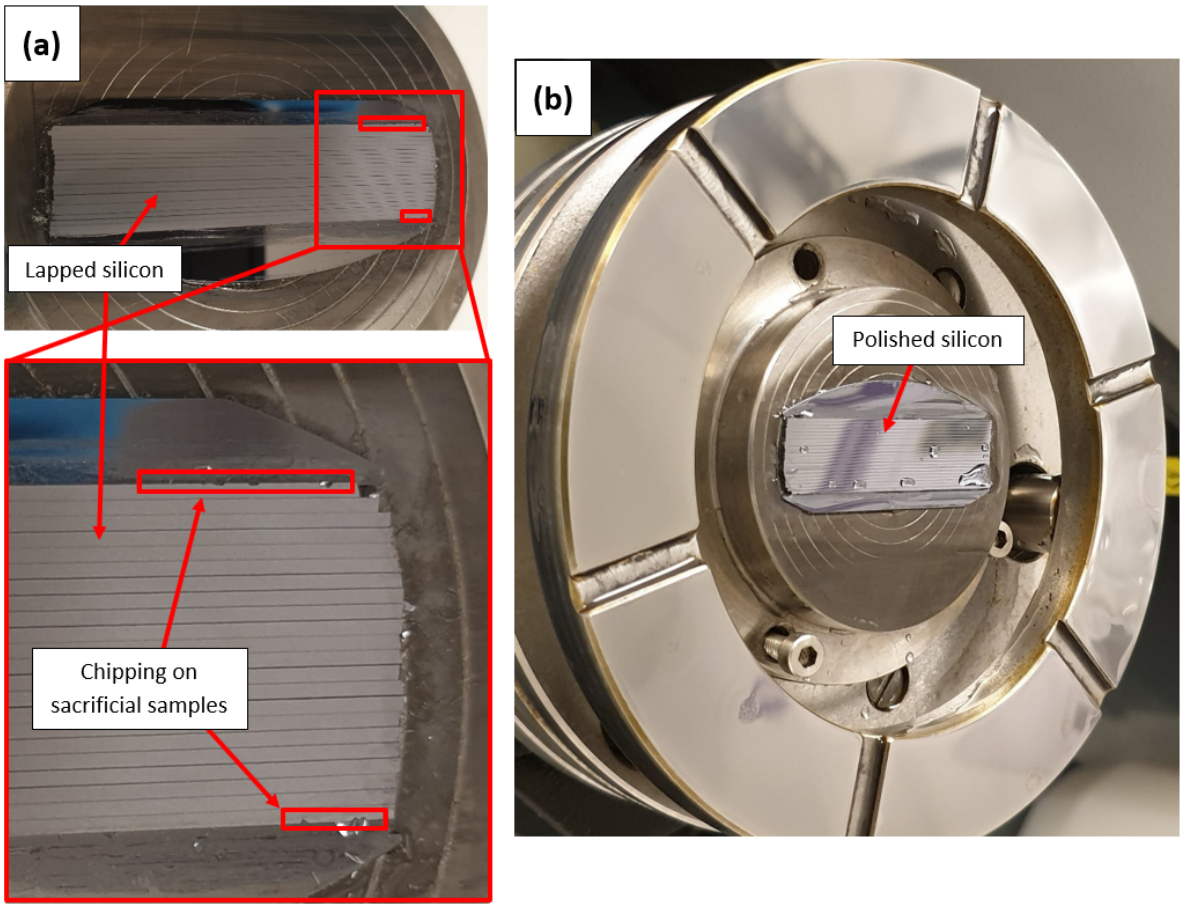

Figure 3.28: (a) Result of the top edges lapped for 2 hours at 50 RPM, resulting in approximately 1100 µm of material removal. Inset of (a) shows chipping on vertices of sacrificial samples. (b) Samples after Syton polishing for 1 hour.

The lapped samples were then attached to the jig. The samples were then initially polished for 40 minutes at 30 RPM since previous experimentation implied this would be sufficient to achieve the required surface finish. The result of this polish can be seen in figure 3.28(b). One can see the metal ring of the PP5 jig has also been polished as expected.

The sample holder was detached and the samples were then imaged, while still on the sample holder, under the Veeco surface profiler. It was found the surface roughness was 12–19 nm R_a, slightly above the desired single figure Å result, as shown in figure 3.29(a).

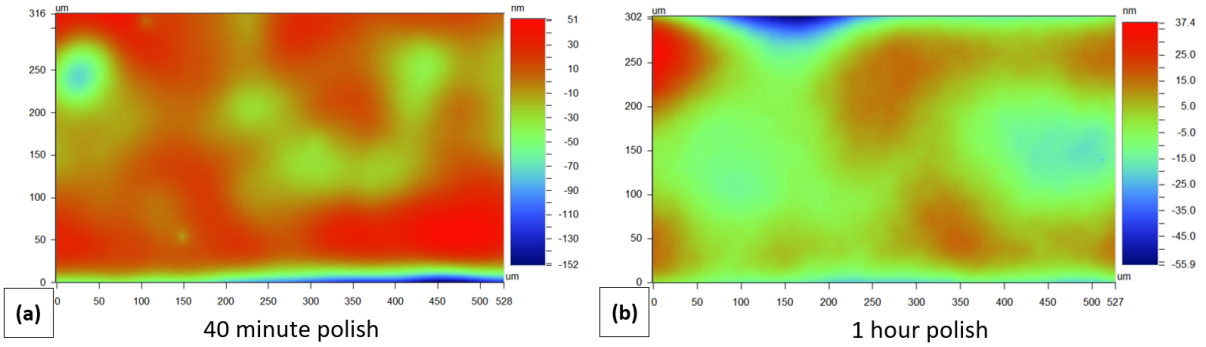

Figure 3.29: (a) Surface roughness measurement of top edge of sample 3068 after 40 minutes of Syton polishing, producing a surface roughness of 16 nm, R_a. (b) Surface roughness measurement of top edge of sample 3068 after 1 hour of Syton polishing, producing a surface roughness of 8 nm, R_a. Colours correspond to hill or valley scales.

The samples were polished for a further 20 minutes at 30 RPM, totalling 1 hour polish time, yielding a roughness of 4–8 nm R_a, as shown in figure 3.29(b), which was deemed sufficient. One can see the extended polishing time helped to bring surface roughness values down.

The entire procedure above was repeated to polish the bottom edges, but this time the samples were initially lapped for only 10 minutes at 50 RPM in order to measure lapping removal rate. Measurement with callipers showed that around 70um of the sample's edge had been removed. Combined with the previous removal rate, this gives an average lapping rate of around 7–9 µm min^{-1}. With 80 µm still to remove, the samples were lapped for a further 10 minutes at 50 RPM, removing approximately 150 µm. The samples were then polished for 1 hour at 30 RPM producing a surface roughness of 2–5 nm R_a.

It should be noted, the sample holder was too large to fit into the SEM chamber. As such, SEM characterisation was not able to be carried out in between the lapping and polishing phases or in between top and bottom edge treatments. Instead, surface roughness measurements were used until the treatment process was completed and the samples could be individually unmounted.

The improved surface quality of the sample edges after polishing is clear from SEM imaging, shown in figure 3.30 and the final R_a numbers as shown in figure 3.31.

SEM imaging showed some chipping along the vertices of the samples. This is possibly due to insufficient slurry flow. Adequate slurry flow is required to ensure removal of lapped silicon from the lapping plate before coming into contact with the samples and damaging the surfaces. Another possible reason for the chipping is due to the pressure placed on the face of the samples by the PP5 jig, where the vertices act as stress concentrators causing chips to occur which relieve stress. This chipping may decrease the failure stress of the samples.

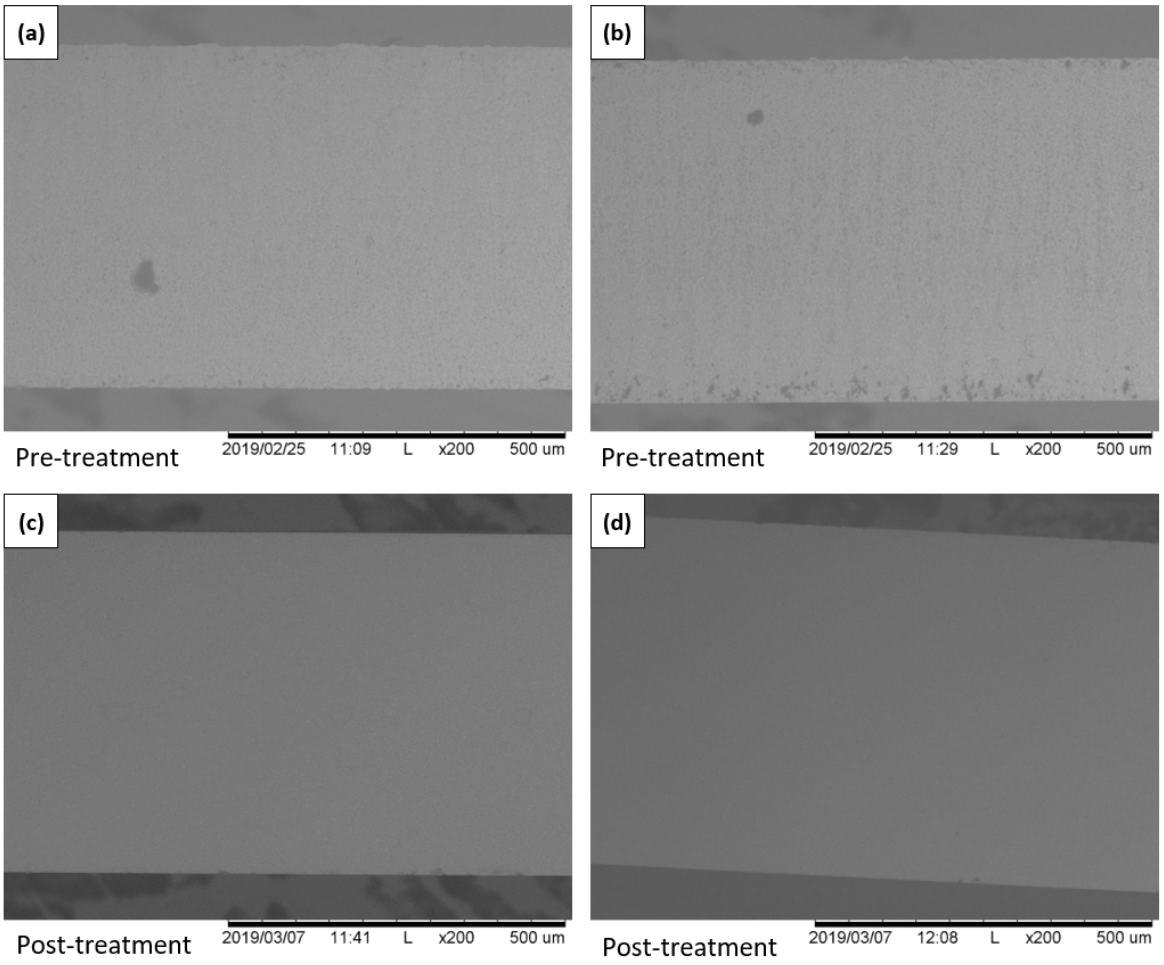

Figure 3.30: SEM imaging of samples pre and post-treatment. (a) New laser technique cut top edge of sample 3064 pre-treatment. (b) New laser technique cut bottom edge of sample 3077 pre-treatment. (c) Top edge of sample 3064 post-treatment. (d) Bottom edge of sample 3077 post-treatment. All images are magnified 200×.

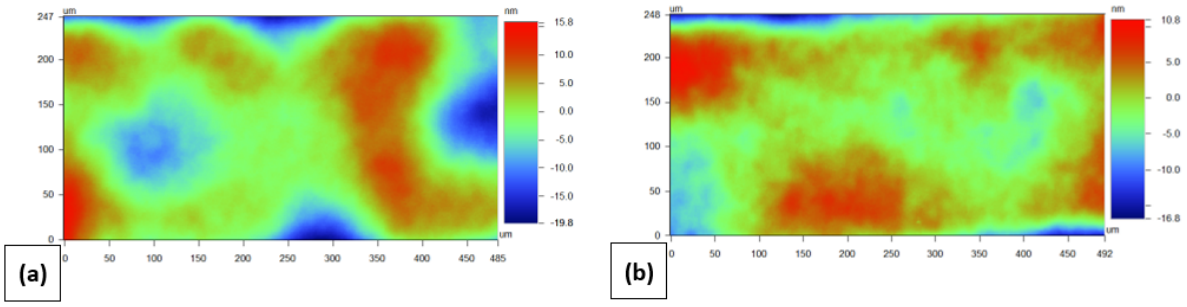

Figure 3.31: (a) Top edge of sample 3059 after polishing, producing a surface roughness of 4.4 nm, R_a. (b) Bottom edge of sample 3064 after polishing, producing a surface roughness of 2.7 nm, R_a.

After surface roughness and SEM characterisation, the samples were cleaned. This consisted of rinsing in DI water, followed by an acetone flush, followed by a further DI water rinse

and finally a methanol flush. No wiping was conducted on the samples to avoid contacting or damaging any treated edges or polished faces.

Finally, the samples were then profiled on the LIGO fibre profiler in both dimensions.

3.6.2.2 Chemomechanical polishing treatment #2

Due to the inadvertent over-lapping in section 3.6.2.1, it was decided that CMP treatment #1 be repeated, with the aim of lapping the intended 150 µm from each edge. The samples treated were similar to those used previously. To improve lapping accuracy, this time each sample was measured on the thin fibre profiler before treatment. Mounting, alignment and conditioning was carried out in line with section 3.6.2.1.

Using treatment #1 as a guide, the top edge of the samples were lapped for 10 minutes at 50 RPM. The samples were then removed to measure the material removal rate.

This time a digital micrometer was used, with a quoted accuracy of 2 µm [291]. Similar to section 3.4.2.5, a random error experiment was conducted on the waxed samples using the micrometer. It was found the random error was 50 µm. Due to the fact the micrometer's measuring area covered approximately half of the 20 samples at one time, this larger uncertainty arises from three errors. Firstly, the samples not being the exact same starting width due to cutting tolerance. Secondly, the samples not being seated properly with varying levels of wax underneath each one. Thirdly, the fact the samples were measured at the very ends (to avoid damaging the exposed treated areas) leading to angular misalignment.

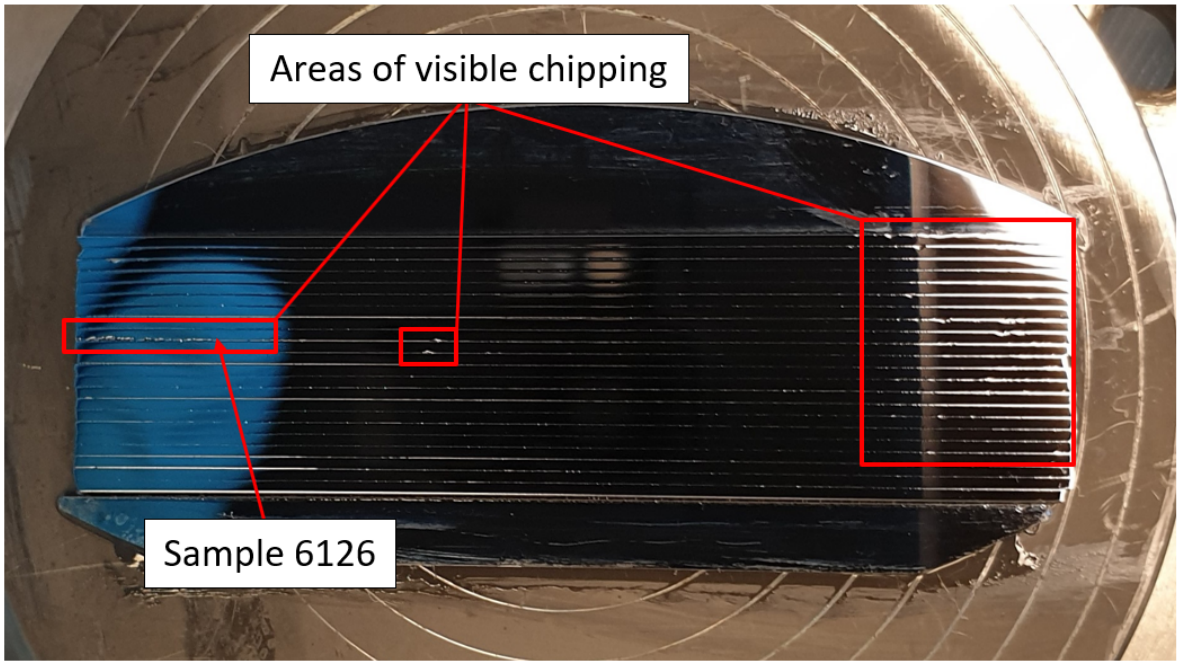

Figure 3.32: Image showing areas of visible chipping of treated samples with particular damage to sample 6126.

142 µm of material was measured to have been removed. With a material removal rate of 14 µm min^{-1} and 8 µm to remove, the samples were replaced and lapped for a further 30 seconds. After conditioning for polishing, the samples were polished for 1 hour at at 30 RPM producing a surface roughness of 3–8 nm R_a.

The bottom edge of the samples were lapped for 10 minutes at 50 RPM. 137 µm of material removal was measured. As such, the samples were replaced and lapped for a further 1 minute. The samples were polished for 1 hour at at 30 RPM producing a surface roughness of 3–9 nm R_a.

This time chipping was visible by eye, not just on the sacrificial samples but also on the treated samples, as shown in figure 3.32. Sample 6126 suffered from particularly bad damage and the reason for this is not known. Misalignment by greater than the 150 µm lapped amount is highly unlikely and would be visible by eye before treatment. Any initial damage due to, for example a vertex sitting higher than the rest, would be lapped away during the first stage of the treatment. Slurry rates were increased slightly for this experiment; but this does not appear to have corrected this issue.

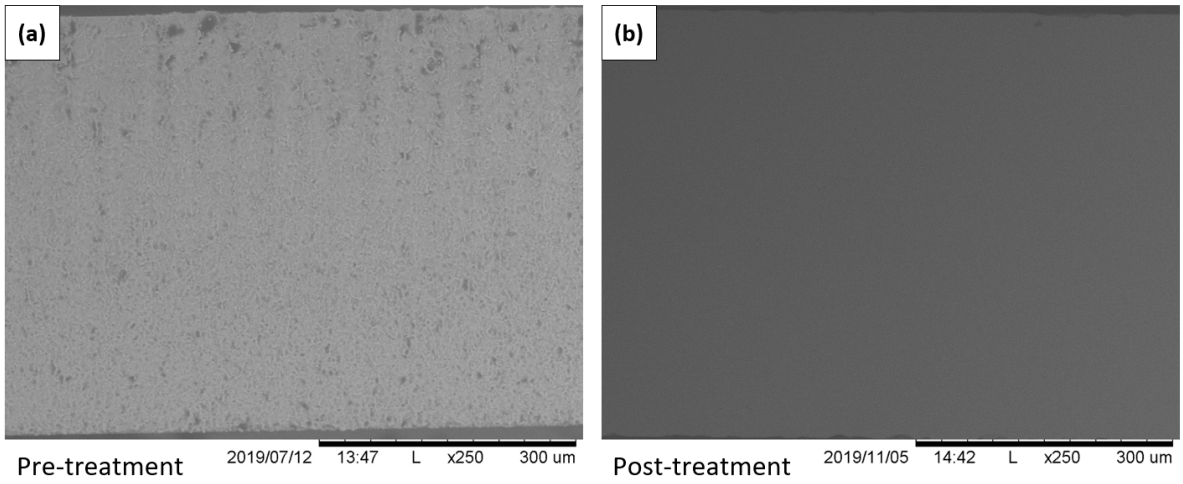

Figure 3.33: (a) New laser technique cut top edge of sample 6122 pre-treatment. (b) Top edge of sample 6122 post-treatment. Both images are magnified 250×.

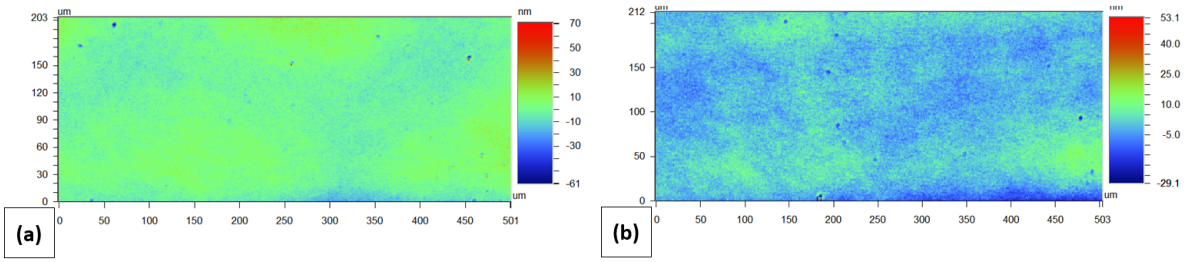

Figure 3.34: (a) Top edge of sample 6119 after polishing, producing a surface roughness of 3.7 nm, R_a. (b) Bottom edge of sample 6133 after polishing, producing a surface roughness of 3.1 nm, R_a.

As in section 3.6.2.1, one can see the improved surface quality of the above process, both visually and under SEM in figures 3.32 and 3.33. Typical surface roughness of polished top and bottom edges are shown in figure 3.34.

The samples were measured on the thin fibre profiler post-treatment in order to calculate the amount of material removed.

3.6.3 Strength results and error analysis

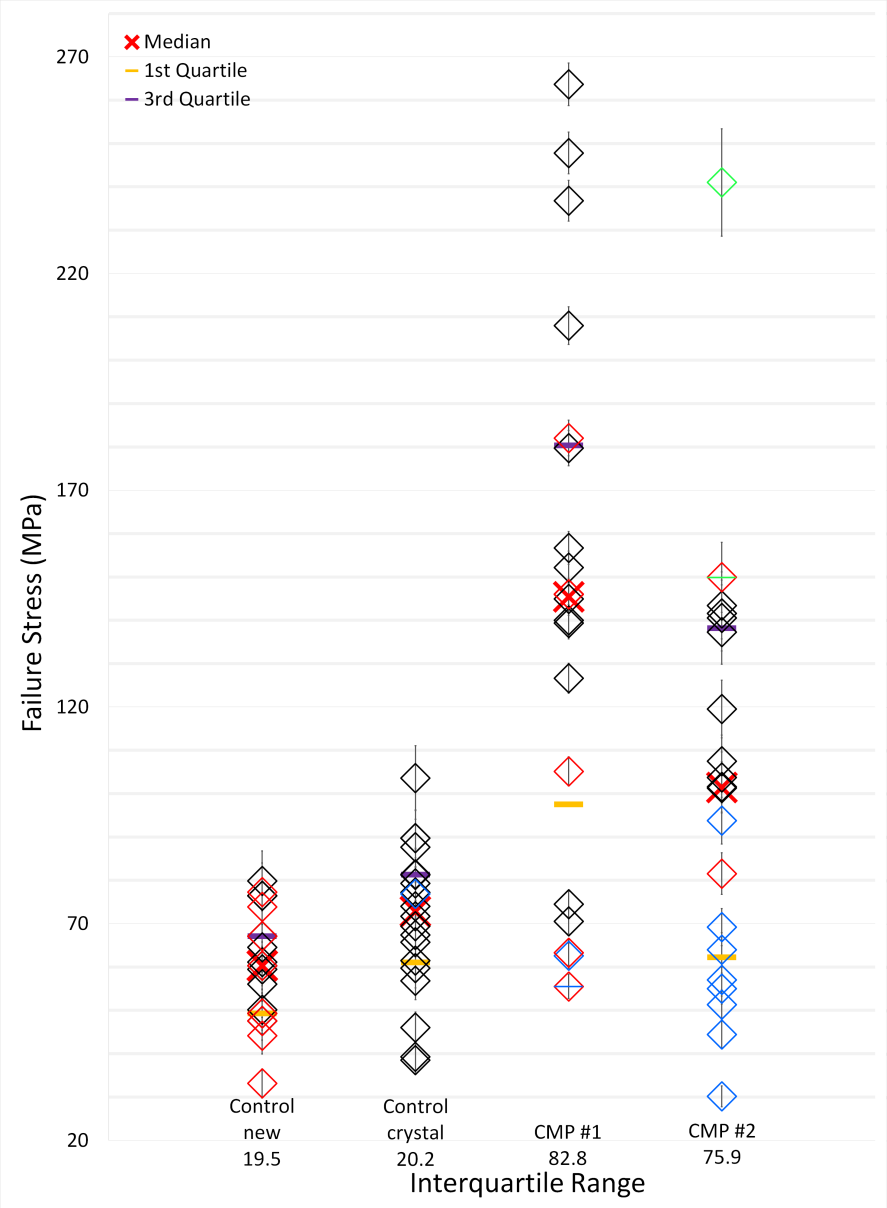

Figure 3.35: Results of tensile strength testing on both CMP treatments. ◇ indicates sample failed at fuse end interface. ◇ indicates a sample was noticeably misaligned in the fuse end assembly. ◇ indicates the sample pulled free from a fuse end due to epoxy failure so the sample was re-epoxied into a new fuse end and re-tested. Where more than one anomaly exists on a sample, the second anomaly is marked with a corresponding horizontal coloured bar.

Both CMP treatments were strength tested and results are shown in figure 3.35.

For CMP treatment #2, sample 6126 is marked as an epoxy failure/re-epoxied sample as the sample was accidentally broken at the top fuse end interface when inserting the fuse end assembly into the strength tester. The sample was recovered and epoxied onto a new fuse end. 8 of the CMP treatment #2 samples were found to be misaligned due to an offset from the centreline of the fuse ends, likely due to the wrong thickness of alignment block being attached to the alignment jig.

Both chemomechanical polishing treatments improved the median strength over the *control new* and *control crystal* samples.

3.6.3.1 Chemomechanical polishing treatment #1 surface removal amount and errors

All 20 samples for CMP treatment #1, were taken from one batch of samples. A subset of 5 samples was measured by calliper for width and thickness. The samples were however individually profiled on the LIGO fibre profiler before strength testing. The appropriate LIGO profiler errors were applied for the measured cross-sectional area of each sample, in line with section 3.4.2.5.

In order to confirm the estimated removal rate of both sides, one can take upper and lower bound measurements, with errors applied to check the material removal rate is consistent with the estimate. The removal amount estimated using the LIGO fibre profiler measurements was 1200–1410 µm, consistent with the experimental procedure.

The errors bars in figure 3.35 arise from the fibre profiling carried out after the experiment. The standard LIGO fibre profiler errors, detailed in section 3.4.2.5, were applied for width and thickness.

3.6.3.2 Chemomechanical polishing treatment #2 surface removal amount and errors

Due to thin fibre profiler measurements before and after treatment, one can measure the material removal amount, as shown in table 3.6. Material removal was under the intended amount of 150 µm.

The errors applied to figure 3.35, are applied from the fibre profiling carried out after the experiment. The standard thin fibre profiler error, detailed in section 3.4.2.5, was applied for width only. The thickness of each sample was not individually measured so the wafer manufacturing tolerance of ± 25 µm was applied for this dimension.

Sample	Width pre-treatment (µm)	Width post-treatment (µm)	Material removed (µm)	Material removed, one side (µm)
5099	2202.0	2000.3	201.7	100.8
6119	2201.1	2014.2	186.9	93.4
6120	2204.8	2005.7	199.1	99.5
6121	2203.2	1981.1	222.1	111.1
6122	2205.6	1985.7	219.9	110.0
6123	2201.6	1969.5	232.1	116.1
6124	2204.8	2017.1	187.7	93.8
6125	2206.8	1999.9	206.9	103.5
6126	2205.0	2011.3	193.7	96.8
6127	2202.9	2032.4	170.5	85.3
6128	2201.9	1993.4	208.4	104.2
6129	2205.5	1913.9	291.6	145.8
6130	2208.5	1936.4	272.1	136.1
6131	2202.1	1952.8	249.4	124.7
6132	2203.0	1933.9	269.1	134.6
6133	2199.5	1930.6	268.9	134.4
6134	2201.2	1986.6	214.6	107.3
6135	2195.8	1917.0	278.9	139.4
6136	2195.6	1934.3	261.3	130.6
6137	2195.2	1921.4	273.8	136.9
		Average	230.4	115.2

Table 3.6: Results of width and material removal before and after CMP treatment #2 as measured on the thin fibre profiler. An error of 0.08 % is applicable to measured values.

3.6.4 Discussion

It is evident in this data set that the misalignment of the 8 samples skewed the results of CMP treatment #2. One sees a 36 % difference between both CMP treatment median values. If one removes the 8 misaligned samples from the data, CMP treatment #2 has a median value of 128 MPa, leading to only a 12 % difference. This shows consistency with the 150 µm surface removal amount.

Figure 3.36 shows a comparison of an aligned to misaligned sample in the tensile strength set-up and the difference in resulting strength that this can have.

The median values from both polished sets increased compared to control samples. The median strength from CMP treatment #1 was 2.4× higher than the *control new* set and 2× higher than the *control crystal* set. The median strength from CMP treatment #2 was 1.7× higher than the *control new* set and 1.4× higher than the *control crystal* set. If one excludes the misaligned samples, these comparative median values increase to 2.1× and 1.8× higher respectively.

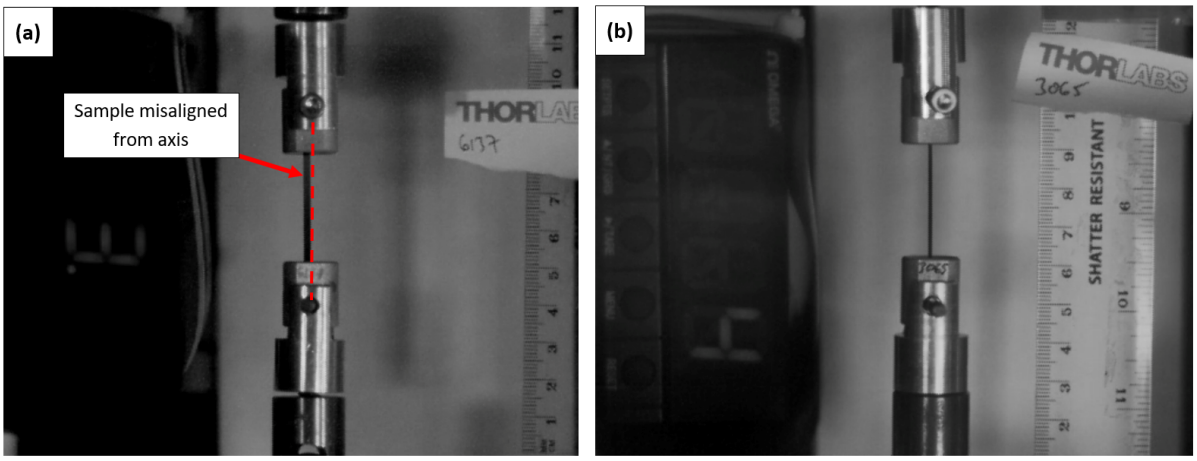

Figure 3.36: (a) Misaligned sample 6137 from CMP treatment #2. (b) A well-aligned sample 3065 taken from CMP treatment #1. Due to the over-lapping from CMP treatment #1, the sample has been noticeably thinned in width. The sample in (b) broke at a tensile strength 149 % higher than in (a).

For aligned samples from both treatments, it is notable that the strongest samples yielded similar strengths with a significant improvement compared to control samples. The strongest sample from CMP treatment #1 yielded a 3.3× increase in strength, with the strongest sample from CMP treatment #2 yielding a 3× increase. Comparing to the *control crystal* set, this reduces to 2.5× and 2.3× respectively.

Both sample sets are consistent with previous work shown in section 3.3, where the control sets were CMP treated [136].

Sample 6126 was shown to have significant chipping, shown in section 3.6.2.2. Strength testing yielded a stress of 101 MPa well within the IQR of CMP treatment #2 sample set, with or without the misaligned sample data removed. This is unexpected considering the level of surface damage exhibited on the sample vertices, however it is possible the surface chipping does not produce significant stress concentrators due to the relatively large radius of curvature of the damaged areas.

By taking the range of stresses for each sample set, one can determine the estimated critical crack size range. For CMP treatment #1, the stress range was 56–264 MPa, resulting in critical crack sizes of 55.5–2.3 µm respectively. For CMP treatment #2, after discounting the misaligned samples, the stress range was 82–241 MPa, resulting in critical crack sizes of 24.5–2.8 µm respectively. If one includes the misaligned samples, the weakest result of 30.2 MPa, yields a critical crack size of 178.7 µm. This result appears unrealistic, as none of the 8 misaligned samples exhibited damage on this scale which would have been clearly visible by eye. This further reinforces that significantly misaligned samples likely break via a non-tensile mechanism due to significant off-axis forces.

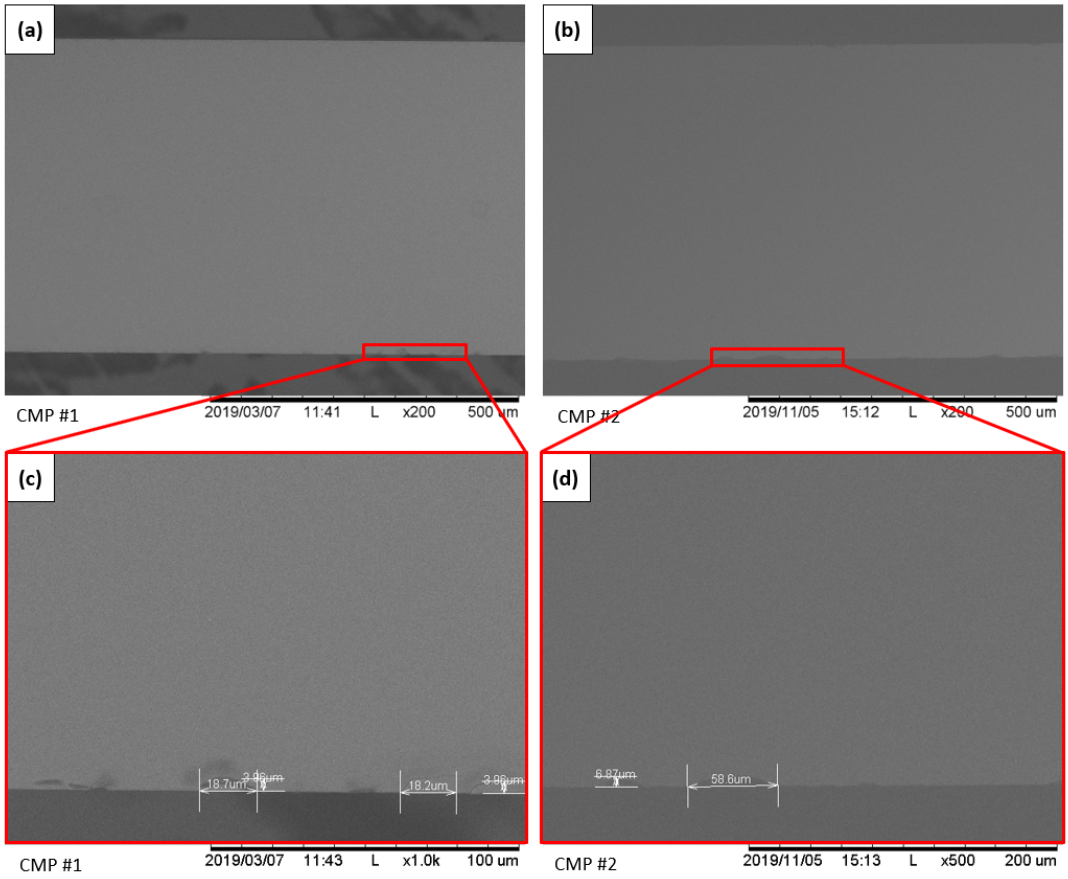

Figure 3.37: (a) Top edge of sample 3064 after CMP treatment #1. (b) Bottom edge of sample 6122 after CMP treatment #2. (a) and (b) are magnified 200×. (c) Bottom edge damage from (a), magnified 1000×. (d) Bottom edge damage from (b), magnified 500×.

These critical crack ranges given by the aligned samples appear to correlate with the order of magnitude of chipping observed on the vertices of treated samples as shown in figure 3.37. It should be noted, these chips were not able to be directly correlated to the failure point of the sample but give a general idea of sample damage depth and critical crack size ranges.

Future experiments should use a more accurate, ideally non-contact, method of measuring samples while attached to the sample holder, such as photography with pixel analysis. Alternatively, one could remove the samples from the sample holder after treating the top edge only in order to more accurately measure the material removal amount on one side. This would both allow for accurate material removal to the desired amount as well as guiding lapping time for treating the opposite edge, providing better uniformity of treatment although care needs to be taken in order to avoid sample contamination.

Further improvement with experimentation lapping and polishing fluid rate may aid in the optimal material removal rate. This would also help flush the lapping plate and polishing pad of any chipped material and remove it before coming back into contact with the samples which is the likely source of the chipping along the sample vertices. Furthermore, using a sample loading gauge to adjust the pressure applied to the samples may aid in reducing vertex damage.

This treatment still relies on substantial physical contact with the surfaces of the silicon ribbons. This is likely to always induce some form of surface cracks or damage compared to non-contact treatments, like CO_2 laser pulling of fused silica fibres, or FZ grown crystal and μ-pulling down techniques specific to silicon. Furthermore, this treatment method does not lend itself well to suspension-scale ribbons, on the order of 60–150 cm. Finally, it may be impossible to apply this treatment method to suspension-scale circular fibres due to their geometry.

3.7 Argon ion etching

Work conducted by R. Birney *et al.* investigated surface treatments on mechanically diced silicon ribbons by way of covering the samples in a variety of coatings [224]. One way of coating the samples was using UHV ion-beam sputter deposition facilities at the University of Strathclyde. This technology is familiar within the gravitational wave community as ion-beam sputter deposition has been used to provide the highly reflective, multi-layer coating on the LIGO test mass mirrors produced by Laboratoire des Matériaux Avancés [292].

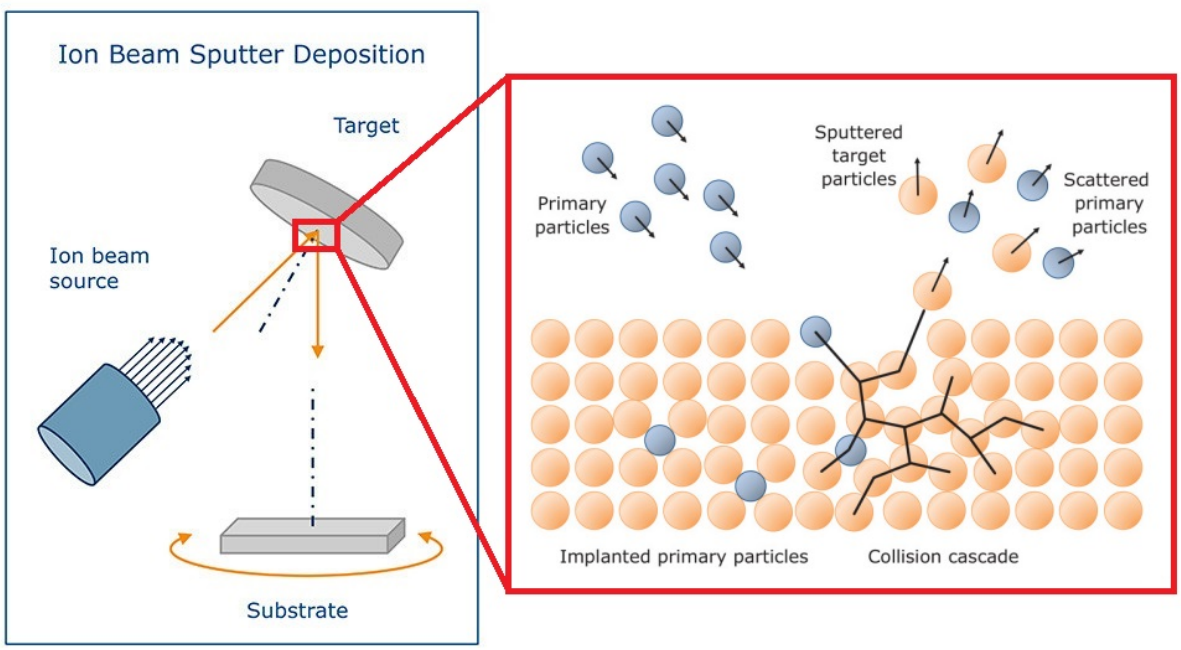

Figure 3.38: Schematic of ion-beam sputtering. In this work, the primary particles are argon ions and the target is silicon. Image modified from [293] [294].

Ions are produced at the ion source and are then accelerated in an electric field within the UHV chamber. As shown in figure 3.38, this accelerates the ions into the target, bombarding the target with enough energy to eject target surface atoms, causing them to sputter material from the target onto the surrounding chamber. The substrate to be coated is then placed within the optimal sputtering path or plume and generally rotated in order to receive an even coating of the

target material's sputtered atoms. Sometimes these positively charged ions are neutralised by electrons from an electron source in order to avoid charging or arcing within the chamber [295].

For the work conducted in this chapter, a novel approach was taken to surface treating silicon with this apparatus. Rather than coat the sample, it was decided that the silicon samples themselves would become the target in the set-up. The damaged edges would be bombarded with argon ions sputtering them off the surface of the silicon ribbons with the intention of using this process to ablate away the surface damage on the sample edges.

3.7.1 Argon ion etching experimental set-ups

Two different ion etching set-ups were available to use at the University of Strathclyde.

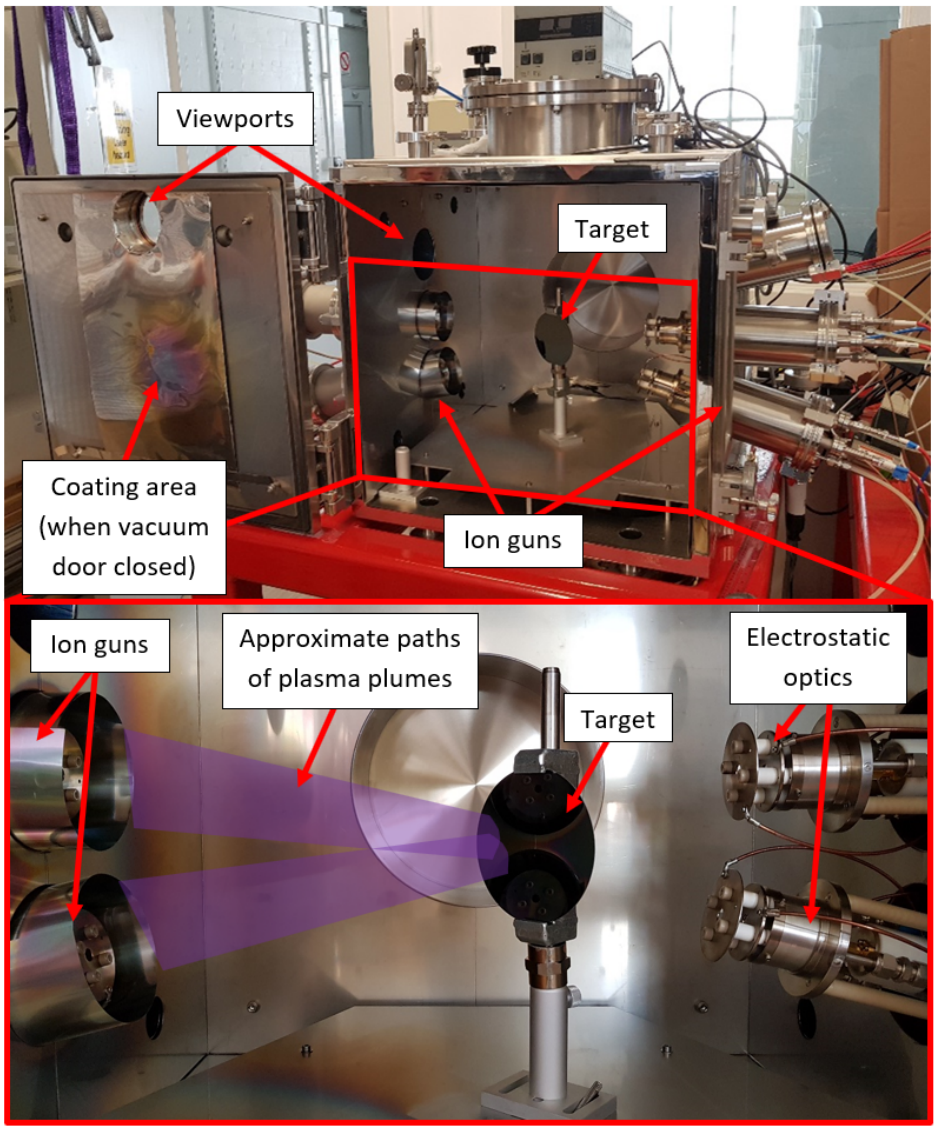

Figure 3.39: The set-up of the argon ion etching experiment within the small UHV chamber. The 2 ion guns on the right hand side of the image are not used during this experiment.

The first set-up is a small UHV chamber, approximately 60 cm^3, shown in figure 3.39. Argon gas is bled into the cavity which is then ionised. These ions are generated in an ultra-high frequency (UHF), 2.45 GHz, resonance cavity via electron cyclotron resonance. The ions are then accelerated through an aperture, via electrostatic optics that produce a divergent beam. The extraction potential is around 10 kV, higher than standard ion sources used for sputtering and etching [296]. While one can vary the etch rate via a number of parameters, the primary etch rate for these experiments was determined by the number of ion sources, the focussing of the beams and their resultant plasma plume. One can reasonably expect the etch rate to be proportional to the number of beams used, assuming all beams have the same divergence and overlap on the same location [296].

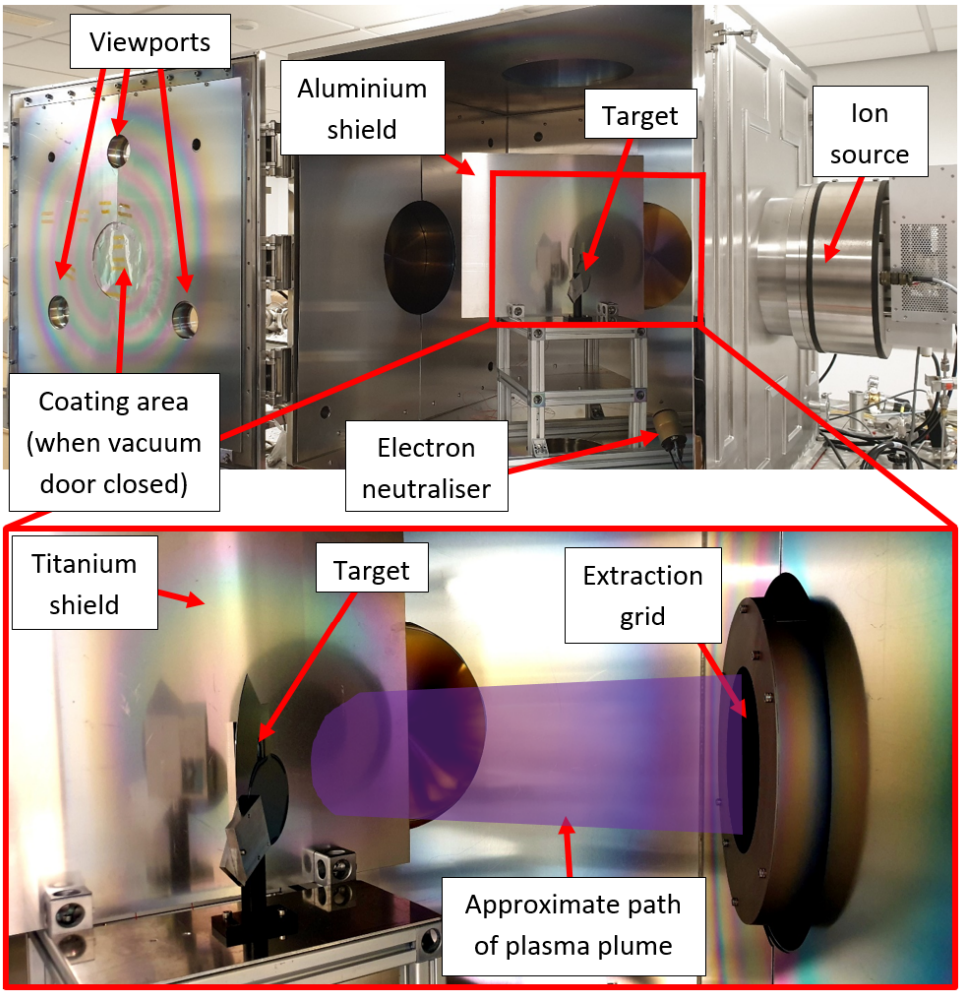

Figure 3.40: The set-up of the argon ion etching experiment within the larger UHV chamber.

The second set-up was a larger UHV chamber, approximately 1.2 m^3, shown in figure 3.40. Ions are generated via radio frequency ion source, producing inductively coupled plasma. These ions are accelerated through an extraction grid held at a lower potential of around 1.5 kV, giving ion energies lower than that typically produced via electron cyclotron resonance but operating at a much higher current, which ultimately produces a much higher etch rate. This also pro-

duces a higher temperature within the system, leading to a greater material plume for coating samples [296]. An aluminium shield has been placed into the centre of the chamber to protect the opposite wall from being etched by the ion source. This was not required for the smaller chamber set-up due to smaller beam size compared to the size of target.

3.7.2 Argon ion etching surface treatment

As with CMP treatment in section 3.6, the intention was to treat the damaged edges of a batch of 20 samples. The small chamber set-up was used initially due to availability. Experimentation on silicon offcuts provided expected etch rates on the edges of the silicon ribbons, however as this was a novel technique, a few technical obstacles had to be overcome before initial experimentation could take place.

A way of clamping silicon samples together without adhesive and exposing only the edges had to be devised. Wax or adhesives can not be used within the system as it is likely to out-gas within a UHV environment. Heating of the samples when exposed to the plasma may also cause further out-gassing from the adhesive or cause it to liquefy. Furthermore, as mentioned in section 3.4.2.6, mechanical clamping of silicon is difficult without damaging the samples by chipping or fracture. This presents yet another challenge as loose clamping of the samples may lead to the polished faces of the silicon ribbons being exposed to the etching process which is not part of the intended treatment process. One must also keep contamination in mind. Anything holding the samples within the ion beam area will also be etched by the ion beam. The material design must be sufficient to withstand both the etching process and any subsequent heat from the plasma without failing or dropping the samples. This material will also sputter itself throughout the chamber system so must be compatible with, or acceptable to, future operators of the coating chamber system.

A pure silicon clamp would be an ideal solution but this is beyond the possibility of current manufacture. Instead, aluminium of 6061-T6 grade was deemed acceptable for use in the chamber, provided it was sufficiently cleaned before and after each etching procedure. The final ion etch jig design can be seen in figure 3.41. The jig was designed to expose the edges of the silicon ribbons to the argon ion etching beam, securing them in place by aluminium clamping blocks. These blocks have had their edges rounded and smoothed to minimise any damage on the silicon ribbon faces. Retaining bolts and washers are then used to secure the samples in place. These bolts are able to run along a channel in order to hold a variable number of samples and clamp them sufficiently. A recess is cut into the back of the jig to aid in aligning the samples vertically. It is worth noting that the full 45 mm exposed edge of the silicon ribbon will not be exposed to the beam as both ends are partly shielded by the retaining bolt channels. This was designed so that no more than 5 mm was shielded at either end, which is acceptable considering 8 mm of either end of the sample will be encased within epoxy and fuse ends, so any untreated edges are not a factor during strength testing. An example of silicon offcuts clamped in an etching

orientation is shown in figure 3.41(b).

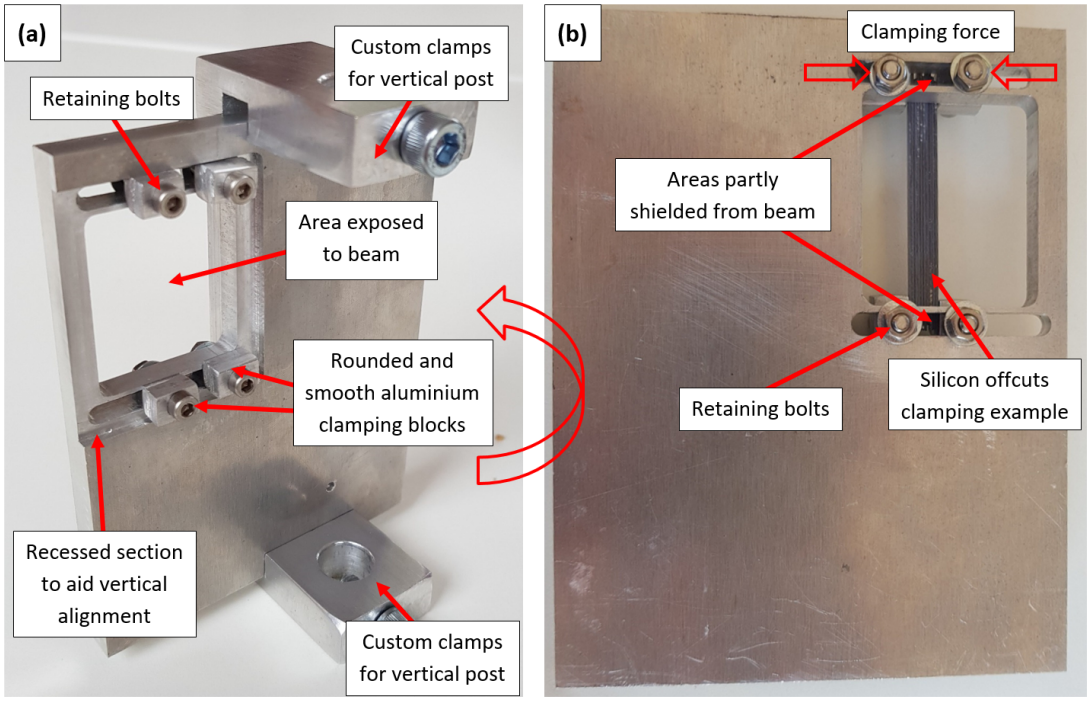

Figure 3.41: (a) Rear of the jig showing the clamping arrangement and custom clamps designed to allow the jig to be secured to a vertical post. (b) Front of the jig, which faces into the etching beam, showing an example of clamped silicon offcuts.

3.7.2.1 Argon ion etching surface treatment #1

Before a batch of 20 samples were treated in the small chamber, it was necessary to etch some test offcuts in order to estimate the etching rate.

Estimation of etch rate - small chamber

20 silicon offcuts of similar dimensions to the treated samples underwent a preliminary etch, although some were shorter than the standard 45 mm length due to being cleaved out of a wafer. 8 of these offcuts were taken from an old laser cut technique wafer. 12 of these offcuts were taken from a new laser cut technique wafer.

Due to the higher voltage and lower current of the small etching chamber etch rates were expected to be on the order of single-figure microns [296]. Samples were weighed collectively before and after the etching process to determine an approximate etching rate. In case the different cutting techniques and edge finish are impacted differently by the etching process, the 8 old laser cut technique samples were weighed collectively but separately from the 12 new laser cut technique samples.

The samples were weighed on a Gemini GR-202 microbalance [297]. This microbalance has

a repeatability (standard deviation) of 0.1 mg, a readability of 0.1 mg and a linearity of 0.2 mg. Thus, the combined error on any *single* weight measurement made is 0.24 mg.

Before treatment, the clamping mechanism, vertical post and jig were dissembled and cleaned in an admixture of IPA and DI water in an ultrasonic bath. The 20 samples were then loaded into the etching jig and sandwiched together, as shown in figure 3.42(a).

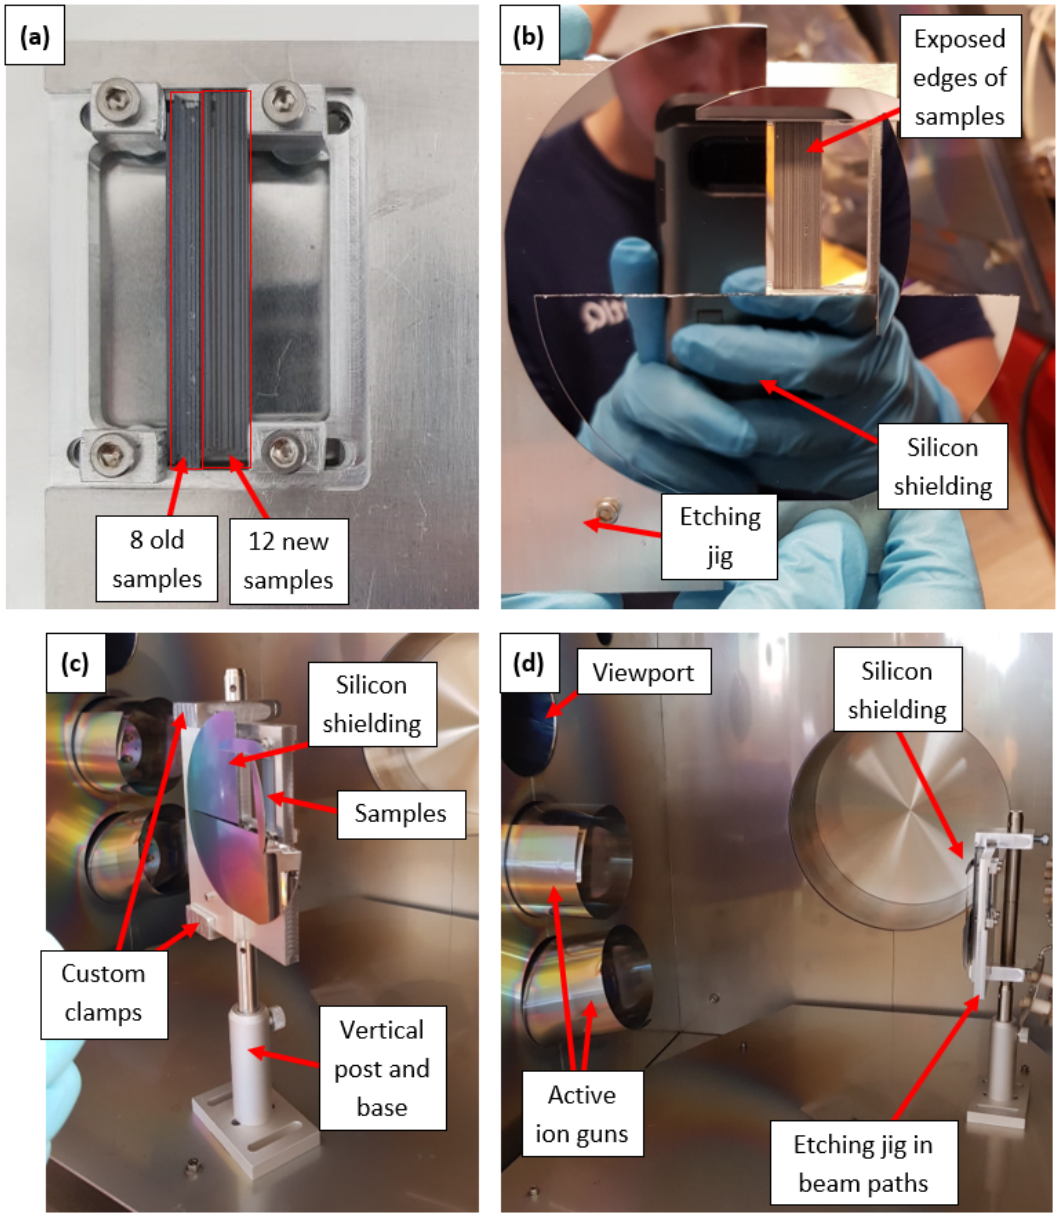

Figure 3.42: (a) Back of the etching jig with the samples affixed. (b) Silicon shielding used to minimise aluminium sputtering from the jig. (c) Vertical, adjustable post and custom clamps used to hold the jig assembly. (d) Final position of the assembly ready to be etched.

In order to minimise excessive aluminium sputtering from the jig around the chamber, it was decided that the exposed areas of the etching jig be shielded by silicon, as shown in figure 3.42(b). Offcuts of silicon were cleaved from spare wafers and affixed to the front of the etching

jig using Kapton tape. Kapton tape is UHV compatible, does not out-gas and is well known to survive at temperature extremes [298].

As shown in figure 3.42(c), the jig was then affixed onto a vertical, adjustable post with a base allowing the entire assembly to be sat in the chamber in any position or angular orientation. Finally, the assembly was placed in the estimated combined beam focal point from the 2 ion guns on the left hand side of the chamber, as shown in 3.42(d). The 3-D position of the jig was noted in order to reproduce similar etch rates for future etching treatments.

Parameter	Value	Unit
UHF power	5	W
Beam potential	10	kV
Electrostatic focus potential	5	kV
Gas inlet pressure	1	$cm^3\,min^{-1}$
Chamber pressure	7×10^{-5}	mbar

Table 3.7: Parameters of the argon ion etch for initial etch rate estimation experiment.

The parameters in table 3.7 were applied to provide ignition and stability to both ion guns. After the beams were ignited, the silicon sample edges appeared to fluoresce which was unexpected. It was not possible to properly image this fluorescence due to the relative brightness of the ion beams.

An etch time of approximately 100 hours was suggested by the chamber operators in order to see appreciable etching [296]. Approximately 72 hours into the etch, the ion beams extinguished themselves. It was suspected that instability within the chamber pressures or ion guns may have been the cause. After a downtime of approximately 3.5 hours, the beams were re-ignited for another 14.5 hours. The beams were then deliberately extinguished in order to modify the parameters by increasing the gas inlet pressure and doubling the UHF microwave power to make up for lost etching time. These new parameters are recorded in table 3.8. This led to a second downtime of approximately 7 hours. The beams were re-ignited for the last time and run for a further 42 hours. This lead to a total etching time of just over 129 hours or 7764 minutes, 5216 minutes with the parameters in table 3.7 applied and 2548 minutes with the parameters in table 3.8 applied.

Parameter	Value	Unit
UHF power	10	W
Beam potential	10	kV
Electrostatic focus potential	5	kV
Gas inlet pressure	1.65	$cm^3\,min^{-1}$
Chamber pressure	7×10^{-5}	mbar

Table 3.8: Updated parameters, applied for the final 42 hours of the argon ion etch for initial etch rate estimation experiment.

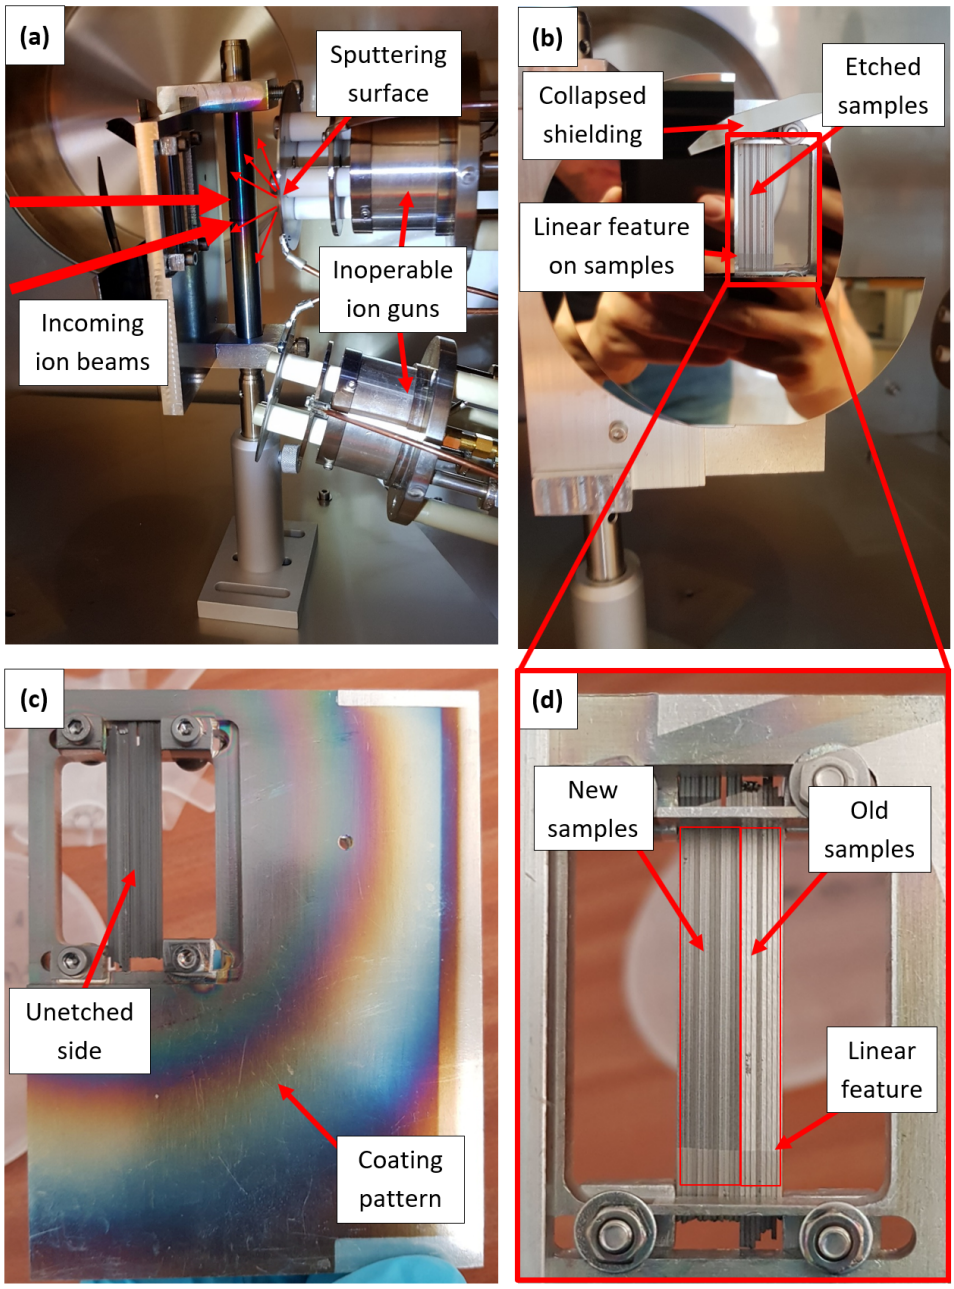

Figure 3.43: (a) Back of the etching jig with the samples affixed. (b) Silicon shielding used to minimise aluminium sputtering from the jig where a small piece has shifted position. (c) Back of the etching jig once removed from the clamping assembly. (d) Close-up of the aperture exposed to the beam, showing the etched edges of the sample with surrounding silicon shielding removed.

The etching jig was notably warm to the touch, after re-pressurising and opening the chamber, implying significant heating had taken placed when exposed to the ion beams. As shown in figure 3.43(a) and 3.43(c), a coating pattern was observed on the vertical post and over the back of the etching jig. This pattern likely arises from the incoming beams passing through the sample aperture and striking the circular structures on the inoperable ion guns shown in figure

3.43(a). These metallic structures then sputter atoms directly onto the vertical post and back face of the clamping jig creating a coating pattern of differing thickness, producing a rainbow spectrum as shown most clearly in 3.43(c). It is also worth noting, there is a high chance of the exposed back edges of the samples being coated in this same layer of material.

Further evidence of substantial warming on the jig was found, as shown in 3.43(b), when a small piece of silicon shield at the top of the etching jig has partially collapsed due to the warping of a loop of Kapton tape affixing it in place. As shown in 3.43(d), the most notable feature on the samples was the presence of a linear feature across the bottom of the samples. This feature will be a result of the position of the silicon shielding. One notes that the beam paths shown in 3.43(a) exposes the very bottom of the samples to the top ion beam, whereas the silicon shielding will block the line-of-sight path from the bottom ion gun. As such, etching has still taken place at the bottom of the samples but at a different rate due to the fact this was only exposed to one ion beam throughout the etching process. This is of no particular concern for strength measurement as it is well within the 8 mm epoxied fuse end assembly.

The above factors make etch rates difficult to estimate from this experiment due to varying length of samples, deposition on the samples from unwanted sputtering, uneven etching on the sample edges and a change in etching parameters. Nonetheless, an idea of etch rate can still be obtained from measuring the weights of the samples post-etch and inferring the etch rate from this, as shown in table 3.9.

Pre-etch	8 old samples	12 new samples
Total weight (mg)	1286.6	1391.8
Sample bag weight (mg)	623.0	621.3
Combined sample weight (mg)	663.6	770.5
Cumulative error (mg)	0.3	0.3
Post-etch		
Total weight (mg)	1283.4	1385.7
Combind sample weight (mg)	660.4	764.4
Weight lost to etch (mg)	3.2	6.1
Weight lost per sample (mg)	0.4	0.5
Volume lost per sample (μm^3)	1.7×10^{14}	2.2×10^{14}
Approximate area of sample (μm^2)	2.4×10^7	2.4×10^7
Approximate etch depth (µm)	7.3	9.2
Average etch depth (µm)	8.3	
Approximate etch rate ($nm\,min^{-1}$)	1.1	
Error (µm)	1.6	

Table 3.9: Pre and post-etch weights of old and new laser cut technique samples with estimated etch rate.

One notes the very large error of 145 % on the etch rate due to the aforementioned factors making etch rate estimation difficult. The approximate etch rate is equivalent to $1.6 \pm 2.3\,\mu m$

per day.

Argon ion etch #1 - small chamber

With the etch rate of the system estimated, a batch of 20 samples was surface treated. These samples were cut with the new laser cutting technique from 525 µm thick wafers, the same as the *control new* samples. In order to fully investigate this novel etching technique, all 20 samples underwent SEM imaging on both laser cut edges.

Samples were weighed on the same microbalance equipment used previously, though this time each sample was weighed individually.

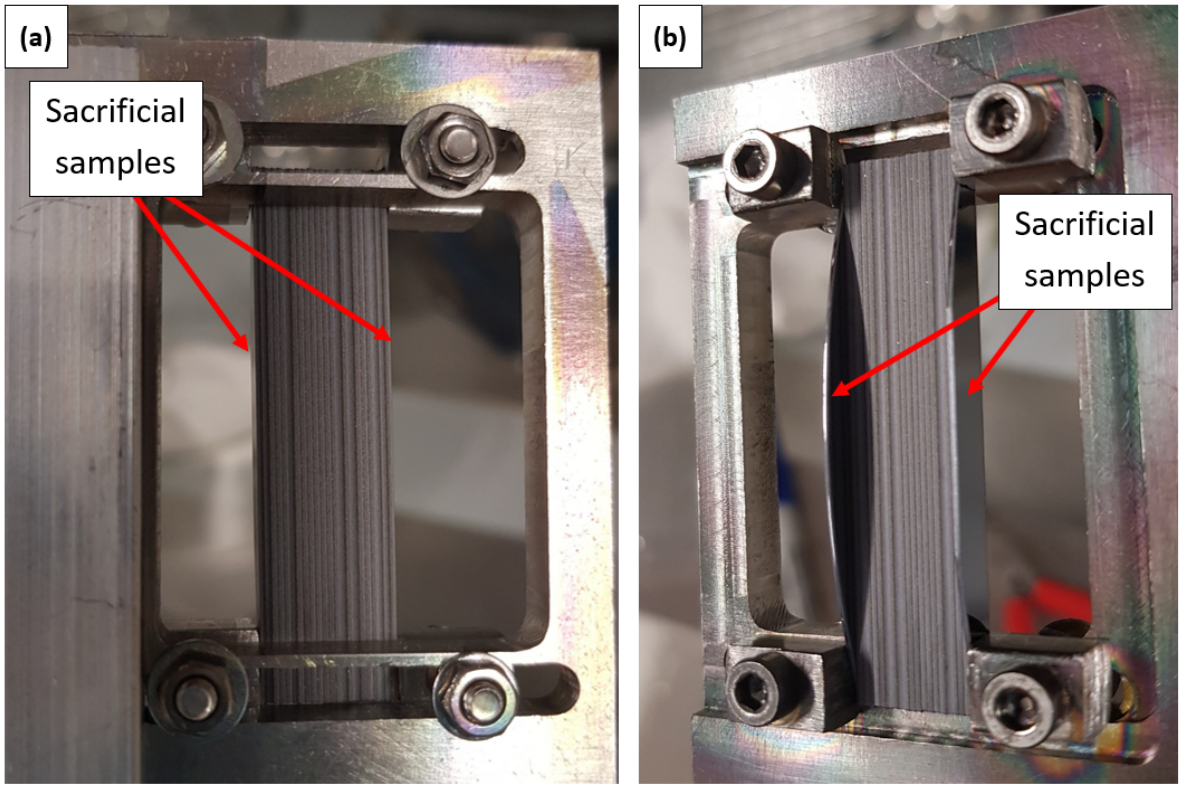

Figure 3.44: (a) Front of the etching jig that will be exposed to the argon ion beam. (b) Back of the etching jig, angled to show sacrificial offcuts.

As shown in figure 3.44, the samples were then placed into the etching jig with sacrificial silicon offcuts placed at either end. The purpose of these sacrificial samples was to both protect the polished faces of the treatment samples from the aluminium clamps and minimise sputtered material coating the faces of the samples.

The only change between etches was the fact the jig itself was clamped more centrally by the custom clamps rather than offset as before. This does not change the relative position of the samples to the argon ion beam. Consistent positioning of samples in the etching chamber is important for reproducible etch rates.

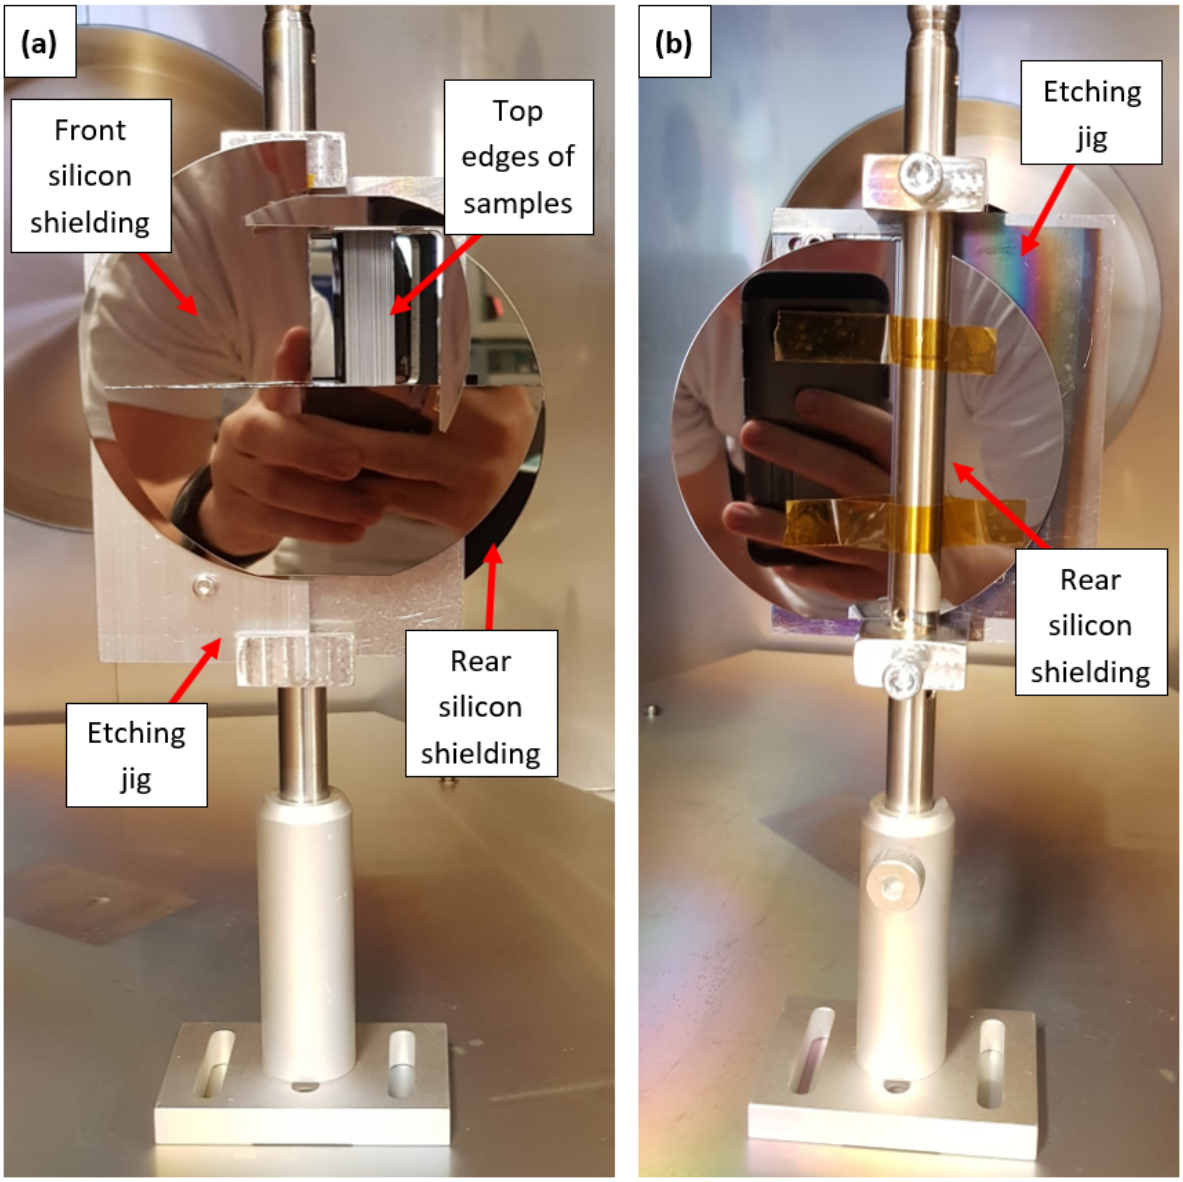

Figure 3.45: (a) Front of the etching jig with silicon shielding. (b) Back of the etching jig with a whole silicon wafer as a rear shield.

In order to negate sputtering of metallic atoms from the area around where the inoperable ion guns were, a whole 4" silicon wafer was attached to the back of the etching jig, meaning only silicon should be sputtered onto the back edges of the sample. The full etching jig set-up, complete with silicon shielding front and back is shown in figure 3.45.

Due to potential damage from further etching of the opposing ion guns, these were removed for the etching treatment.

Parameter	Value	Unit
UHF power	10	W
Beam potential	10	kV
Electrostatic focus potential	5	kV
Gas inlet pressure	1.65	$cm^3\,min^{-1}$
Chamber pressure	7×10^{-5}	mbar

Table 3.10: Parameters of argon ion etch #1 applied to top and bottom edge of the samples.

The intention was to etch the top edge of the samples for just over 3 days, or 80 hours, using the parameters set in table 3.10. This should give an etch depth of approximately 5 µm from the initial etching experiment etch rate estimation.

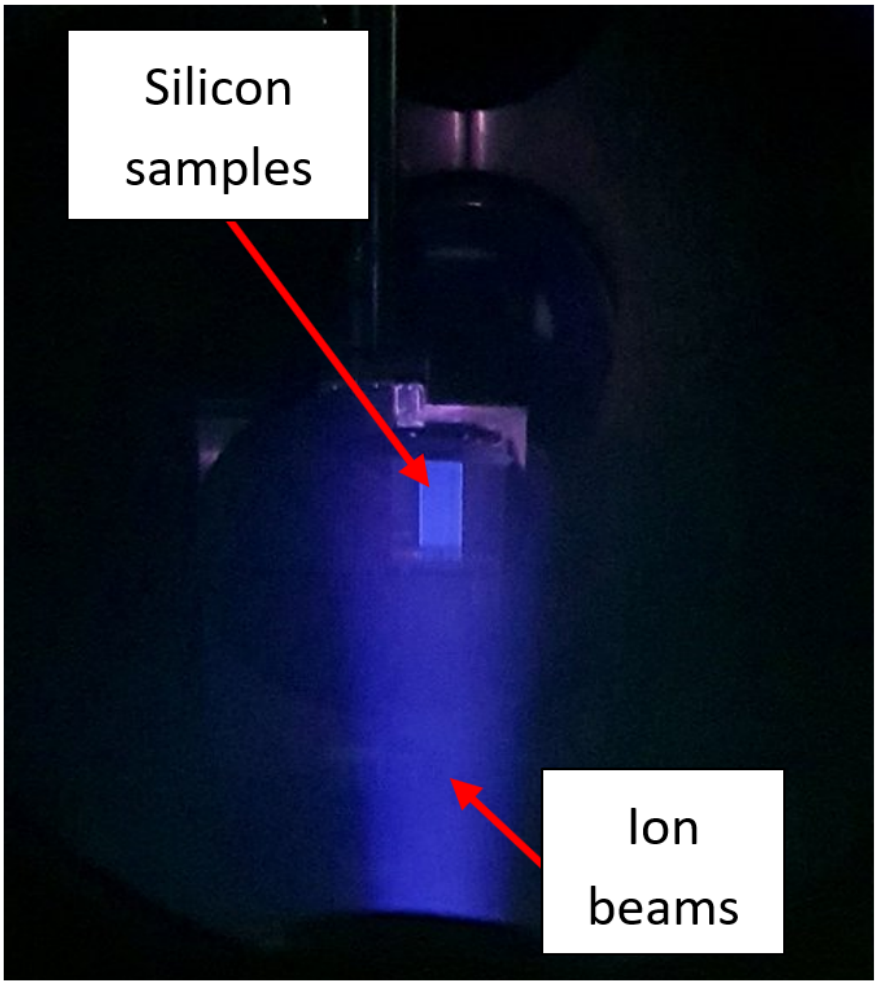

Figure 3.46: Image taken from the viewport of the top edge of the samples undergoing argon ion etching.

This etch depth is far short of the 150 µm etching target for surface treatments. Due to the nature of argon ion bombardment this is an unrealistic etch amount. Firstly, it is difficult to keep stable conditions in the chamber required for plasma production for extended periods of time.

Secondly, the sheer length of experimental time would be on the order of 3 months for treating one side of the samples to achieve this target amount. This chamber was required for other coating experiments and so the overall etching experiment was limited to around one week of operation before the experiment had to be concluded.

An image of the top edge of the samples undergoing argon ion etching is shown in figure 3.46. One noted the recurrence of fluorescence of the sample edges however this effect faded after approximately one day implying this was either due to surface contaminants burning off or some form of reaction with the native oxide. The actual etching time applied to the top edge of the samples was just over 115 hours, or 6915 minutes, due to scheduling and staff availability.

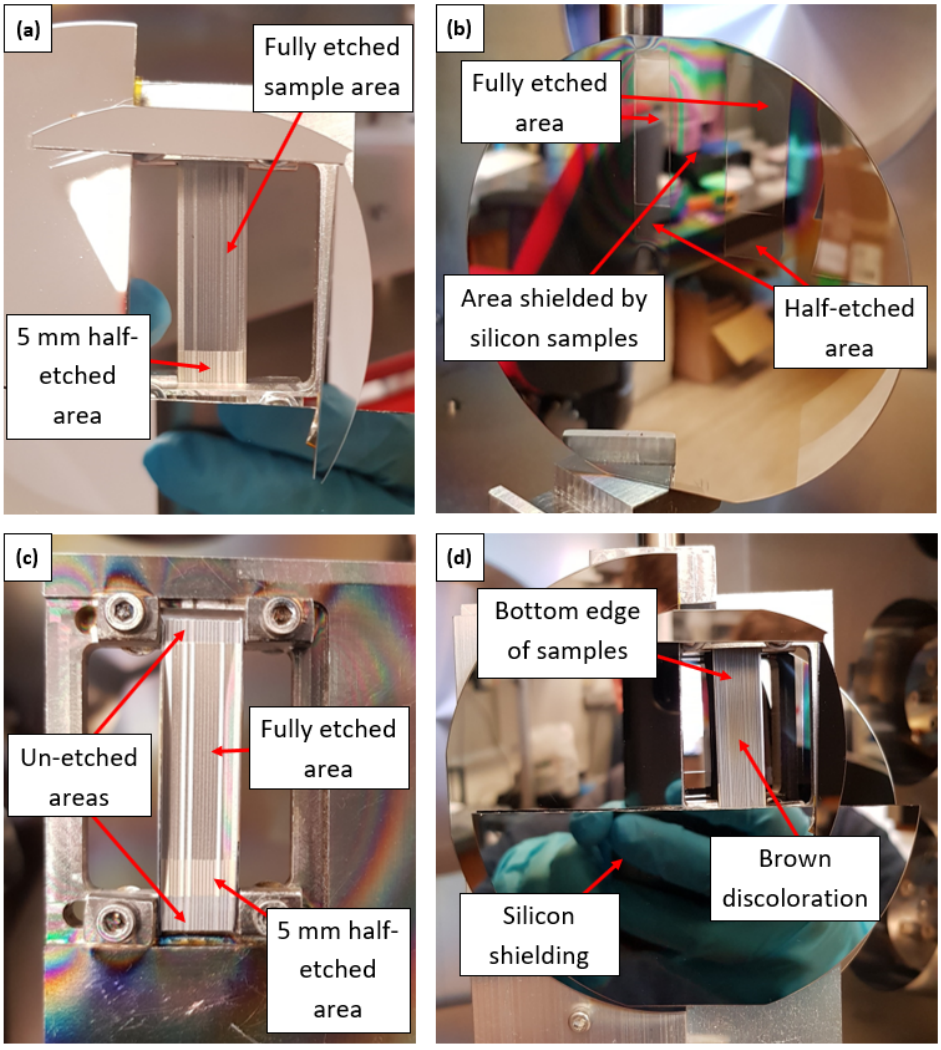

Figure 3.47: (a) Top edges of the samples after treatment, highlighting the area that was exposed to only one ion gun. (b) Rear silicon wafer shield after exposure to the etching process. (c) Samples flipped, highlighting the half-etch and un-etched areas on the top edges of the samples.(d) Bottom edges of the samples ready to be treated, highlighting the area of brown discolouration.

As shown in figure 3.47(a), it was noted due to the position of the silicon shielding and lower ion gun that the samples appeared to be shielded from the lower ion gun at the bottom by

approximately 5 mm, which means a half-etched area may be exposed during strength testing, as this sits above another 5 mm that is not exposed to any etching due to it being fully shielded by the etching jig. This area is termed "half-etch" due to the fact it will still have been exposed to the upper ion gun during the etching process. This half-etched area is also shown clearly on the rear silicon wafer shield, as shown in figure 3.47(b).

The samples were flipped in the etching jig to expose the bottom edges and this time the samples were aligned closer to the top of the etching aperture in order to minimise the half-etched area at the bottom of the samples, as shown in figure 3.47(c). The bottom edges of the samples were then shielded as before and ready to be treated as shown in figure 3.47(d). One can also seen a slight brown discolouration on the bottom edge of the samples, highlighted in figure 3.47(d). It is unknown what the cause of this feature is, though it may arise from localised heating.

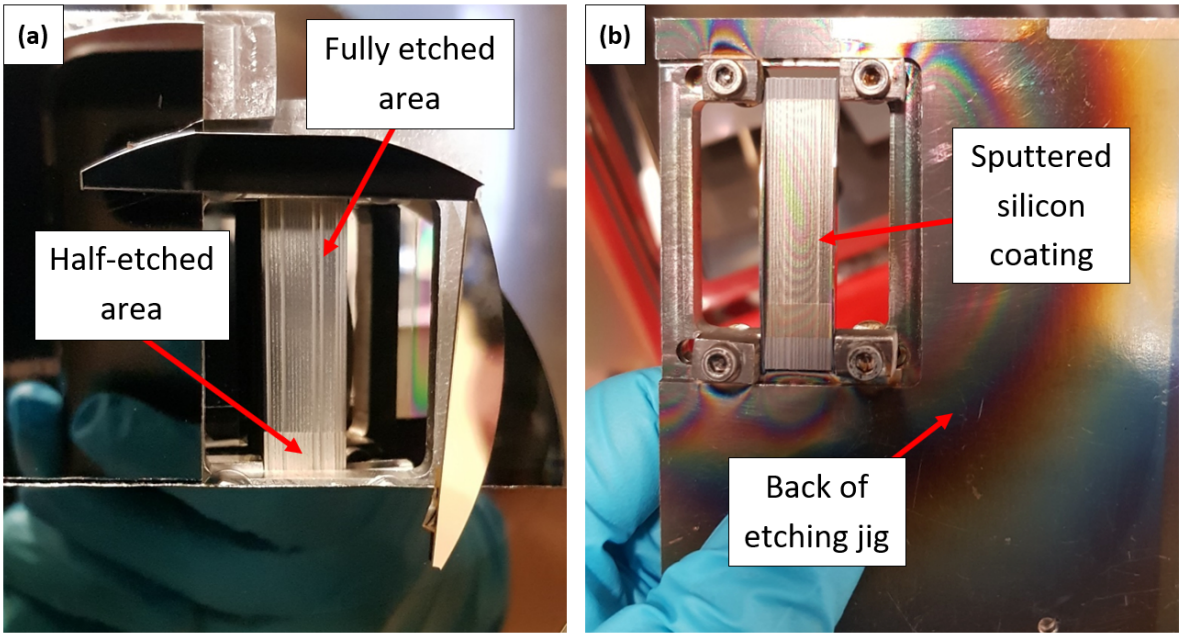

Figure 3.48: (a) Bottom edges of the samples after treatment, highlighting the area that was exposed to only one ion gun. (b) Rear of the etching jig and previously treated top edges of the samples, highlighting the uneven silicon sputtered coating.

The etching jig was placed back into the small UHV chamber in the same position and set-up. The etching time applied to the top edge of the samples was just over 117 hours, or 7034 minutes, to align with normal staff working hours. Once complete, the etching jig and samples were inspected again, noting the half-etch area was once again observed on the bottom edges but this time over a smaller area as expected, as shown in figure 3.48(a). Figure 3.48(b) shows a coating of sputtered silicon onto the top edge of the samples. An interference pattern is present owing to how thin the coating is. This coating is likely sputtered silicon from the rear silicon shield.

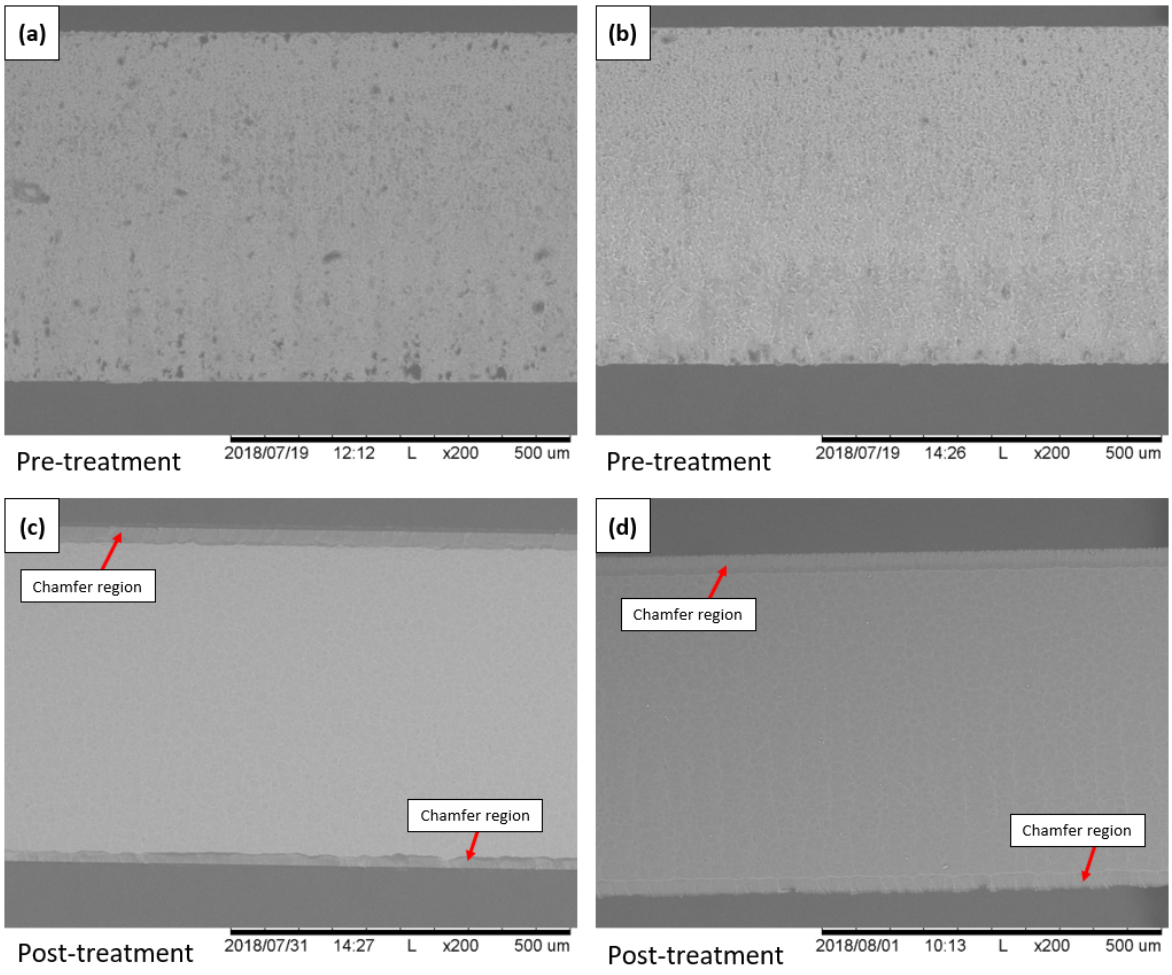

Figure 3.49: (a) SEM imaging of the bottom edge of sample 1006, before argon ion etching. (b) Top edge of sample 1013, before argon ion etching. (c) Bottom edge of sample 1006, after argon ion etching. (d) Top edge of sample 1013, after argon ion etching. All images are magnified 200×.

Although the samples appeared to show a shiny surface on the treated edges, the roughness was too high to measure on the surface profiler. SEM images in figure 3.49 compare typical sample edges before and after the etching process. An improvement on both top and bottom treated edges can be seen in figures 3.49(c) and (d) where these regions have been fully etched by both argon ion guns. Chamfer regions highlighted in figures 3.49(c) and (d) are discussed below.

One notes two particular features unique to argon ion etch treated samples. The first is a particular cell-like structure as highlighted in figures 3.50(a) and (b) on both top and bottom treated edges. It is assumed these structures form from the direct ion bombardment and are on the order of 20 µm or smaller in size across the sample edges. Also highlighted in figures 3.50(c) and (d) are the half-etch and full-etch boundaries on the edges of the sample. A clear etch step is visible in both images as well as a change in the surface quality on either side of the boundary.

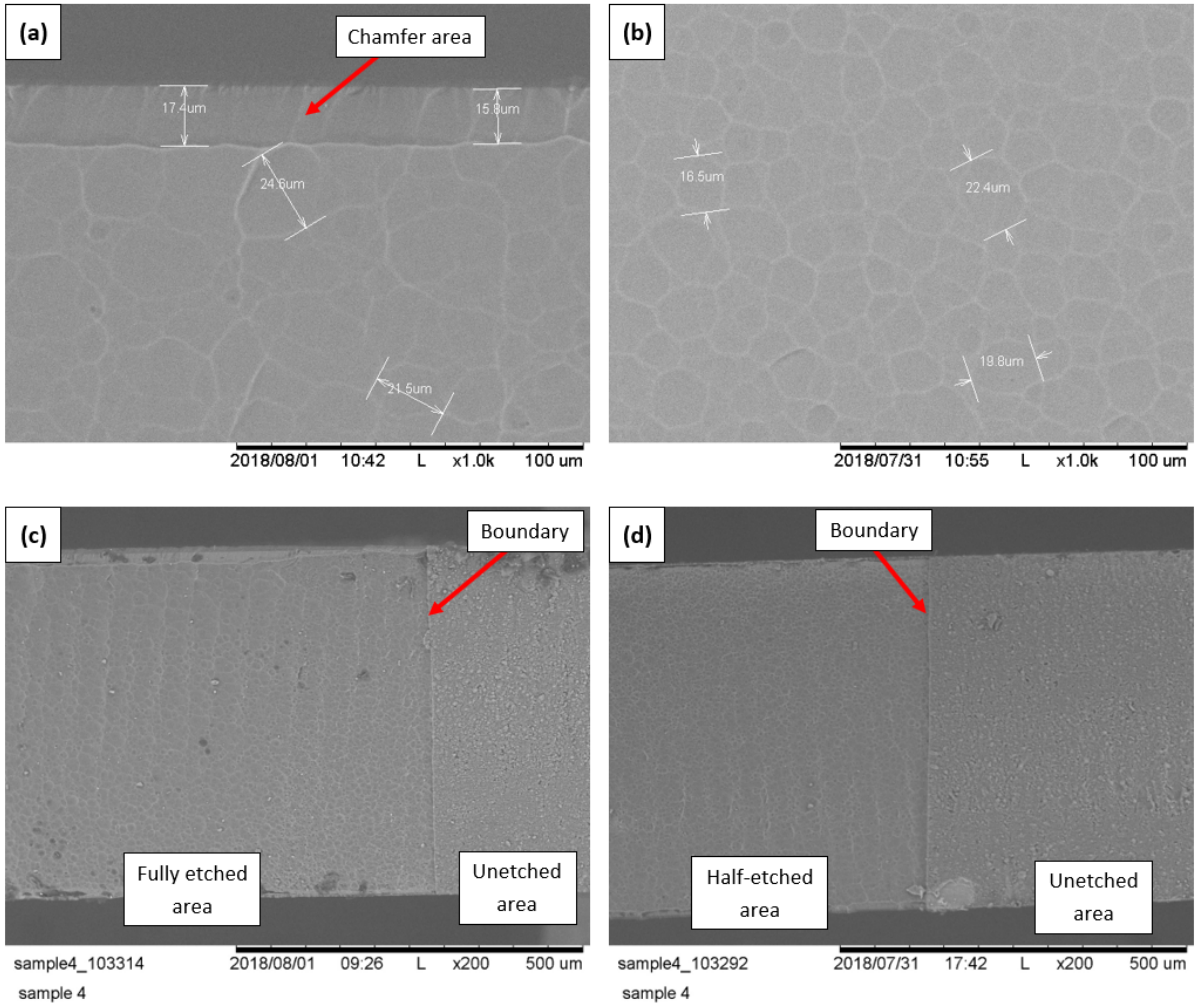

Figure 3.50: (a) SEM imaging of the top vertex of the bottom edge of sample 1013, after argon ion etching. (b) Centre of the bottom edge of sample 1001, after argon ion etching. (a) and (b) are magnified 1000×. (c) Near the top of sample 1009, at the etching boundary. (d) Near the bottom of sample 1009, at the etching boundary. (c) and (d) are magnified 200×.

The second feature appears along the vertices. One can see some form of chamfer forming into the page as shown in figure 3.50(a) and previously in figures 3.49(c) and (d). This appears to be where argon ions have entered in between the edges of the samples and potentially ricocheted off the polished faces between each sample. It is likely this increases the chances of trapping ions within this region as they collide with further incoming ions [296]. This appears to have caused a form of chamfering where these ricocheting ions etch perpendicular into the vertex where the treated edge meets the polished face. A schematic of the suspected perpendicular etching mechanism is shown in figure 3.51.

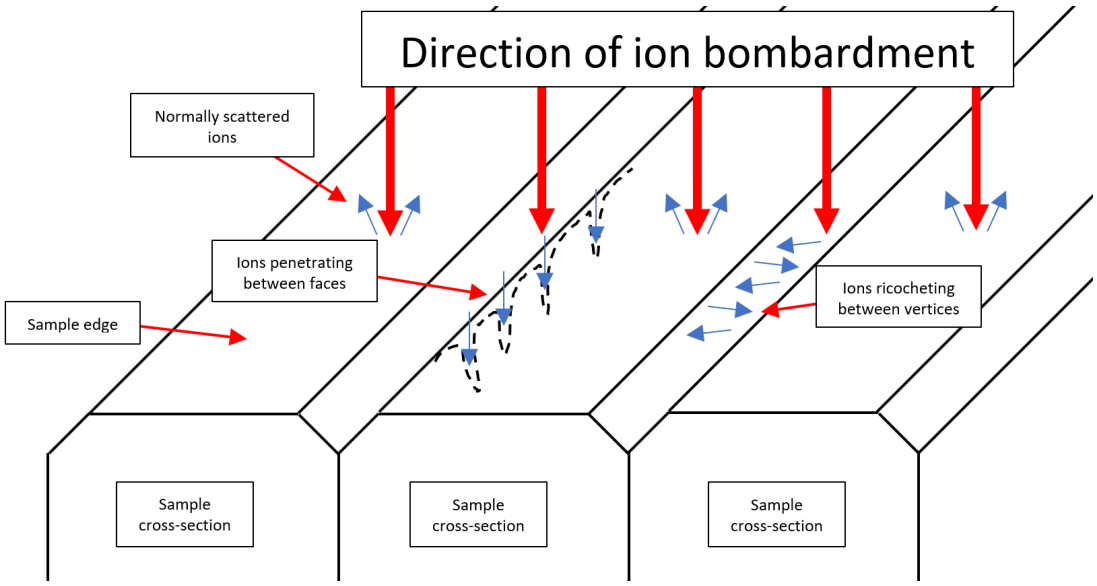

Figure 3.51: Schematic of suspected etching mechanism responsible for etching down between polished faces of the samples.

It is believed these same features are seen in SEM imaging in figure 3.52 from a different angle. Figure 3.52(a) appears to show etched material removal approximately 50 µm, relative to the bottom edge, into the face of the sample along the vertex. One can clearly see the appearance of channels formed in the direction of argon ion bombardment. This appears to suggest that as some ions enter directly into the gap between samples polished faces, they penetrate further down in between the samples, but likely lose energy due to ricocheting back and forth between the polished faces, creating this channel effect. Figure 3.52(b) shows the back polished face, looking along the bottom edge of the sample and focusing on the vertex. Once again, similar etching is observed along the vertex here although the depth appears to be reduced to around 20 µm, relative to the bottom edge of the sample.

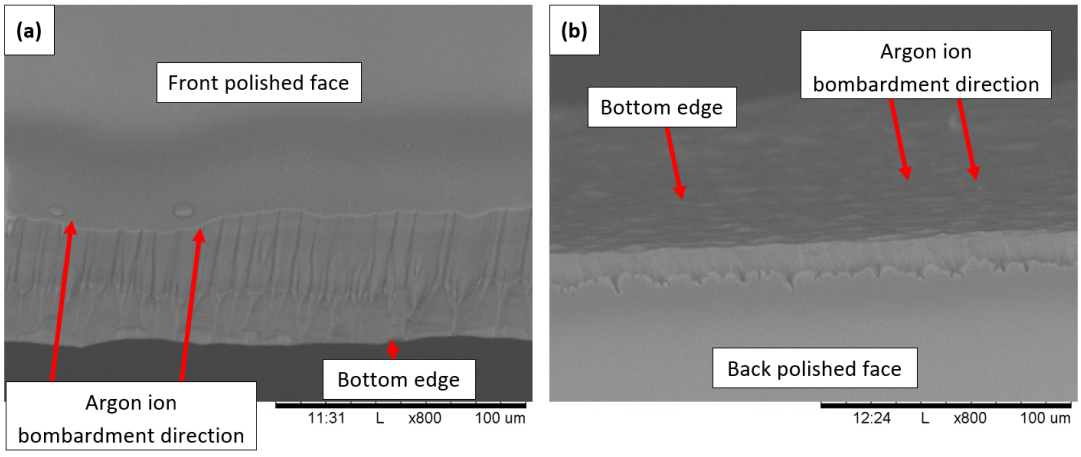

Figure 3.52: (a) SEM imaging of the front polished face of sample, magnifying the fully etched bottom edge. (b) Back polished face of sample, zoomed on the fully etched bottom edge. Both images are magnified 800×.

3.7.2.2 Argon ion etching surface treatment #2

As discussed in section 3.7.1, the large UHV chamber is an entirely separate experimental set-up with lower extraction potential applied to the ions but operating at a much higher current which should significantly increase the etch rate.

The second advantage of this second treatment, other than the significantly higher etch rates, was the removal of a second ion source that lead to areas of half-etching when exposed to only one beam. With all samples sitting in the line-of-sight of one ion source, with a larger beam radius, one should expect more uniform etching across the sample set.

Estimation of etch rate - large chamber

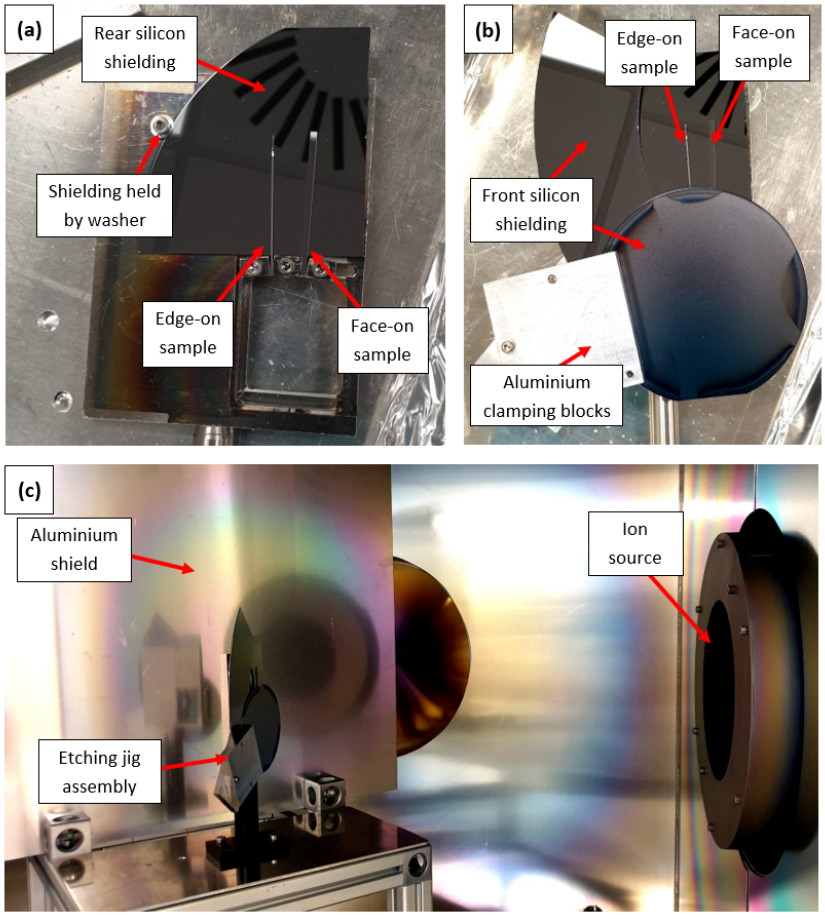

Figure 3.53: (a) Modified orientation of the etching jig used for argon ion etch #1 with two silicon samples. (b) Front silicon shielding and clamping blocks used to hold it. (c) Etching jig assembly set-up in the large UHV chamber.

As with section 3.7.2.1, an initial experiment had to be conducted in order to estimate the etch rate. Due to the higher etch rate, and potentially higher temperatures, Kapton tape could not be used for affixing any silicon shielding, and shielding had to be extensive to minimise any

contamination from sputtering from other materials. An unpolished wafer and a large curved piece of silicon offcut were used as front shielding, as shown in figure 3.53(b). A quarter of a cleaved silicon wafer was also to be placed directly behind the samples to avoid any line-of-sight path of the ion beam through the aperture to minimise sputtering from the aluminium target onto the back of the samples and throughout the chamber.

These requirements led to a modification of the etching jig, as shown in figures 3.53(a) and (b), in order to affix silicon shielding by means of clamping with aluminium blocks and washers that needed to be outside the approximately 16 cm ion beam diameter. A temperature sensor was also attached onto the rear of the aluminium shield.

As shown in figure 3.53, two silicon samples, identical to the *control new* samples, were to be etched, one edge-on and one face-on to the ion beam, in order to probe if etching rates were affected by initial surface quality.

Sample	Dimension (μm)	Error (%)
Face-on	517.1	0.05
Edge-on	2214.8	0.25

Table 3.11: Pre-etch dimensions of silicon samples before initial etch test.

Due to the large uncertainties from weighing samples it was decided that fibre profiling would be a better method of measuring material removal from etching and the corresponding etch rates. The face-on sample was measured on the thin fibre profiler and the edge-on sample was measured on the LIGO fibre profiler due to availability of each profiler at the time. The dimension of each sample exposed to the etching process and the corresponding error is recorded in table 3.11.

Parameter	Value	Unit
Beam current	300	mA
Beam potential	1.5	kV
Acceleration voltage	100	V
Acceleration current	8	mA
Ion source emission voltage	41	V
Ion source emission current	450	mA

Table 3.12: Parameters of the initial etch test for argon ion etch #2.

The samples were etched with the parameters recorded in table 3.12. An image of the samples undergoing the initial etch test is shown in figure 3.54. The samples were exposed to the ion beam for a total of over 16 hours, or 1000 minutes. The maximum temperature recorded on the aluminium plate was 180 °C.

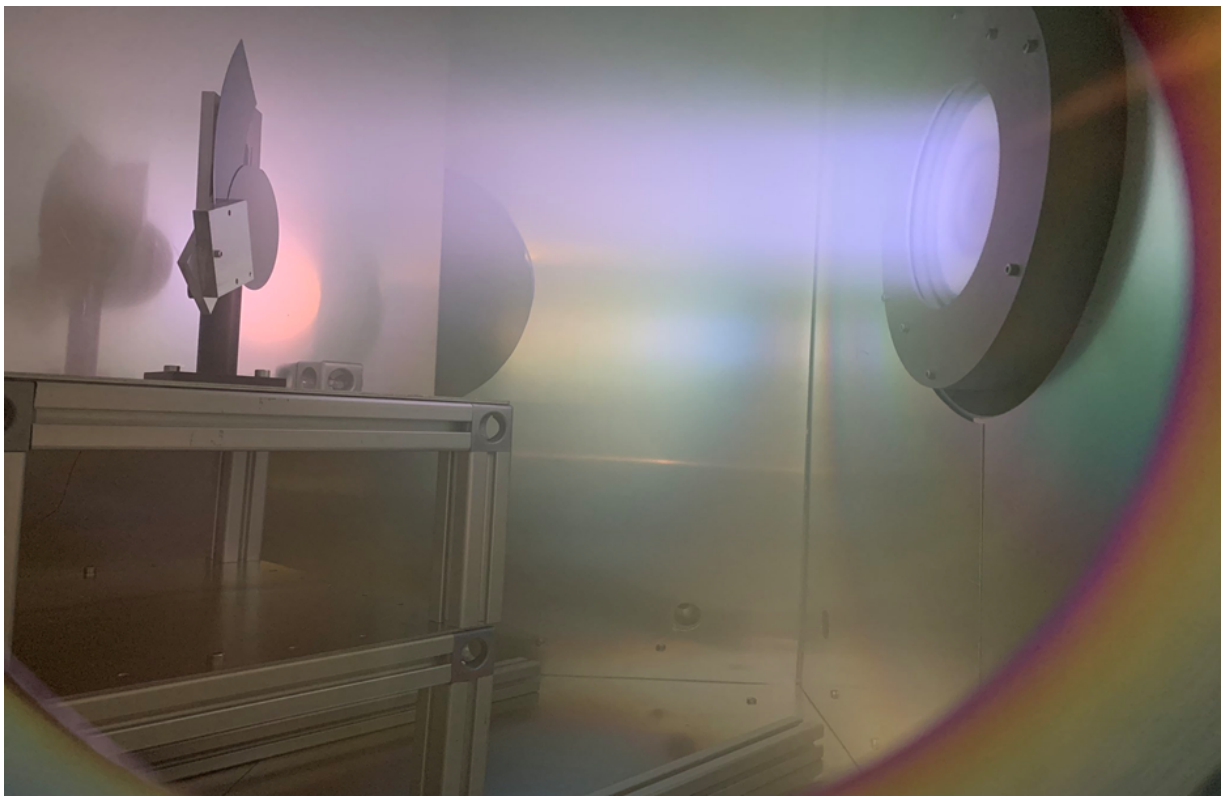

Figure 3.54: Image of the ignited argon ion beam during the initial etch experiment in the large UHV chamber. Image courtesy of Prof Stuart Reid [296].

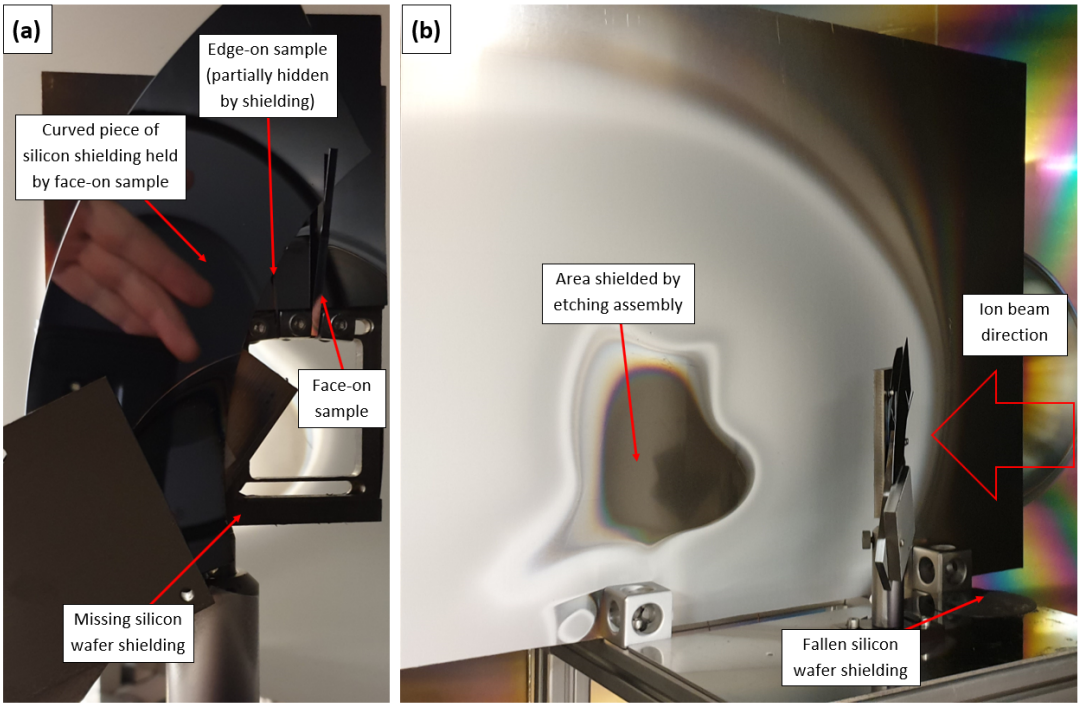

Figure 3.55: (a) Collapsed etching assembly post-etch. (b) Resultant etching and sputtering pattern on the aluminium shield after etching.

Upon opening the chamber, it was noted that the front silicon wafer had fallen off the etching jig assembly. The large arc of silicon had fallen forward and was stopped only by the face-on sample getting caught behind it. This subsequently shielded the majority of the top of the edge-on sample from the ion beam, as shown in figure 3.55(a). It is possible that due to the high temperatures experienced within the chamber than thermal expansion of the clamping structures may have caused this collapse. As shown in figure 3.55(b), one also notes the substantial etching and sputtering patterns visible on the aluminium shield, especially the area that appears to be untouched which would have been directly shielded by etching jig assembly.

Due to the collapsed shielding, the etch rate will be difficult to estimate due to the fact the edge-on sample was shielded for an unknown amount of time from the ion etching being. The face-on sample was continuously exposed to the beam, although the orientation of the sample was no longer perpendicular to the ion beam due to the silicon shielding catching on the back of the sample and pushing it forward. Nonetheless, it is not expected this would drastically change the etching rate on the face-on sample and the bottom edge-on sample was always exposed to the ion beam throughout.

Sample	Dimension (μm)	Etch amount (μm) / Etch rate ($nm\,min^{-1}$)	Etch amount error (μm) / etch rate error ($nm\,min^{-1}$)
Face-on	496.4	20.7	0.4
Edge-on	2187.3	27.5	7.8

Table 3.13: Post-etch dimensions of silicon samples after initial etch test and corresponding etch rates.

The samples were re-measured, with only the bottom area of the edge-on sample being measured, and recorded in table 3.13. Since the etch time was 1000 minutes, both etch amount in μm and etch rate in $nm\,min^{-1}$ are recorded in the same table.

By taking measurement from only the areas on the ribbons that were fully exposed to the beams throughout, the etch rates are in agreement with one another which implies starting surface quality does not affect etching rates. This, combined with the smaller error, means the face-on sample results will be taken as the etch rate estimate from this experiment, $20.7 \pm 0.4\,nm\,min^{-1}$, around 14 times faster than argon ion etch #1 treatment.

Argon ion etch #2 - large chamber

A batch of 20 samples were used, cut with the new laser cut technique, the same as the *control new* samples. Unlike section 3.7.2.1, a subset of 5 samples were SEM imaged on both faces and edges before being etched.

All 20 samples were individually profiled across their width on the thin fibre profiler since this is the direction in which they will be etched. The samples were then placed into the etching

jig, exposing the top edges for treatment first, with 2 sacrificial samples of the same dimension placed at either end.

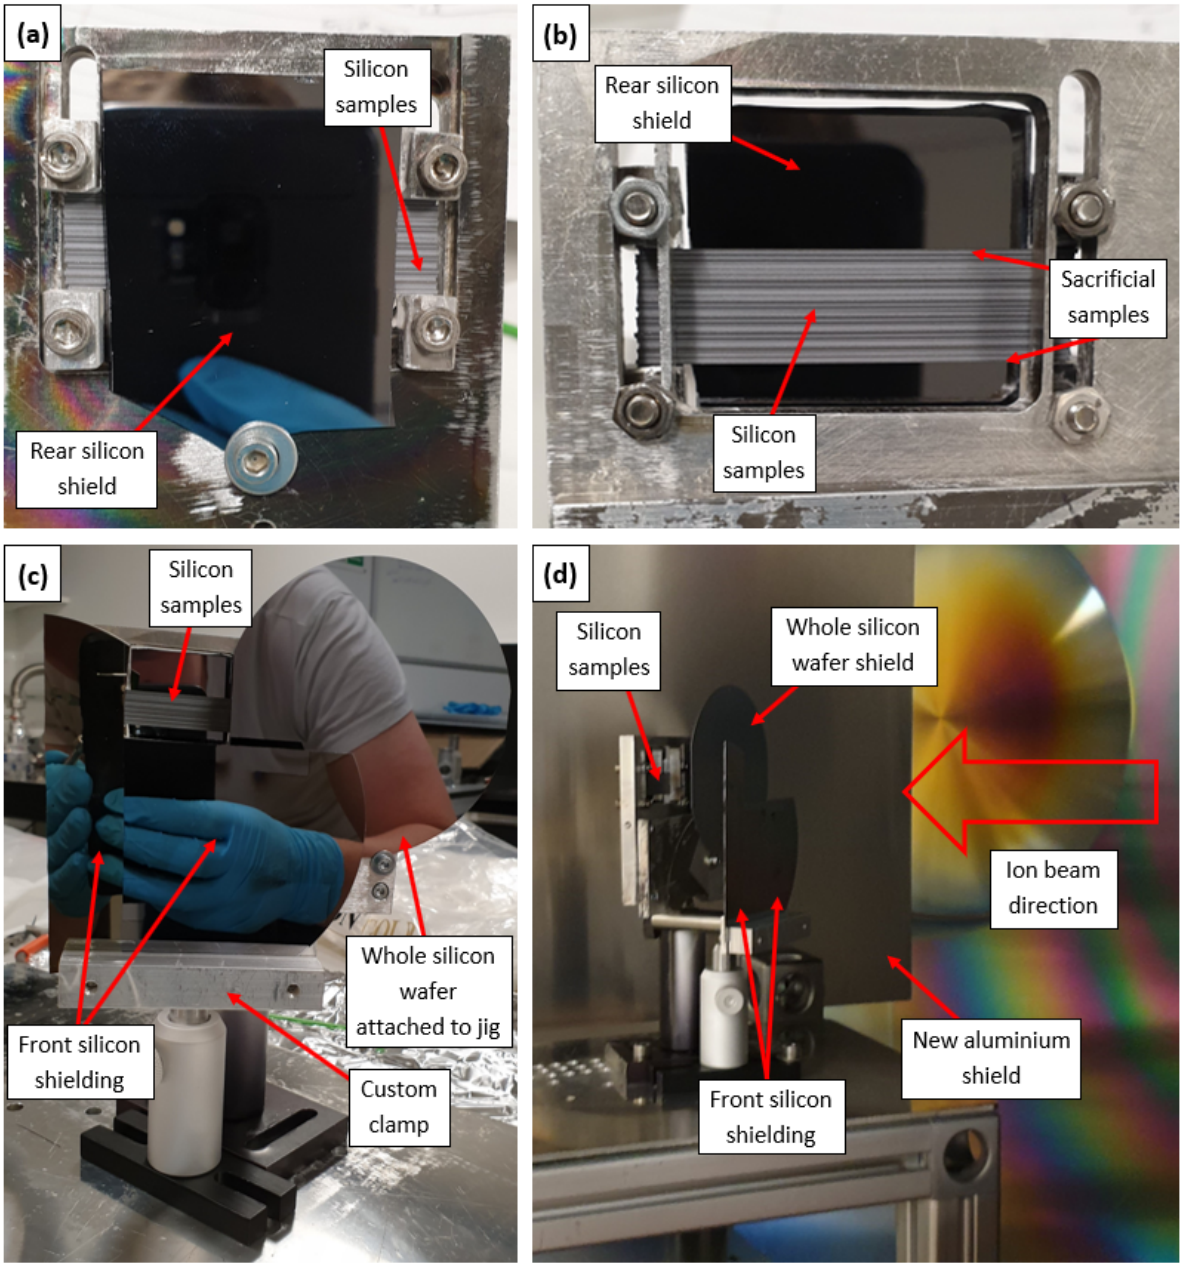

Figure 3.56: (a) Rear of the etching jig with silicon shield. (b) Front of the etching jig and top edges of the samples to be treated. (c) Silicon shielding arrangement. (d) Etching jig assembly in the large UHV chamber, showing the separation of the front and back shielding components.

Due to the collapse of the previous etching assembly, a new shielding configuration was developed. As shown in figures 3.56(a) and (b), this involved turning the jig on its side and securing a piece of silicon material at the back with washers, to rest again the back of the samples and block the ion beams line-of-sight through the aperture and minimise sputtering onto the back of the samples. Shown in figure 3.56(c) is a custom clamping set-up devised in order to hold

a straight silicon and quarter silicon wafer shield in front of the bottom and left hand side of the jigs. Due to this being held on a different clamp, these two pieces at approximately 5 cm forward of the whole wafer and etching jig assembly, as shown in 3.56(d). From the line-of-sight perspective of the ion gun, this should still shield the jig sufficiently from the ion beam. This design afforded more secure clamping of the silicon shielding. A small section of the top of the etching jig was left unshielded but this was deemed acceptable.

Parameter	Value	Unit
Beam current	300	mA
Beam potential	1.5	kV
Acceleration voltage	100	V
Acceleration current	8	mA
Ion source emission voltage	41	V
Ion source emission current	450	mA

Table 3.14: Parameters of argon ion etch #2 applied to top and bottom edges of the samples.

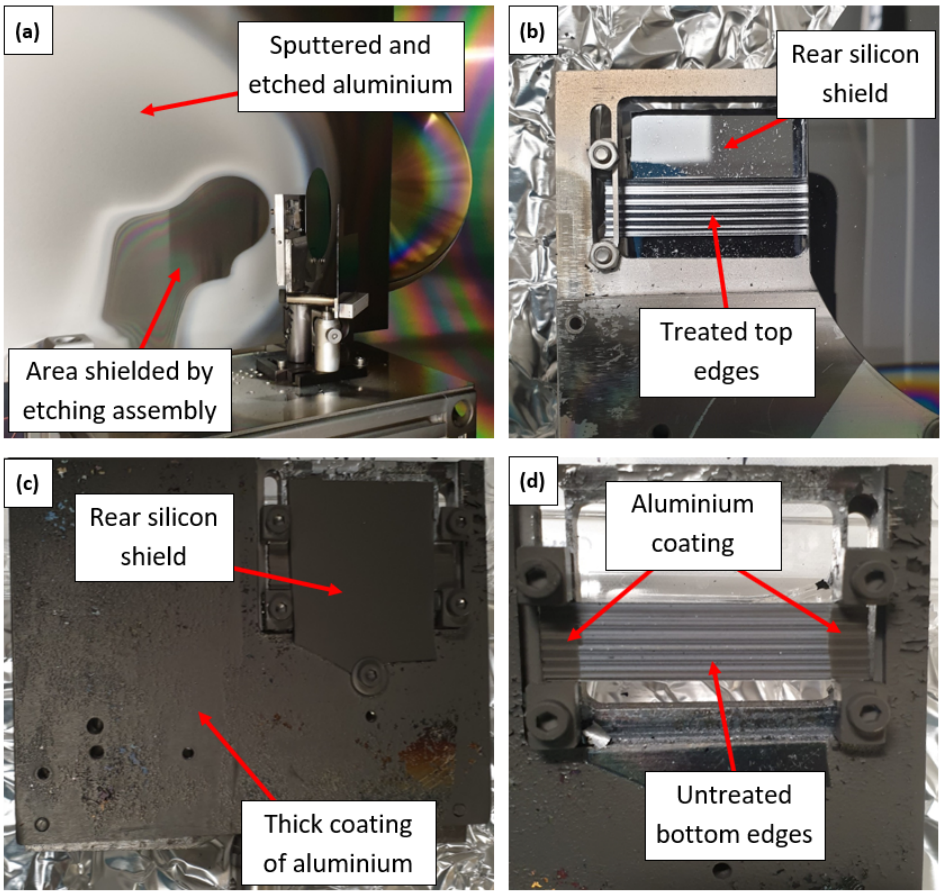

Figure 3.57: (a) Etching jig assembly and aluminium shield after the first half of etching. (b) Top edges of the samples after etching. (c) Back of etching jig covered in a thick coating of sputtered aluminium. (d) Untreated bottom edges of the samples with a coating of aluminium at either end.

A new, clean, 1 mm thick aluminium shield was placed into the chamber for this etch. The parameters in table 3.14 were set and the samples were etched. The samples were exposed to the ion beam for 17 hours, or 1020 minutes. A maximum temperature of 195 °C was recorded on the aluminium plate. This managed to partially buckle the aluminium plate. As show in figure 3.57(a), sputtering and etching of the plate was observed, as seen before in the initial etching experiment. So before starting the second etch, a new, clean, 3 mm aluminium shield replaced this.

As shown in figure 3.57(b), the top edges of the sample appeared to be shiny after the etching treatment, consistent with argon ion etch #1 results. A particularly thick layer of sputtered aluminium was observed, as shown in figures 3.57(c) and (d). After removing the rear silicon shield, a ruler measurement confirmed this extended approximately 5 mm from either end of each sample and so was of no concern for the strength testing process.

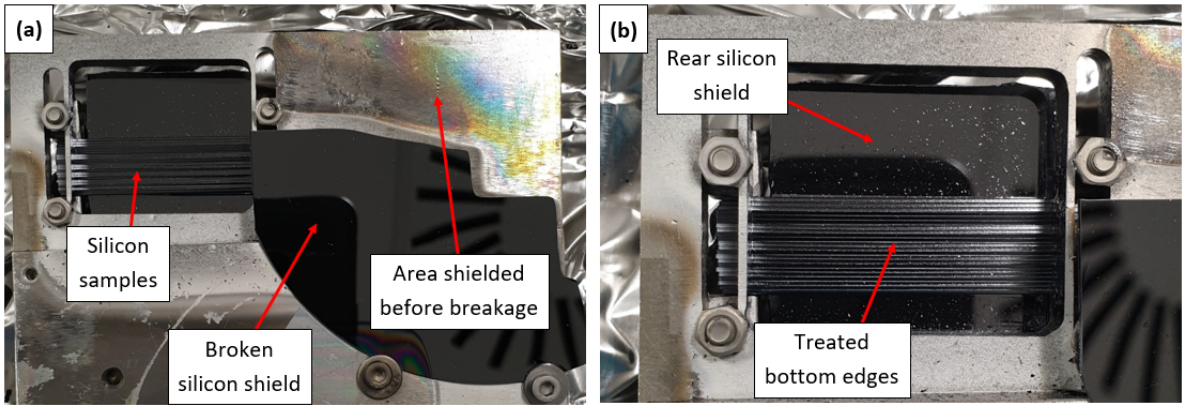

Figure 3.58: (a) Front of the etching jig and broken silicon shield. (b) Bottom edges of the samples after etching. (c) Back of etching jig covered in a thick, uniform coating of sputtered aluminium. (d) Treated top edges of the samples with a coating of aluminium at either end.

The bottom edges of the samples were exposed to the ion beam for 1020 minutes with the parameters set in table 3.14. Due to the thicker aluminium plate, a maximum temperature of only 145 °C was recorded. However, enough heat was generated to likely cause the failure of the whole silicon shield attached to the etching jig. Figure 3.58(a) shows what appears to be heat stress damage to the wafer causing it to cleave close to the two crystal axes aligned with and perpendicular to the wafer flat. It was not known when this breakage occurred but this did not cause any issues with the etching of the samples themselves. As shown in figure 3.58(b), like the top edges, the bottom treated edges appear shiny. A more uniform thick aluminium coating was observed from further aluminium sputtering onto the back of the etching jig. The previously treated top edge samples were coated at either end with this aluminium coating, as with the bottom edges, similar to what is shown in figure 3.57(d). Again, this coating only extends to approximately 5 mm from the sample ends and is not a concern.

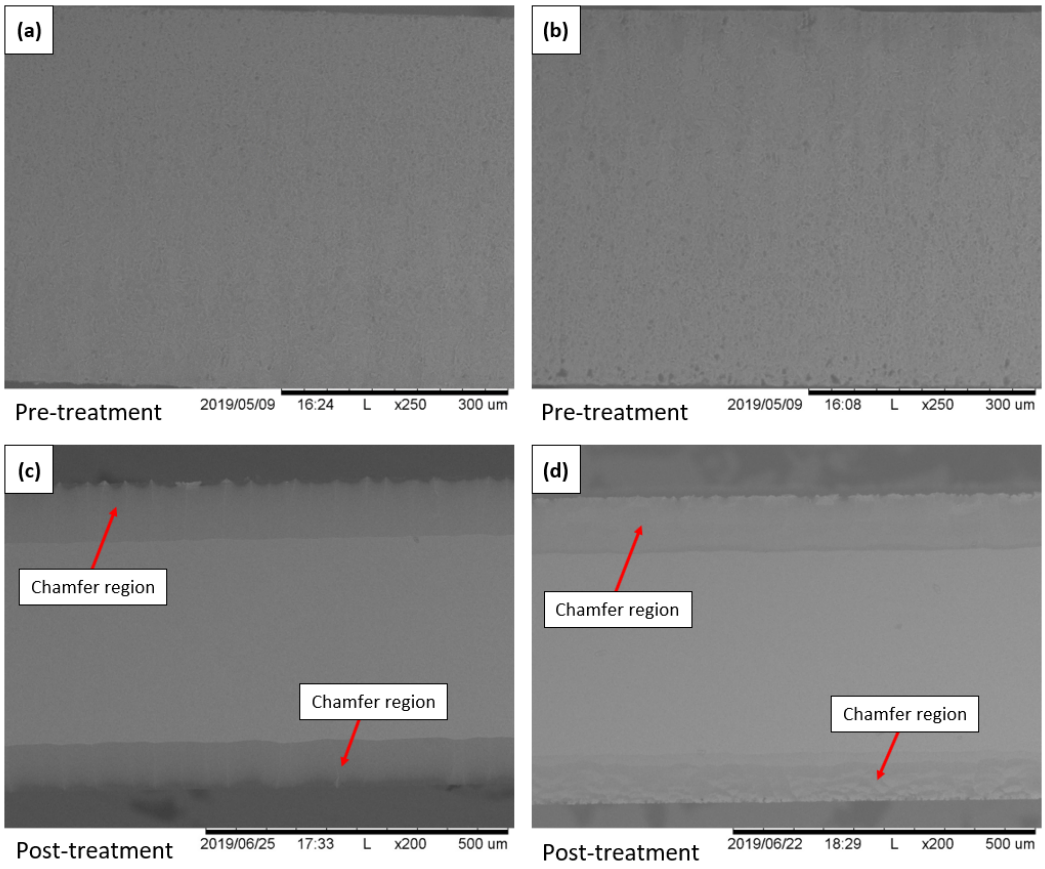

Figure 3.59: (a) SEM imaging of the bottom edge of sample 4084, before argon ion etching. (b) Top edge of sample 5094, before argon ion etching. (a) and (b) are magnified 250×. (c) Bottom edge of sample 4084, after argon ion etching. (d) Top edge of sample 5094, after argon ion etching. (c) and (d) are magnified 200×.

A subset of 5 samples were imaged under the SEM in order to see the effect of the argon ion etch #2 treatment on the edges of the samples. Figure 3.59 compares typical sample edges before and after the etching process. Again, an improvement on both top and bottom treated edges was observed, alongside chamfering features, as seen in figures 3.59(c) and (d).

Both argon ion etching features seen on the sample edges etched in the smaller chamber, in section 3.7.2.1, appeared again. As shown in figure 3.60(a), the etching into the vertices of the samples, due to the suspected ricocheting of ions between the polished faces, is far more pronounced in argon ion etch #2 compared to argon ion etch #1. This effect appears to extend to around 100 µm from the top and bottom vertices, leaving only approximately 300 µm of the edges perpendicular to the face, with the rest subsumed into the chamfered region. This larger chamfering effect is potentially due to the substantially higher etching rate of this process.

One also notes the return of the cell-like structure as highlighted in figure 3.60(b). The cell size appears to have approximately doubled or tripled in size to between 40–60 µm. Larger cell-like structures are generally consistent with deeper etches [296].

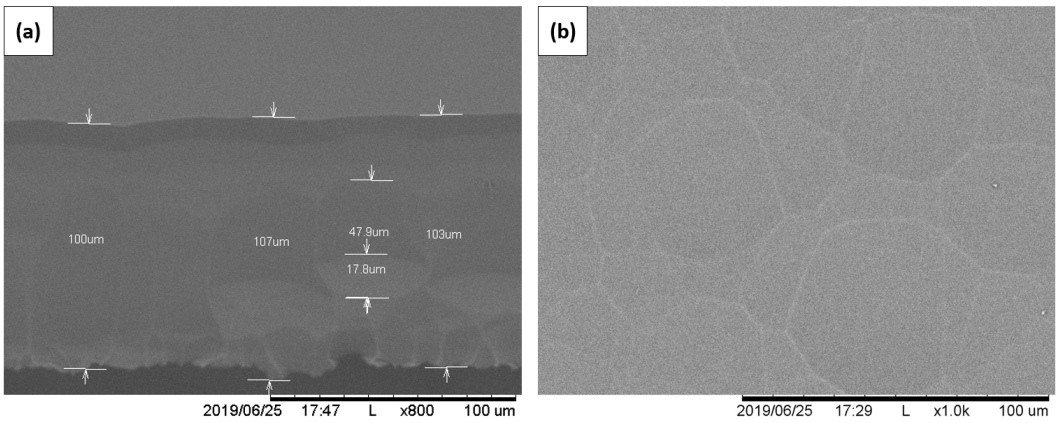

Figure 3.60: (a) SEM imaging of the bottom vertex of the bottom edge of sample 5094, after argon ion etching, magnified 800×. (b) Bottom edge of sample 4079, after argon ion etching, magnified 1000×. (b) has been edited to increase the contrast to clarify the cell-like structure.

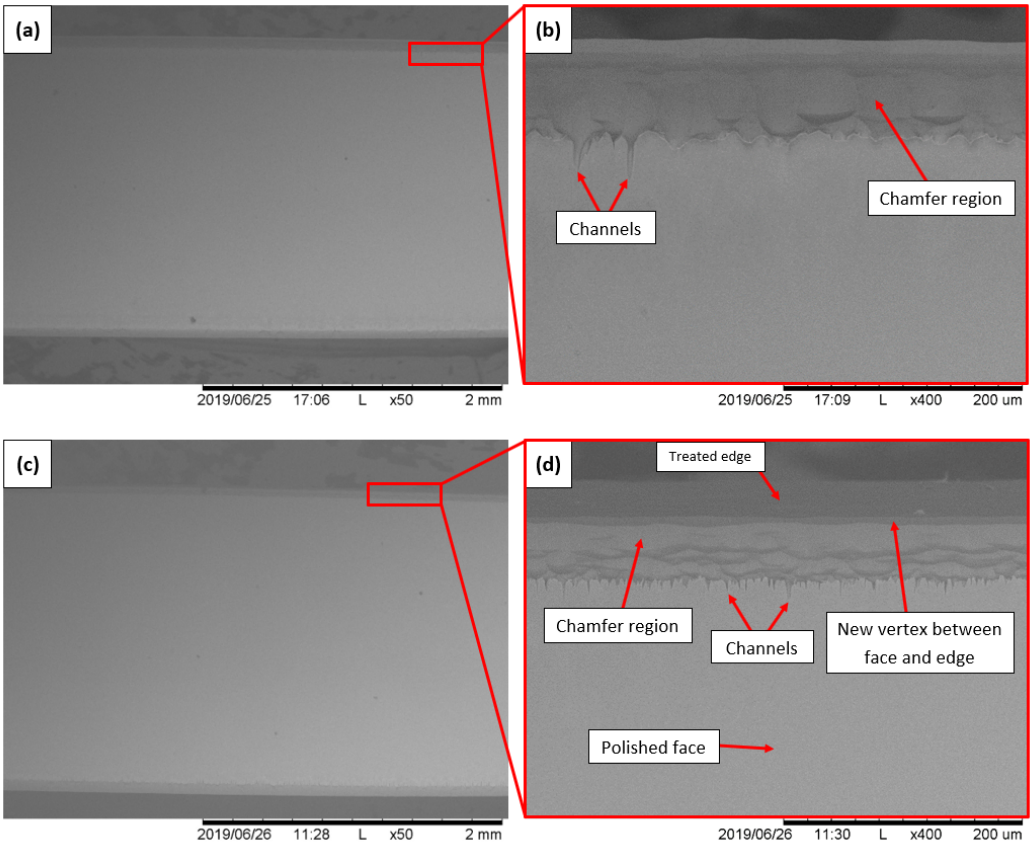

Figure 3.61: (a) SEM imaging of the back face of sample 4090. (b) Back face of sample 4090, zoomed in on the bottom treated edge. (c) Front face of sample 4090. (d) Front face of sample 4090, zoomed in on the top treated edge. (a) and (c) are magnified 50×. (b) and (d) are magnified 400× and edited to increase contrast.

Investigating the faces, one witnesses similar etching on the vertices of the samples when looking perpendicular to the faces as shown in figures 3.61(a) and (c). The depth of this etching

appears to be similar to that observed for argon ion treatment #1. Figure 3.61(b) appears to show a depth of around 100 µm, giving a chamfer angle of approximately 45° into the page. Figure 3.61(d) appears to show a depth of around 50 µm, giving a chamfer angle of approximately 30° into the page. Similar channels, presumably formed by the most energetic ions that penetrate between the sample edges and along the faces, are seen on the top and bottom areas of the faces as shown in figures 3.61(b) and (d).

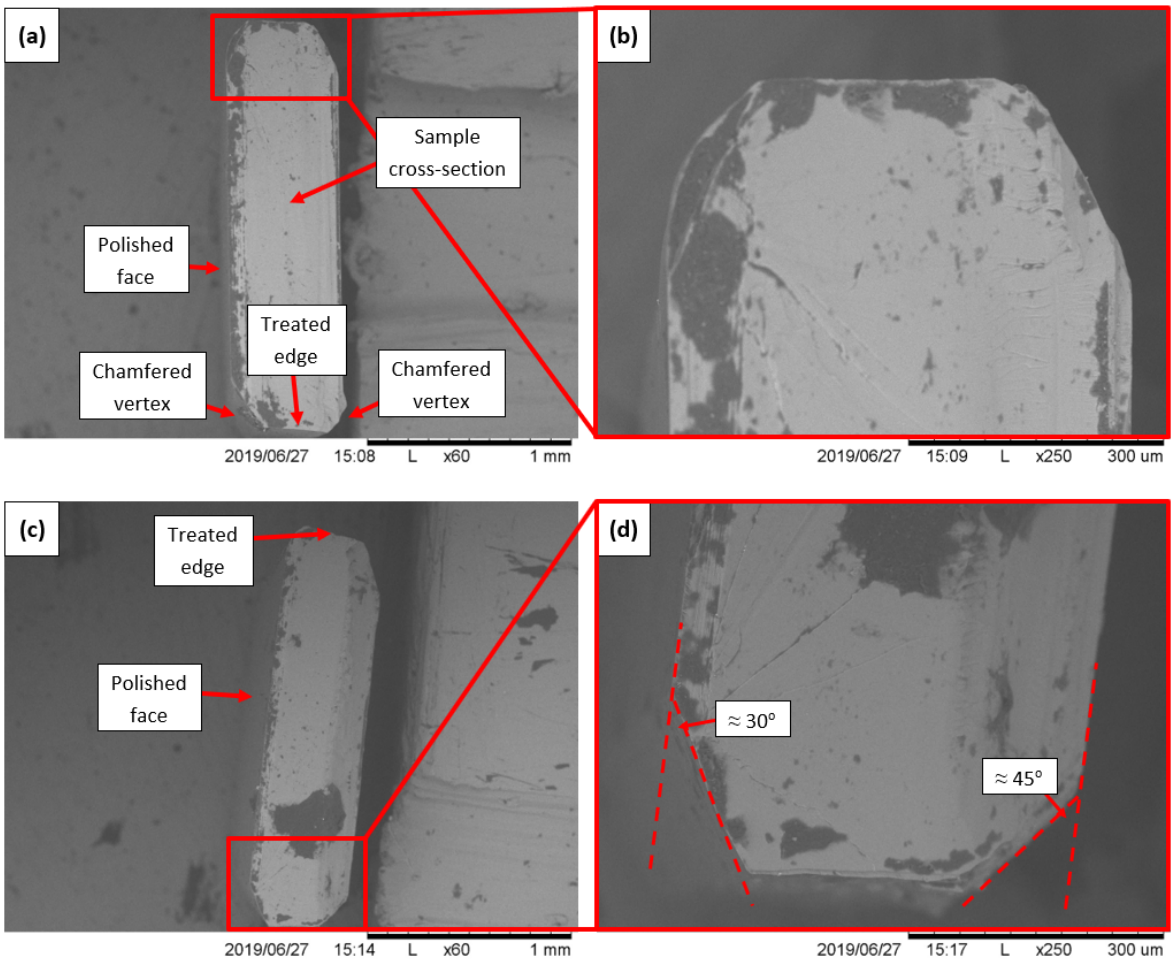

Figure 3.62: (a) SEM imaging of the cross-section of one of the sacrificial samples. (b) Zoom in on the top section of (a) to highlight chamfering effects along with SEM calibrated measurements. (c) Cross-section of the other sacrificial sample. (d) Zoom in on the bottom section of (c) to highlight chamfering effects. (a) and (c) are magnified 60×. (b) and (d) are magnified 250×.

To investigate this apparent chamfering effect, both sacrificial samples were deliberately broken in order to SEM image their cross-section. As can be seen in figure 3.62, a chamfering of the vertices has indeed taken place during the argon ion etching process. Figures 3.62(a) and (c) shows the cross-section of both of the sacrificial samples. The vertices of the sample have been chamfered although the chamfering does not appear to be even on each vertex. It appears that one vertex has a shallower chamfer at an approximately 30°, with the other at approximately

45°, consistent with observations made above. The appearance of this effect on diagonal vertices implies a potential angular offset within the etching jig assembly which biases one edge over the other. Figures 3.62(b) and (d) appear to show chamfer lengths and depths consistent with edge and face SEM imaging. It is not unreasonable to assume the same chamfering effect is present on the argon ion etch #1 samples, albeit to a lesser extent.

One might expect an increase in strength due to chamfering reducing stress concentrators on the sample vertices. Regardless, it is interesting to note this technique's ability to round or smooth a ribbon geometry at μm scales and below.

3.7.3 Strength results and error analysis

Both argon ion treatments were strength tested in the same way as the control samples. Both argon ion treatments improved the median strength over the *control new* and *control crystal* samples, as shown in figure 3.63.

For argon ion etch #1, sample 1007 slipped from the epoxy without breaking. However, this sample recorded the second highest value of the whole set. An asterisk (*) has been placed on the corresponding strength marker to denote that this sample did not fail but the sample is still valid as it experienced *at least* a stress of 175 MPa before slipping. The sample was re-epoxied and re-tested yielding a result of 129 MPa. This lower result is not included in any analysis or plotted on figure 3.63.

For argon ion etch #2, sample 4081 was dropped during post-treatment cleaning. No visible damage was observed on the sample, however this was the weakest sample in the set. A total of 8 samples were misaligned within the fuse end assembly, however this misalignment was much less in comparison to figure 3.36(a), for example.

Both argon ion surface treatments improved the median strength over the *control new* and *control crystal* samples.

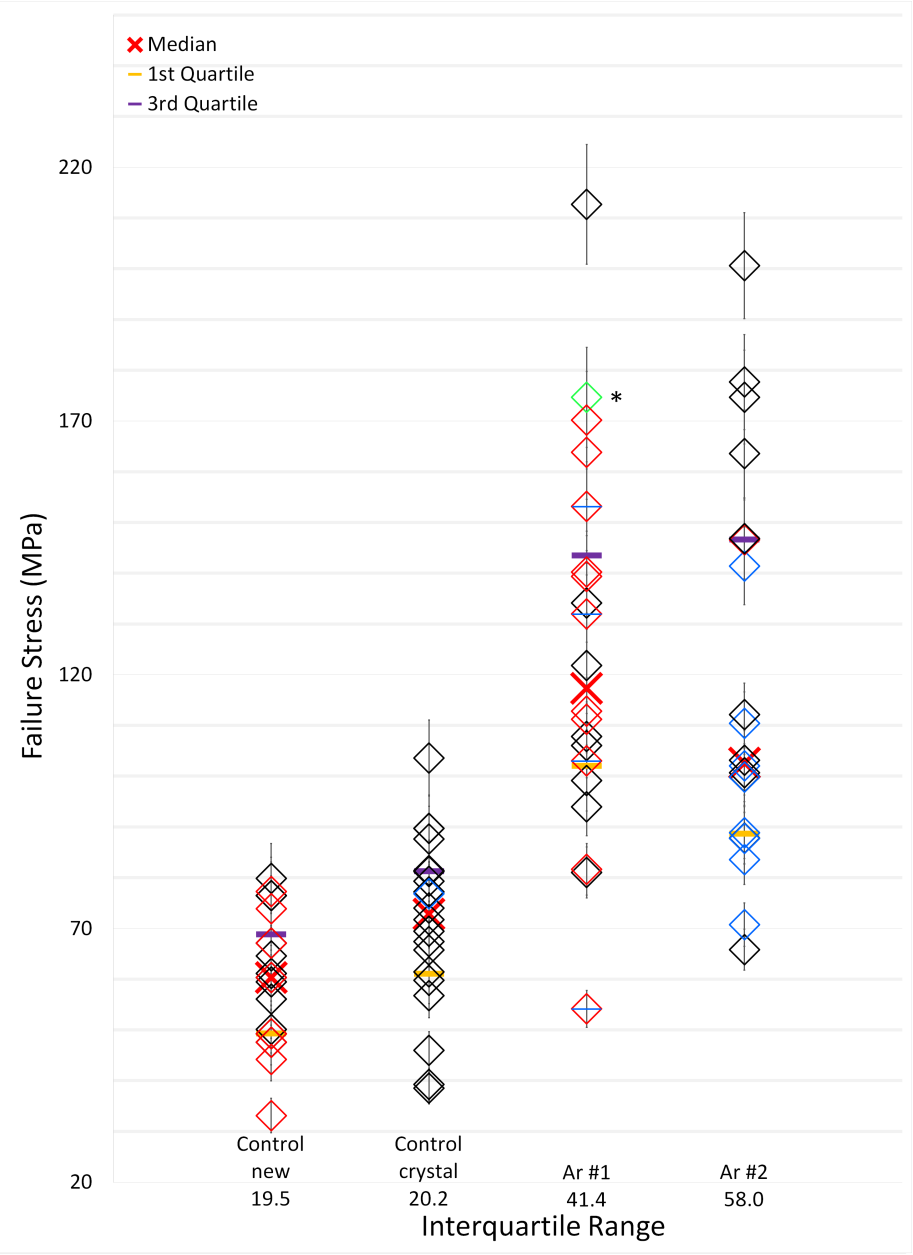

Figure 3.63: Results of tensile strength testing on both argon ion treatments. ◇ indicates sample failed at fuse end interface. ◇ indicates a sample was noticeably misaligned in the fuse end assembly. ◇ indicates the sample pulled free from a fuse end due to epoxy failure so the sample was re-epoxied into a new fuse end and re-tested. Where more than one anomaly exists on a sample, the second anomaly is marked with a corresponding horizontal coloured bar.

3.7.3.1 Argon ion etching surface treatment #1 removal amount and errors

Sample	Pre-etch weight (mg)	Post-etch weight (mg)	Weight lost to etch (mg)	Volume lost per sample (μm^3)	Etch depth total (μm)	Etch rate ($nm\,min^{-1}$)
1001	116.2	114.8	1.4	6.0×10^{14}	31.0	2.2
1002	115.9	114.7	1.2	5.2×10^{14}	26.5	1.9
1003	116.6	114.9	1.7	7.3×10^{14}	37.6	2.7
1004	115.9	114.9	1.0	4.3×10^{14}	22.1	1.6
1005	116.3	114.9	1.4	6.0×10^{14}	31.0	2.2
1006	116.1	115.1	1.0	4.3×10^{14}	22.1	1.6
1007	116.2	114.9	1.3	5.6×10^{14}	28.7	2.1
1008	116.2	115.1	1.1	4.7×10^{14}	24.3	1.7
1009	116.3	115.3	1.0	4.3×10^{14}	22.1	1.6
1010	116.1	115.1	1.0	4.3×10^{14}	22.1	1.6
1011	116.1	115.2	0.9	3.9×10^{14}	19.9	1.4
1012	116.3	116.0	0.3	1.3×10^{14}	6.6	0.5
1013	116.3	115.4	0.9	3.9×10^{14}	19.9	1.4
1014	116.2	115.2	1.0	4.3×10^{14}	22.1	1.6
1015	115.6	115.4	0.2	8.6×10^{13}	4.4	0.3
1016	116.1	115.4	0.7	3.0×10^{14}	15.5	1.1
1017	116.1	115.3	0.8	3.4×10^{14}	17.7	1.3
1018	116.1	115.4	0.7	3.0×10^{14}	15.5	1.1
1019	116.3	115.6	0.7	3.0×10^{14}	15.5	1.1
1020	116.0	116.0	0.0	0.0	0.0	0.0
					Average	1.5
					Error	0.8

Table 3.15: Pre and post-etch weight of samples with corresponding weight lost to etch.

Table 3.15 shows the pre-etch and post-etch weights of the individual samples, once both sides have been treated.

One notes from table 3.15, sample 1020 appears to have lost no weight during the etching process. It is possible that sputtering onto the opposing edge of the sample was equal to the amount etched off the exposed edge, especially if misalignment or offset from the argon ion beams occurs. Alternatively, the weight etched may be within experimental error.

One also notes that the lower numbered samples, located on the left hand half of the etching aperture as exposed to the argon ion beam, appear to have been etched at a greater extent than the samples located on the right hand side. This may be further evidence of misalignment or offset of the etching jig or beams, where the profile of the beams is expected to be Gaussian in shape [296].

Table 3.15 also details the estimated etch rate and amount etched on each sample. The figures work on the assumption that the samples were exposed to the ion beams along 37 mm of their length as 3 mm at the top and 5 mm at the bottom of the samples were not etched at all

due to being fully shielded by the silicon and jig apparatus. The etch depth total corresponds to total material removal amount from both top and bottom edges of the silicon sample. As such, the etch rate is calculated by dividing this value by 2, since both edges were exposed to the ion beams for a similar length of time, and dividing this by the average etch time from the top and bottom edges, 6975 minutes.

One can see the error has been reduced by individually weighing the samples before and after the etch, although the error still remains very high at 53 % of the average etch rate determined. However, the etch rate is consistent with the experimental results shown in the initial experiment, even with the large associated errors and is of the order of single figure $\mathrm{nm\,min^{-1}}$. The average total etch depth is 20.2 µm with a combined error of 3.4 µm.

Material removal was measured by weighing only and no dimensional measurement was taken after etching. In order to determine the cross-sectional area, a measurement of the width of 5 random samples of the 20 samples treated in this batch was obtained prior to etching. This was part of the standard check of laser cutting tolerance, as measured by the digital callipers. The average width at either end of the 5 samples was determined as 2.19 µm. The associated error from these 10 measurements, taking the standard 130 µm calliper error for width measurement, yields an error on this value of 41 µm. The combined weighing errors carried through to the total etch depths listed in table 3.15 yields a total etch depth error of 11 µm for all samples. This error can then be combined with the average width and its error, to calculate a final post-etch width error of 42 µm for all samples. The standard wafer manufacturing tolerance of ± 25 µm was applied to the thickness of the samples.

3.7.3.2 Argon ion etching surface treatment #2 removal amount and errors

Table 3.16 shows the pre-etch and post-etch widths of the individual samples as measured on the thin fibre profiler. The standard error of 0.08 %, taken from table 3.2, is applied to each measurement. This is then used to calculate the material amount etched from each sample and the corresponding etch rate.

The data from table 3.16 implies relatively even etching across all samples although some variability in etch rate and etch depth exists between samples. The average etch depth for argon ion etch #2 is 140.6 µm with a corresponding error of 0.5 µm. The average etch rate is $68.9\,\mathrm{nm\,min^{-1}} \pm 0.3\,\mathrm{nm\,min^{-1}}$.

As with section 3.7.3.1, the standard wafer manufacturing tolerance of ± 25 µm was applied to thickness of the samples.

Sample	Pre-etch width (µm)	Post-etch width (µm)	Etch depth total (µm)	Etch depth error (µm)	Etch rate ($nm\,min^{-1}$)	Etch rate error ($nm\,min^{-1}$)
4079	2214.7	2084.9	129.8	2.4	63.6	1.2
4080	2212.4	2081.4	131.0	2.4	64.2	1.2
4081	2204.4	2072.5	131.9	2.4	64.7	1.2
4082	2206.9	2074.3	132.6	2.4	65.0	1.2
4083	2204.4	2065.4	139.0	2.4	68.1	1.2
4084	2200.2	2064.6	135.7	2.4	66.5	1.2
4085	2207.1	2064.4	142.7	2.4	69.9	1.2
4086	2200.2	2069.2	131.0	2.4	64.2	1.2
4087	2204.9	2066.0	139.0	2.4	68.1	1.2
4088	2206.4	2056.4	150.0	2.4	73.5	1.2
4089	2204.3	2060.4	143.9	2.4	70.6	1.2
4090	2206.4	2057.4	148.9	2.4	73.0	1.2
4091	2208.9	2065.3	143.6	2.4	70.4	1.2
5092	2207.0	2059.7	147.3	2.4	72.2	1.2
5093	2207.3	2066.7	140.6	2.4	68.9	1.2
5094	2201.3	2058.0	143.3	2.4	70.2	1.2
5095	2204.2	2059.6	144.6	2.4	70.9	1.2
5096	2203.3	2067.8	135.5	2.4	66.4	1.2
5097	2211.9	2061.3	150.6	2.4	73.8	1.2
5098	2213.6	2061.9	151.7	2.4	74.4	1.2
		Average	140.6	0.5	68.9	0.3

Table 3.16: Pre-etch and post-etch widths of silicon samples exposed to argon ion treatment #2 as measured on the thin fibre profiler.

3.7.4 Discussion

Similar to section 3.6.3, the 8 misaligned samples from argon ion etch #2 appear to skew the strength results. If one removes both the 8 misaligned samples and the weakest sample from the set which was dropped, the median value of argon ion etch increased by 43 %, from 103 MPa to 147 MPa. This would push the argon ion etch #2 median value from 13 % below the median value of argon ion etch #1, to 25 % above.

It is difficult to draw conclusions on the advantage to either process over the other due to the large errors inherent from weighing samples in argon ion etch #1 and the misaligned samples from argon ion etch #2. However, one notes that similar median and maximum strengths are attained even though the amount of material removed during argon ion etch #2 is 7 times greater than argon ion etch #1. This implies that greater material removal using this method does not necessarily lead to greater tensile strengths. It is difficult to specify what would cause this effect during this treatment method. Since temperature measurements were not taken during argon ion etch #1, it is not possible to draw conclusions on temperature effects on the samples during etching. Alternatively, excessive etching of unwanted material during argon ion etch #2, such

as aluminium or silicon shielding, may have caused kinetic damage when coating the silicon samples. Indeed, one notes a large amount of particulates on the rear silicon shield in figure 3.57(c) which may have induced damage onto the samples.

The strongest sample from argon ion etch #1, yielded a 2.7× increase in tensile strength compared to the strongest sample from the *control new* set, with the strongest sample from argon ion etch #2 yielding a 2.5× increase. Comparing to the *control crystal* set, this reduces to 2.1× and 1.9× respectively.

The median values from both sets also increased compared to control samples. The median strength from argon ion etch #1 was 1.9× higher than the *control new* set and 1.6× higher than the *control crystal* set. The median strength from argon ion etch #2 was 1.7× higher than the *control new* set and 1.4× higher than the *control crystal* set. If one excludes the misaligned samples, these comparative median values increase to 2.4× and 2× higher respectively.

Sample 1007 is unique in being able to withstand a higher stress before pulling out of the epoxy-fuse end assembly. Due to the design of the strength tester, this high load *must* have been supported through the silicon ribbon itself as no other mechanical piece supports the bottom of the strength tester. As such, it is valid to record this higher stress as a *minimum* stress value for this sample.

The critical crack sizes for argon ion etch #1 stress range 54–213 MPa equate to 55.5–3.6 µm. These values seem to be consistent with visible edge damage observed in figure 3.49(d).

For argon ion etch #2, and after discounting the misaligned and dropped samples, the stress range was 88–201 MPa, resulting in critical crack sizes of 21.0–4.0 µm respectively. If one includes the dropped sample, the weakest result of 65.8 MPa, yields a critical crack size of 37.5 µm. A critical crack of this size would not have been clearly visible by eye as was noted during inspection after the sample was dropped. These values seem to be consistent with visible edge damage observed in figure 3.60(a).

There appeared to be no correlation between the exposed half-etched boundary at the bottom of the argon ion etch #1 samples and the failure point of the samples. Of the 20 samples tested, including sample 1007 that was re-tested, only 2 samples broke near the bottom of the sample, near to this half-etch boundary.

It is notable that the argon ion etch #2 rate appears to be 3.3× higher for the actual etch rate compared to the initial etch rate test. This is consistent with the argon ion etch #2 etched amount which is 3.4× higher than the initial etched amount. Parameters applied in both etches were confirmed to be the same and no explanation can be offered for this increase in etch rate for the latter experiment.

It is clear from the large uncertainties of the argon ion etch #1 results that weighing samples with such a small amount of material removed could be drastically improved by the use of dimensional non-contact imaging, such as the fibre profiling used for argon ion etch #2. If one were to continue with weighing samples, the error could be reduced by using a larger mass

to measure alongside the samples in order to reduce the percentage error. One could incorporate this larger mass into two separate masses separated such that one could weigh the sample ensuring only the end 8 mm contacted the weights, reducing any potential damage that may be inflicted when placing samples onto the microbalance directly. One could also weigh the samples in between treating the top and bottom edges. The risk of extra contamination would likely outweigh the benefit of determining a more accurate etch rate, especially if the etch times applied to both sides are close enough that such a slow etch rate would make little to no difference in the amount of material removed, as is the case for this treatment.

This treatment has the advantage of being non-contact, reducing the risk inducing surface damage from the treatment itself. However, due to substantial heating observed in argon ion etch #2, there is a risk of inducing inherent stress within ribbons undergoing treatment due to subsequent cooling after etching. Surface quality is not ideal, as cell-like features on the surface leave a non-uniform surface finish. These structures will increase the surface area and hence the overall surface to volume ratio of the ribbon, meaning there may be an overall increase in surface loss. This has not been measured or verified during the scope of this work.

This treatment has the added advantage of there being no particular reason why circular fibres could not be treated in this way. Some form of rotational mechanism would have to be employed within the vacuum chamber. This is not unrealistic as these chamber typically contain rotary mechanisms, and even planetary stages, on the doors to provide an even coating for samples when the apparatus is used conventionally. While there is no apparent reason for circular fibres not to be treated via this method, one can not easily scale ion sources to the sizes required for 3G suspension element sizes. Furthermore, even if scaling is possible, the nature of the ion beam is approximately Gaussian which may lead to practical issues of evenly etching large surface areas. However, this beam profile may be advantageous if one wishes to shape the geometry of silicon fibres or ribbons by varying the exposure time to the ion beam and so should not be discounted. It is unclear if this can be scaled up to larger ion sources with larger beam diameters which would be required for suspension-scale ribbons and fibres. Otherwise either the ion source itself or the ribbon or fibre to be treated would need to be moved lengthwise through the ion beam in order for the sample to be evenly treated along its length.

Some of these issues may be overcome with the manufacture of bespoke chambers for the purpose of etching suspension-scale fibres and ribbons however it does not seem likely that such systems would be able to reach the strengths required ultimately for 3G detectors due to the fundamental surface finish left behind. Cell-like structures will ultimately create stress concentrators that would not exist under smooth, uniform finished surfaces.

3.8 Wet chemical etching

Chemical etching of silicon is a well understood technique, commonly used for the purposes of shaping silicon at microscopic or nanoscopic scales, called microfabrication and nanofabrication respectively. This technique can produce miniaturised geometries that would be impossible using standard mechanical machining. The development of these technologies has allowed for the advancement of mechanical, electrical and optical devices that can perform complex tasks such as micro-electromechanical gravimeters [299].

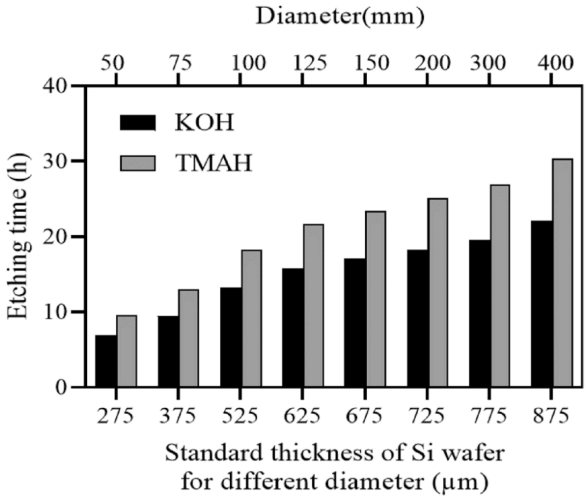

Figure 3.64: Etching time to form through holes in different silicon wafer thicknesses for two common silicon chemical etches, potassium hydroxide (KOH) and tetramethylammonium hydroxide (TMAH) [300].

As shown in figure 3.64, chemical etching can also be used on macroscopic scales due to very high material removal rates, with the ability of some chemical etches to dissolve through a 525 μm thick wafer in under a day [300]. As was shown in section 3.3, silicon ribbons from the previously published study were etched on macroscopic scales and showed the strongest non-oxidised sample results [136].

There are a multitude of factors to consider when determining a suitable etch solution for silicon. The primary considerations are the relative concentrations of each chemical in the solution, the strength or concentration of each chemical, the flow or recirculation rate and the temperature of the etching solution. Etching selectivity must also be considered for anisotropic etching processes. No masking was required for this work in order to remove all surface edge damage.

Practical considerations must also be taken into account, specifically, chemical compatibility of materials for both safety and to stop contamination of the etch solution, chemical toxicity and exposure for the operators conducting the etch, as well as etch solution heating and circulation where applicable.

3.8.1 Proposed chemical etches

Due to large etch rates and etch control, it was decided to pursue two different, but commonly used [301], chemical etches as surface treatments. The two processes selected differed primarily in their etching reaction with the crystallographic planes of silicon to see what effect this had on the tensile strength of the silicon ribbons.

3.8.1.1 HNA etch

The first etch treatment was an HNA etch, consisting of HF acid, nitric acid (HNO_3) and acetic acid (CH_3COOH). This etch is isotropic, meaning it etches silicon equally in all directions regardless of crystallographic orientation [302]. However, practical barriers are present when conducting this etch. Firstly, HF acid is particularly difficult to work with due to its acute toxicity and the corresponding strict laboratory safety controls required. Secondly, nitric acid is a strong oxidiser which comes with a separate set of safety protocols and, along with acetic acid, is also corrosive. Thirdly, the HNA etch is particularly sensitive to temperature [303] and relative acid concentrations, the latter of which is shown in figure 3.65(a) [244], so accurate temperature stabilisation and solution quantities are required.

Schwartz and Robbins 1976 paper has extensive information on this particular etch, including the finished surface quality depending on the ratio of the HNA etch acids, shown in figure 3.65(b) [244].

The HNA etch works via the following two simplified chemical reactions [244] [302]. Firstly, the oxidation of silicon to silicon dioxide (SiO_2) by the nitric acid, as shown in equation 3.10.

$$3\,Si + 4\,HNO_3 \longrightarrow 3\,SiO_2 + 4\,NO + 2\,H_2O \tag{3.10}$$

Secondly, the HF acid then dissolves the newly formed silicon dioxide, exposing a new silicon surface ready to be oxidised, as shown in equation 3.11.

$$SiO_2 + 6\,HF \longrightarrow H_2SiF_6 + 2\,H_2O_2 \tag{3.11}$$

One notes the lack of acetic acid in the above reactions, however the acetic acid plays no part in the etching process itself, rather it acts as a diluent for the etchant. While water can be used as the diluent for this reaction, it is more common to use acetic acid as this also prevents excessive dissociation of HNO_3 in the etching solution [302] [304].

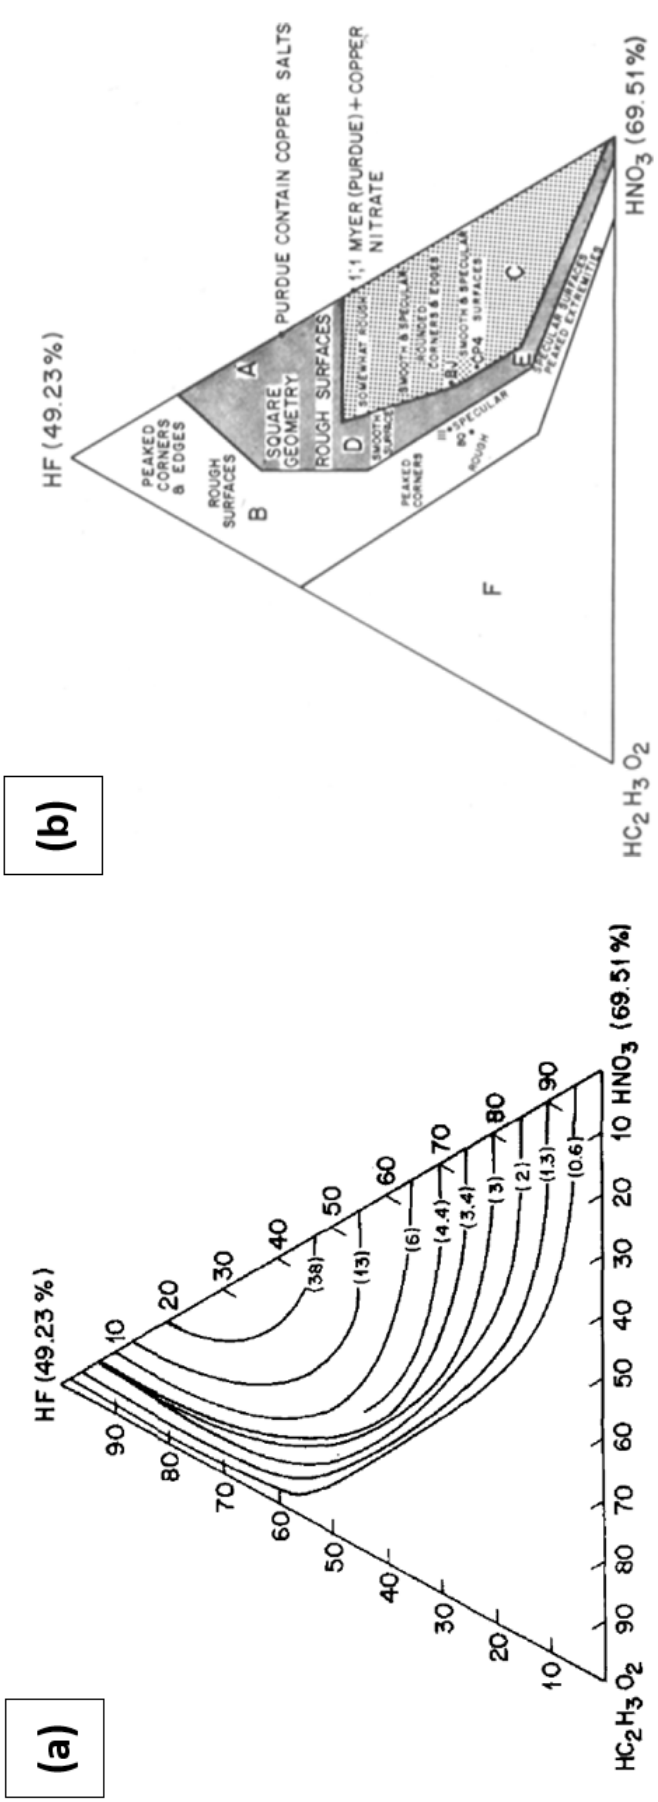

Figure 3.65: (a) Curves of constant HNA etch rate (mils per minute) as a function of etchant composition [244]. (b) Resultant geometry of the HNA etched surface as a function of etchant composition [244].

3.8.1.2 KOH etch

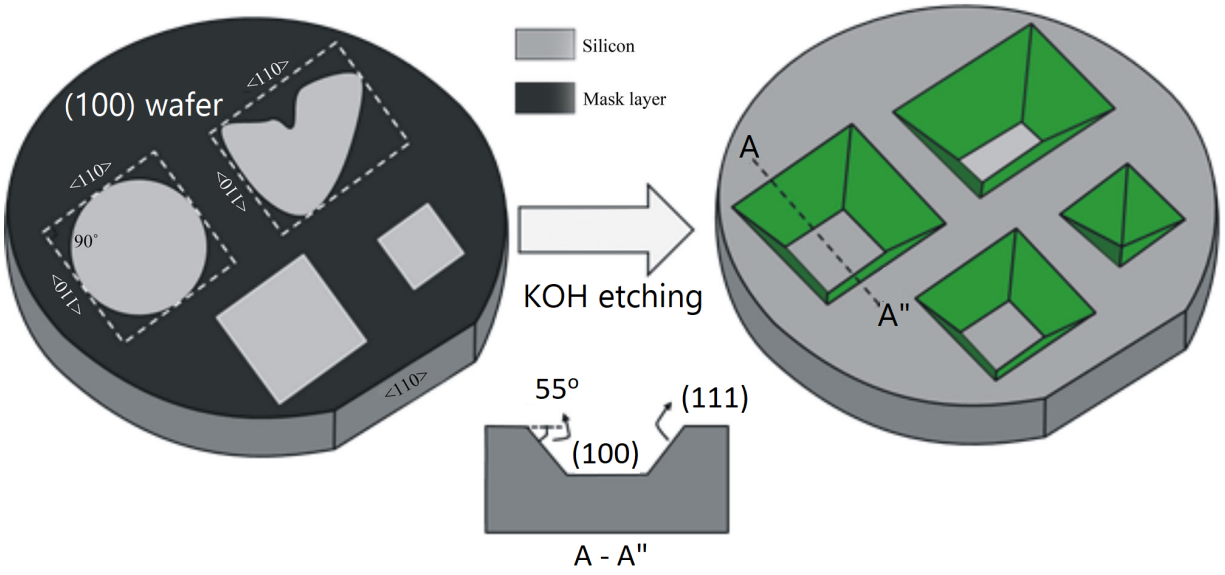

Figure 3.66: Schematic of etch direction of KOH into a (100) silicon wafer. Image modified from [305].

The second etch treatment selected was a potassium hydroxide (KOH) etch. KOH is a strongly alkaline solution. This etch is anisotropic, meaning it etches silicon preferentially dependent on the crystallographic orientation [306]. For (100) silicon, KOH will etch 55° into the face of the (100) silicon wafer, as shown in figure 3.66.

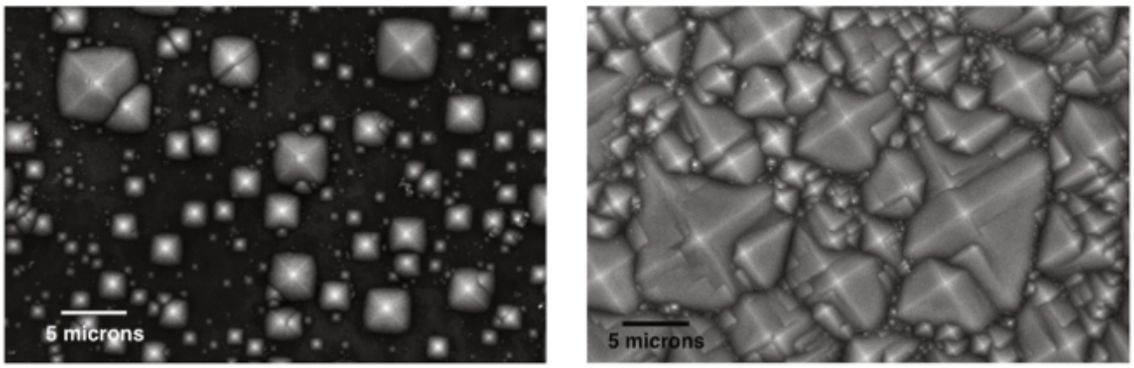

Figure 3.67: Examples of hillocks formed on a (100) silicon surface under different concentrations of KOH. Image modified from [307].

Due to the preferential etch, (100) plane surfaces will likely produce pyramidal surface features, known as hillocks, as shown in figure 3.67. Hillocks can be minimised or completely removed by agitation of the solution, increasing the KOH concentration greater $\geq$45 % and reducing the temperature of the solution [308] [309].

The KOH etch works via a series of chemical reactions shown in equations 3.12, 3.13 and 3.14 [310]. Firstly, chemical oxidation in which silicon becomes hydroxyl-terminated silicon

(Si(OH)). After this termination, Si-Si bonds in the material become significantly polarised due to large electronegativity of oxygen that results in weakening of these bonds. These Si bonds are then easily attacked by polar water molecules, leading to the removal of silicon atoms as a Si(OH) product [306] [309] [310].

$$\mathrm{Si} + 2\,\mathrm{OH^-} \longrightarrow \mathrm{Si(OH)_2^{2+}} + 2\,\mathrm{e^-} \tag{3.12}$$

$$\mathrm{Si(OH)_2^{2+}} + 2\,\mathrm{OH^-} \longrightarrow \mathrm{Si(OH)_4} + 2\,\mathrm{e^-} \tag{3.13}$$

$$\mathrm{Si(OH)_4} + 4\,\mathrm{e^-} + 4\,\mathrm{H_2O} \longrightarrow \mathrm{Si(OH)_6^{2-}} + 2\,\mathrm{H_2} \tag{3.14}$$

The $Si(OH_4)$ molecule is soluble in water. The main by-product of the reaction is hydrogen gas, which forms as bubbles on the surface of the silicon that can hinder etching if they stick to the surface as they effectively stop any further reaction at the bubble site, leading to a non-uniform and rougher surface [311]. It is common to etch with an admixture of KOH and IPA [312]. This greatly increases the surface finish of the etch by providing a Teflon-like surface on the etched silicon meaning it is more difficult for hydrogen bubbles produced from the etch to adhere to the silicon surface [313].

Concentration of the KOH solution, combined with the temperature, will define the etch rate of the solution [312]. These parameters must be balanced for practicality, optimal etch rate and minimisation of any features such as hillocks in order to reduce stress concentrators or defects on the crystal surface.

3.8.2 Chemical etching surface treatments

Initial chemical etching work was conducted at the James Watt Nanofabrication Centre (JWNC), University of Glasgow [314], and an external metal fabrication plant, ITM Mechanical Services Ltd, in Leven, Fife [315]. Both of these organisations have extensive experience with wet chemical etching.

3.8.2.1 Initial HNA experiments

An initial HNA etch was conducted at ITM on 5 silicon offcuts, cut using the new LML laser cutting technique. The emphasis for this experiment, referred to as "ITM HNA", was placed on surface quality rather than etch rate.

By referring to figure 3.65(b), an etch recipe that should leave a surface finish equivalent to the middle area in area C was developed where the surfaces are described as "smooth and specular". The HNA recipe used was 10 % by volume of CH_3COOH at 100 % concentration, 55 % by volume of HNO_3 at 70 % concentration and 35 % by volume of HF at 49 % concentration.

From figure 3.65(a), this would imply an etch rate of $139\,\mu\mathrm{m\,min^{-1}}$ at room temperature [244]; therefore a 1 minute and 4 seconds etch time would yield the desired surface removal

amount of 150 µm.

After etching the samples were wiped with IPA, followed by a rinse in tap water. The samples were then oven baked at 60 °C for 15 minutes to remove any remaining liquid.

The samples appeared to be very similar to when they left except the edges of the samples were slightly shinier. No difference in the amount of material etched was able to be measured. As stated, the etch rate was not the primary aim of this experiment so an accurate measurement of the amount of material removed was not required.

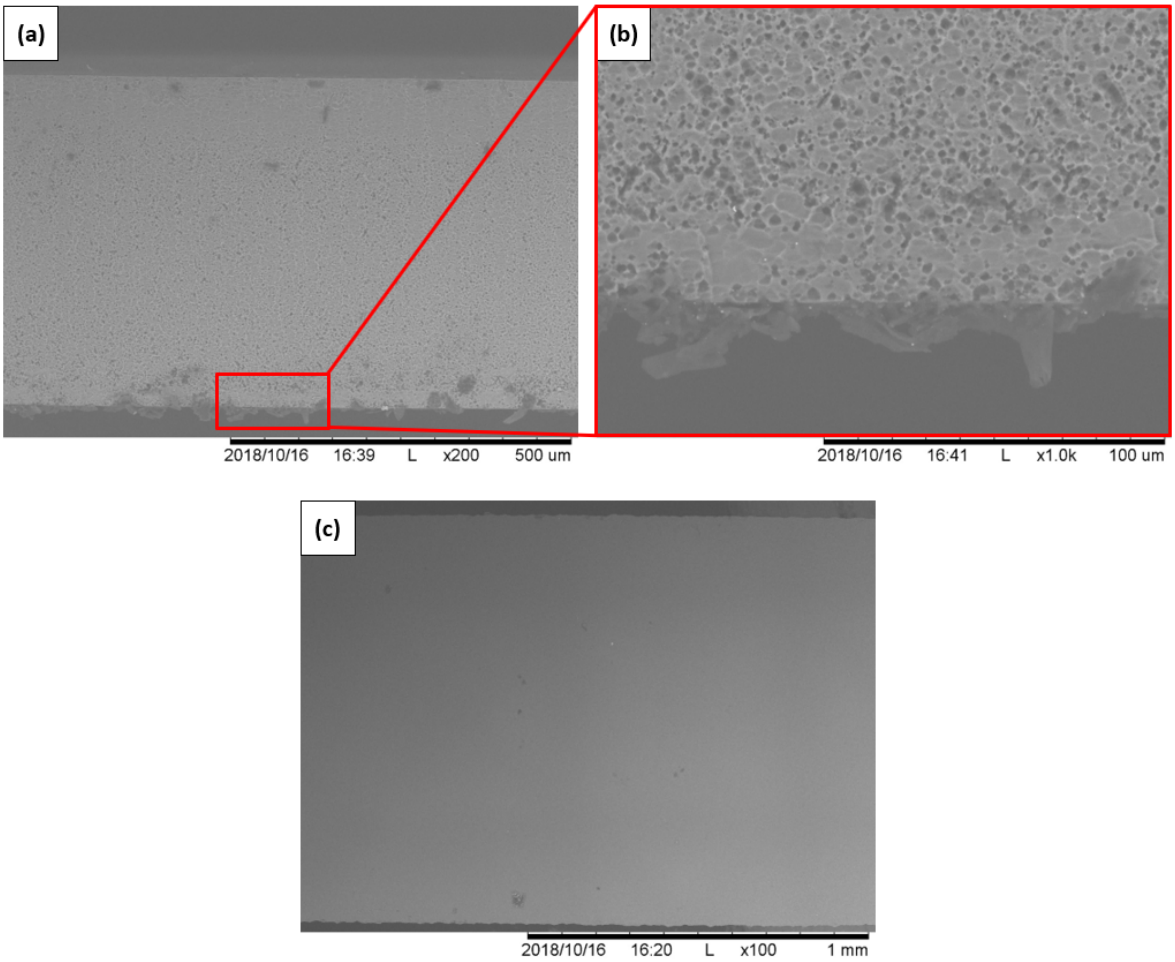

Figure 3.68: (a) SEM imaging of the HNA etched edge of sample 3, magnified 200×. (b) Bottom HNA etched edge of sample 3, magnified 1000×. (c) HNA etched face of sample 3, magnified 100×.

The samples were SEM imaged in order to investigate the effects of this etch recipe on the faces and edges, shown in figure 3.68. As shown in figures 3.68(a) and (b), the HNA etch appears to have removed the linear cutting effects, suspected to be from the pulsed laser, from the edges of the sample, implying some form of surface etching had occurred. However, the edges appears to be pitted throughout their width implying the etch had either caused this directly or has not been given enough time to fully remove the laser cutting damage. One also notes on the

bottom vertex of the edge, a series of features apparent on the face of the sample as viewed into the page in figure 3.68(b), either generated by etching or some form of surface contamination. Imaging of a sample face, as shown in figure 3.68(c), appears to show little to no etching effect. There appears to be contamination near the bottom of the sample face, with a piece of dust or debris visible. Since ITM did not conduct this etching in a cleanroom and ran the sample under ordinary tap water before a DI water rinse, contamination is possible, and a possible explanation for features shown on the faces of the sample.

While this initial experiment suffered from many flaws, due to inexperience of HNA chemical etching, contamination of samples and poor characterisation, these 5 samples were still used for strength testing.

A second etch was conducted in the JWNC from a wafer cut using the old LML laser cutting technique, referred to as "JWNC HNA". All chemical and characterisation work within the JWNC was undertaken by Dr Matthew Smith, with the author observing the work being carried out.

The HNA etch selected is known as the 3:2:1 etch, which consists of a ratio of 3 parts acetic acid, 2 nitric acid and one part HF [316] [317]. The original intention was to etch 5 samples individually for varying periods of time to see etchant effects on surface and material removal amount. Due to time constraints within the JWNC, only one sample was etched.

The recipe used was 150 ml CH_3COOH at 100 % concentration, 100 ml HNO_3 at 69 % concentration and 50 ml HF at 48 % concentration.

Referring to figure 3.65(b), one determines this should develop a geometry within area F. This area and the resultant surface geometry is not covered in the Schwartz and Robbins paper [244]. The reported etch rate of this solution is meant to be 3–5 $\mu m\,min^{-1}$ [316].

This etch is performed at room temperature in a non-glass beaker, due to the presence of HF which etches oxides such as borosilicate and fused silica, and the etching solution was agitated gently with a plastic stirrer throughout the etch. A fizzing reaction was visible on the sample surfaces. The sample was etched for a total of 10 minutes.

The sample was then inspected, photographed, and underwent Nomarski optical microscopy [318] within the JWNC, as shown in figure 3.69. The edges of the silicon ribbon were notably shiny after the HNA etch. They were not smooth enough to image on the surface profiler however. As shown in figures 3.69(b) and (c), one notes bubble-like formations on the surface of the sample edge. It is unknown whether these formed due to the fizzing reaction or not. Figure 3.69(b) shows that this extends across the length of the sample. These formations appeared to be directional along the length of the sample. Most importantly however, all evidence of laser cutting has been removed from the edges.

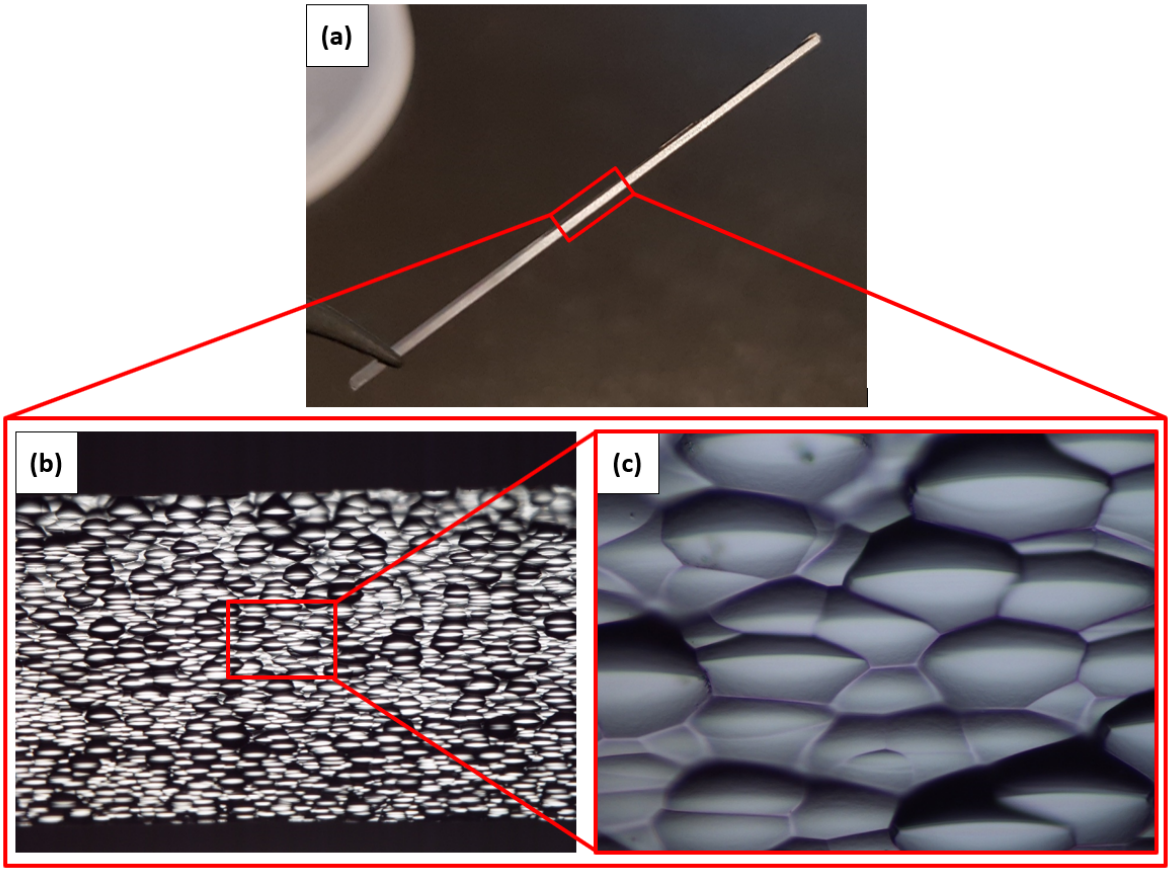

Figure 3.69: (a) Photograph of the edge of sample 1 after HNA etching. (b) Nomarski optical imaging of the edge of the sample, magnified 20×. (c) Area of (b) magnified 100×.

SEM imaging was taken of the edges and faces, as shown in figure 3.70. As shown in figure 3.70(b), these bubble-like features appeared to be depressions in the edge of the samples where the etching process has taken place. Again, comparing to figure 3.70(a), all evidence of laser cutting has been removed. As shown in figures 3.70(c) and (d), no apparent effect was visible on the front or back faces of the sample itself which is unusual considering the isotropy of the etch but consistent with the previous HNA etch conducted at ITM. However, this may be due to the agitation of the fluid and initial smoothness of the polished faces making it difficult for bubbles to adhere to the surface. One however notes the apparent damage and roughness of the vertices where the bubble-like features from the edge intersect with the polished faces. It is possible that bubbles of hydrogen gas have caused this effect on the edges of the sample, especially as equation 3.11 shows that this is a by-product of the reaction. It is also possible this is purely a result of the etching process itself however one can not explain why this effect is not seen on the faces of the sample too.

Due to the COVID-19 pandemic, there was no further time available to pursue this line of work.

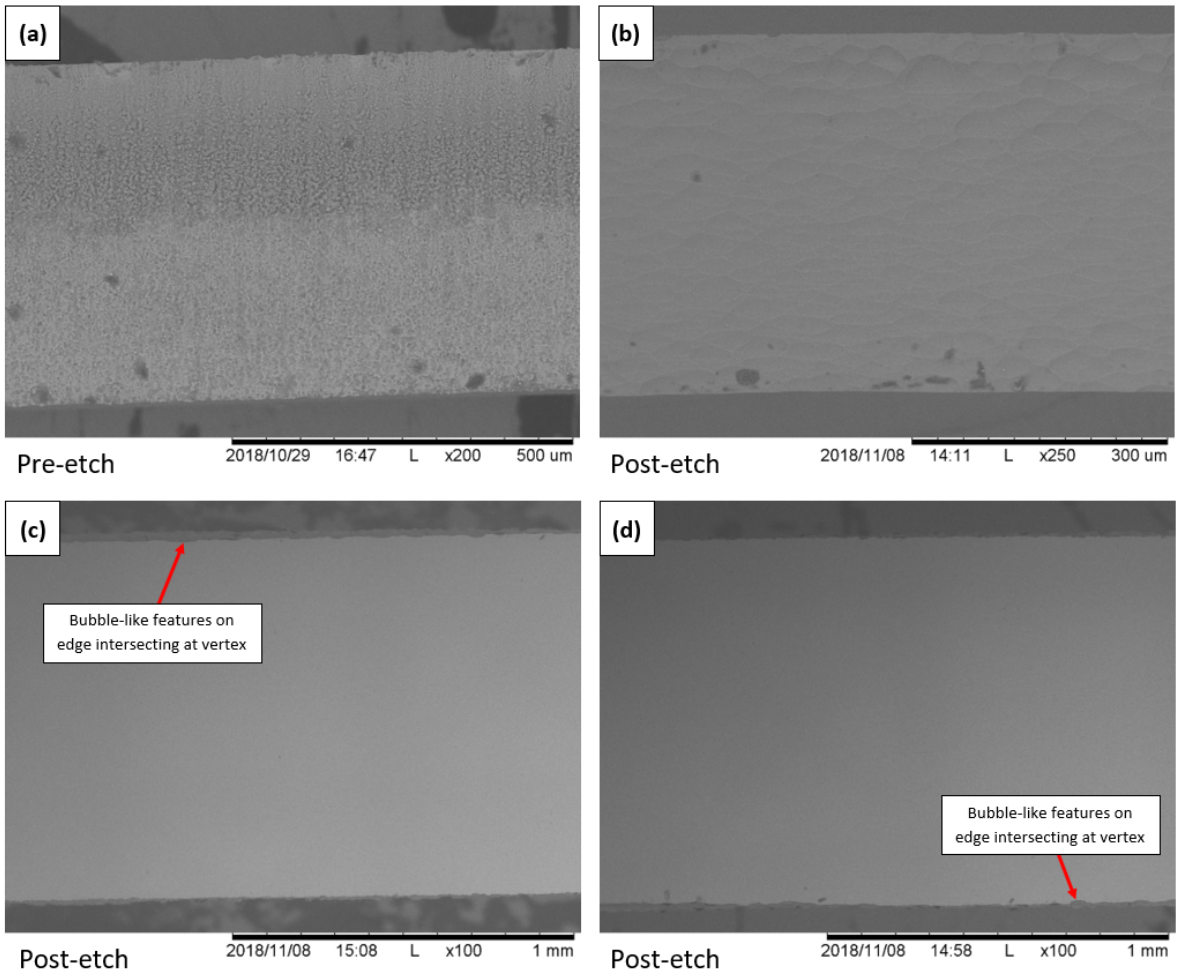

Figure 3.70: (a) SEM imaging of old laser technique cut edge of sample one, before HNA etching, magnified 200×. (b) Old laser technique cut edge of sample one, after HNA etching, magnified 250×. (c) Polished sample face, after HNA etching. (d) Opposite polished sample face, after HNA etching. (c) and (d) are magnified 100×.

3.8.2.2 Initial KOH experiments

This experiment used the identical offcut samples as the HNA etch. This experiment will be referred to as "ITM KOH".

The intention of this first etch was to explore the admixture of IPA into the KOH [307] [312].

The KOH recipe was a ratio of 1 part 20 % KOH by concentration, to 1 part of 4 parts H_2O to 1 part of IPA [309] [319] [320] at a temperature of 80 °C. This recipe is said to produce a smoother surface, with an etch rate of 0.96 µm min^{-1} [320]. Due to this etch rate, the prescribed etch time was 156 minutes to etch the 5 silicon samples to attempt to remove 150 µm in the (100) direction. One would anticipate surface features such as hillocks since the KOH concentration is <45 % [307].

After etching, the samples were wiped with IPA, followed by a rinse in tap water, followed by DI water. The samples were then oven baked at 60 °C for 15 minutes to remove any remaining

liquid.

The samples appeared to be visibly rough with mottling across the entire sample. Measurement indicated that only approximately 27 µm was etched off the width of the sample and 83 µm off the thickness. Although this is an anistropic etch, both the width and thickness edges and faces sit in the (100) plane. This apparent difference could be down to the standard calliper measurement error. Alternatively, the differing etch amounts could be due to starting quality of the surfaces, with the rough laser cut edges of the silicon sample inhibiting etching compared to the polished faces. There was no initial explanation for the etch amounts being well below the 150 µm target.

SEM imaging is shown in figure 3.71. Samples 1 and 3 highlight all of the main features etched onto the samples. All evidence of laser cut damage on the edges appears to have been removed.

As shown in figures 3.71(a), (b) and (c), striations are seen across the etched edge of the samples. In this initial experiment, no agitation was conducted during the etch. Due to this, it is curious to see directional striations on the edges of the etched samples, as it is unlikely these were created by fluid flow over the surface edges, as with the HNA etched samples. There is no apparent explanation for these striations and they can not be from underlying features of laser cutting as they run exactly perpendicular to the laser cutting pattern. Highlighted in figures 3.71(c), (e) and (f), pyramidal structures, or hillocks, are shown on the etched sample. This is as expected due to the low KOH concentration. The edges of these hillocks are approximately 55° into the sample surface due to KOH etching parallel to the (111) crystallographic plane, as shown in figure 3.66. Finally, as highlighted in figures 3.71(b), (d) and (e), the apparent roughness of the samples come from what appears to be an undercutting effect during the etch [313].

It was upon inspection of these samples, it was realised that the KOH etch should be preceded by HF dip to remove the native oxide on the silicon surface. This is because KOH itself etches native oxide at a slower rate than silicon, typically 1 to 2 orders of magnitude [309]. Since native oxide on silicon is typically 1–4 nm thick it has acted as a masking layer for the KOH [182] [321]. Native oxide growth is also typically non-uniform across the silicon surface [321]. This was confirmed to be the case on the samples from ellipsometry measurements. For the length of time of this etch, the native oxide has *eventually* been etched away in the thinnest areas, allowing the KOH to etch the bare silicon underneath at a much faster rate after breaking through. Once this has occurred, undercutting takes place where the KOH etches the bare silicon faster than the native oxide above, creating these features. This is also likely why the etch amount was substantially lower than the target, as much of the etching time will have been used to etch through the native oxide. Furthermore, it was discovered from ellipsometry measurements that approximately 13 nm of thermal oxide has grown near the silicon sample edges, due to exposing the silicon to high temperatures during laser cutting. Therefore it is possible that the silicon edges also have a thicker layer of native oxide which would take longer to etch through which

likely explains the discrepancy.

However, this experiment still produced 5 valid samples for strength testing and provided excellent SEM imaging of typical features that result from KOH etching which was wholly consistent with literature.

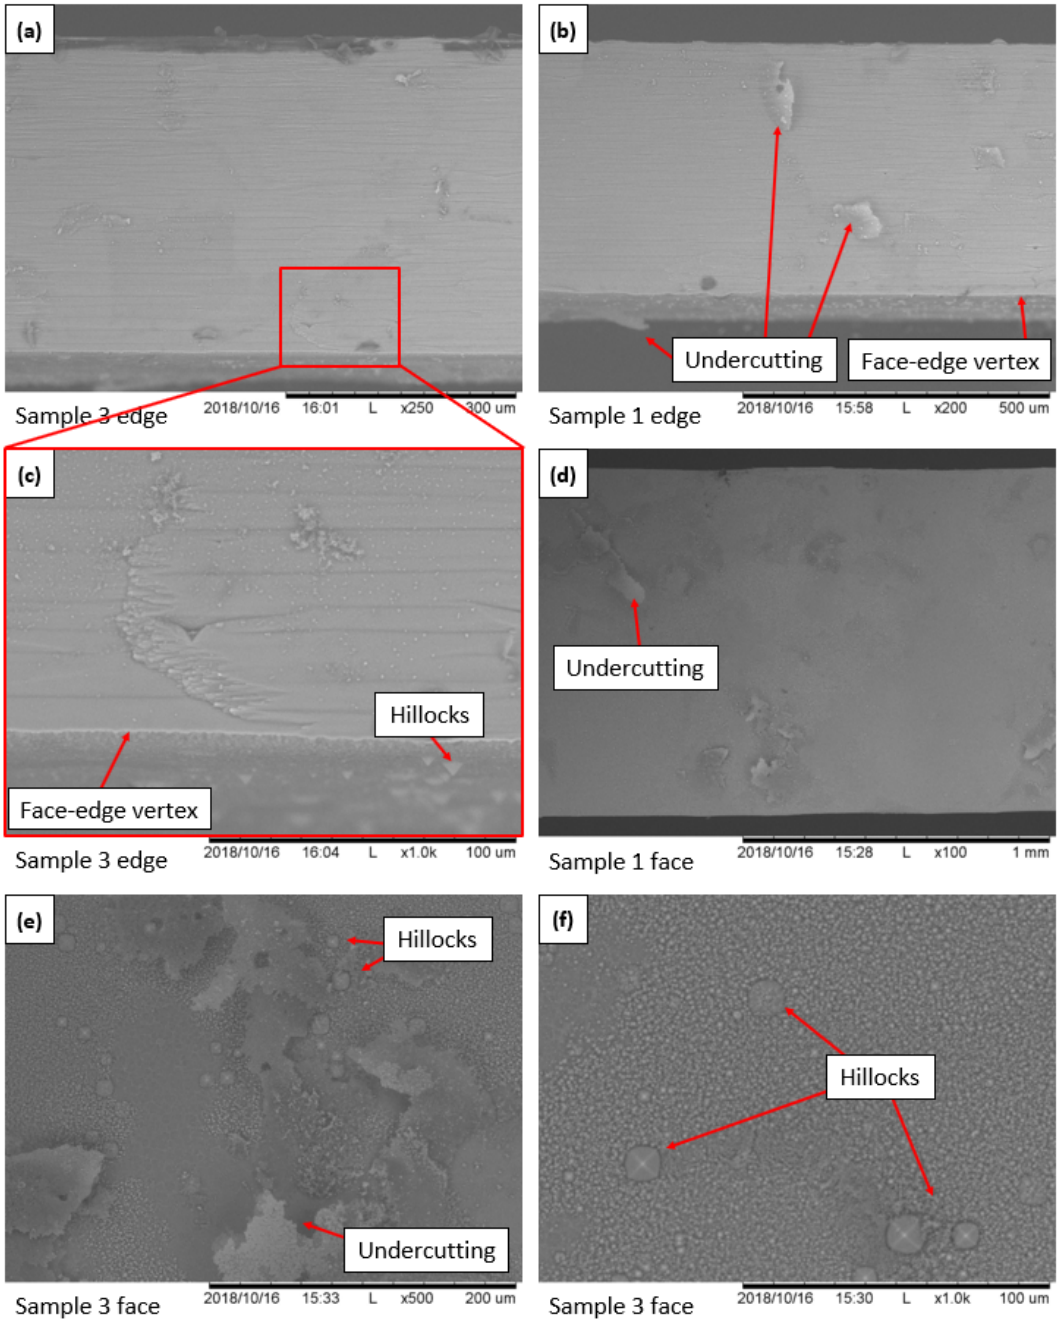

Figure 3.71: (a) SEM imaging of edge of sample 3 after KOH etching, magnified 250×. (b) Edge of sample 1 after KOH etching, magnified 200×. (c) Bottom vertex on the edge of sample 3 after KOH etching, magnified 1000×. (d) Face of sample 1 after KOH etching, magnified 100×. (e) Face of sample 3 after KOH etching, magnified 500×. (f) Face of sample 3 after KOH etching, magnified 1000×.

3.8.3 Strength results and error analysis

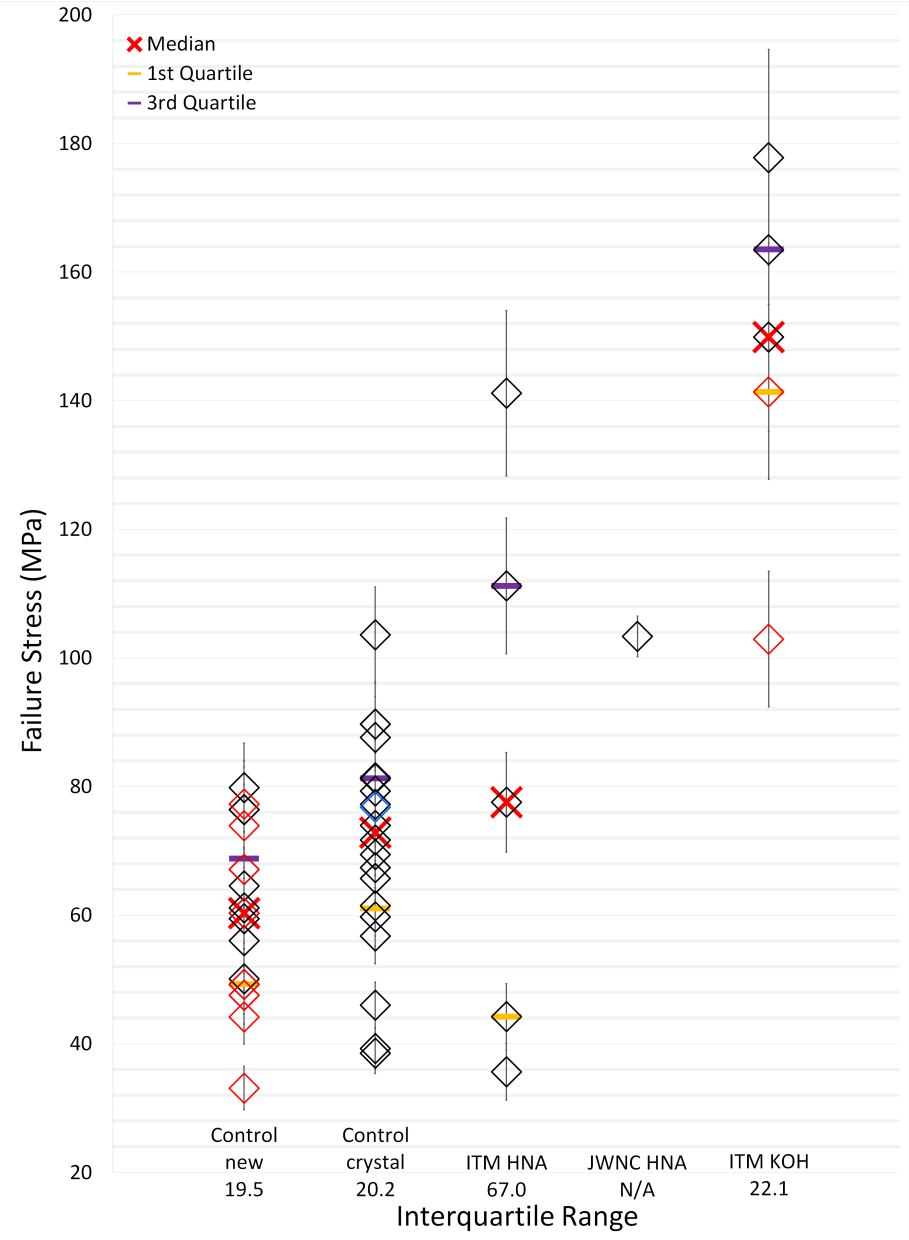

Figure 3.72: Results of tensile strength testing on on wet chemical etching treatments. ◇ indicates sample failed at fuse end interface. ◇ indicates a sample was noticeably misaligned in the fuse end assembly. ◇ indicates the sample pulled free from a fuse end due to epoxy failure so the sample was re-epoxied into a new fuse end and re-tested. Where more than one anomaly exists on a sample, the respective colours are applied to the data marker.

The ITM HNA, JWNC HNA and KOH surface treatments were strength tested, as shown in figure 3.72. All chemical treatments improved the median strength over the *control new* and *control crystal* samples.

Due to only 1 sample being part of the JWNC HNA sample set, median, quartile and IQR values can not be obtained.

3.8.3.1 HNA surface treatment removal amount and errors

Sample	Average width (µm)	Error (µm)	Average thickness (µm)	Error (µm)
1	1230.0	91.9	505.0	15.6
2	1240.0	91.9	510.0	15.6
3	1240.0	91.9	510.0	15.6
4	1240.0	91.9	510.0	15.6
5	1235.0	91.9	500.0	15.6

Table 3.17: Results of the average width and thickness of 5 samples after ITM HNA etching as measured by digital callipers.

For the ITM HNA experiment, table 3.17 shows the width and thickness of the 5 HNA etched samples before strength testing. The values are averaged from two measurements at either end of the silicon ribbon and the combined calliper errors for width, 130 µm, and thickness 22 µm, are applied respectively.

	Width (µm)	Thickness (µm)
Pre-etch	1257.4	516.9
Post-etch	1216.5	499.3
Etch amount	40.9	17.6
Error	4.4	7.2
Etch rate ($\mu m\,min^{-1}$)	4.1	1.8
Error ($\mu m\,min^{-1}$)	0.4	0.7

Table 3.18: Results of the width and thickness of sample 1 before and after JWNC HNA etching as measured on the LIGO fibre profiler. Also included are the etch amounts and etch rate for each dimension.

For the JWNC HNA experiment, table 3.18 shows the width and thickness of the sample before and after HNA etching, including the etch amount, etch rate and associated errors.

One notes that where clear etching has occurred on the edges, corresponding to a reduction in the width of the sample, the etch rate appears to be consistent with the expected etch rate of 3–5 $\mu m\,min^{-1}$ [316]. However the thickness etch rate is notably slower, consistent with imaging of the sample's faces which showed little evidence of change. Literature notes this anisotropy can arise in HNA etching, particularly with high concentrations of nitric acid relating to diffusion limiting reactions [304] [322].

3.8.3.2 KOH surface treatment removal amount and errors

Sample	Average width (µm)	Error (µm)	Average thickness (µm)	Error (µm)
1	1205.0	91.9	435.0	15.6
2	1190.0	91.9	440.0	15.6
3	1205.0	91.9	435.0	15.6
4	1200.0	91.9	445.0	15.6
5	1205.0	91.9	455.0	15.6

Table 3.19: Results of the average width and thickness of 5 samples after ITM KOH etching as measured by digital callipers.

For the ITM KOH experiment, table 3.19 shows the width and thickness of the 5 KOH etched samples before strength testing. These samples are measured in the same way as the 5 HNA samples.

3.8.4 Discussion

Comparing the ITM HNA and ITM KOH etches, one notes the KOH etch median stress value is 93 % higher than the HNA etch.

The strongest sample from the ITM HNA etch, yielded a 1.8× increase in tensile strength compared to the strongest sample from the *control new* set, with the strongest sample from the ITM KOH etch yielding a 2.2× increase. Comparing to the *control crystal* set, this reduces to 1.4× and 1.7× respectively.

The median values from both sets also increased compared to control samples, although this increase was marginal for the ITM HNA etch. The median strength from the ITM HNA etch was 1.3× higher than the *control new* set and 1.1× higher than the *control crystal* set. The median strength from the ITM KOH etch was 2.5× higher than the *control new* set and 2.1× higher than the *control crystal* set. One notes that the weakest 2 samples broke at the fuse end interface. While there is the potential for a non-tensile mechanism to cause this failure prematurely, there are not enough data points to draw any conclusion in this sample set. Furthermore, no particular correlation was seen in sections 3.6 or 3.7.

The single JWNC HNA etched sample appears to be consistent with the ITM HNA treated samples due to the fact it sits within the extremes of the ITM HNA data.

One might reasonably assume chemical etching is less likely to treat samples asymmetrically, for example due to different pressures or misalignments of samples due to the fact that wet etching surrounds the samples evenly and the chemical reaction should be the same throughout the sample geometry. However, due to the small sample sizes it is difficult to draw comparisons from available IQR data.

The critical crack sizes for the ITM HNA etch stress range 36–141 MPa equate to 127.5–8.2 µm. Looking at figure 3.68(a) and 3.68(b), one notes chipping or etched features on the

order of tens of μm which appears to be consistent with the smaller critical crack sizes.

The critical crack sizes for the single JWNC HNA etched sample was 15.2 μm. Looking at figures 3.69 and 3.70, it is unlikely that the bubble-like structures were the cause of failure due to the fact they are typically on the order of 50–60 micron along their length and the direction of tensile testing. Some damage exists along the vertices where this bubble-like pattern meets the polished faces of the sample, so it is entirely possible the failure mechanism came from a feature that arose from this interface.

The critical crack sizes for the ITM KOH etch stress range 103–178 MPa equate to 15.3–5.1 μm. Firstly, although large features exist from undercutting, due to their geometric nature it is unlikely these contribute to the bulk strength of the silicon sample. Hillocks are on the order of tens of μm and, combined with their sharp features acting as stress concentrators, it is possible that a hillock on a vertex of the sample could create a geometry that would lead to the associated critical crack sizes.

Fundamentally, more data points and better characterisation is required to better compare sample sets to each other as well as provide a more conclusive picture of where chemically etched samples sit against other surface treatment regimes. However, the initial aim of this work was to determine the practicalities of different etch recipes and the surface quality left behind as opposed to the etch removal amount or etch rate. Future work will require good characterisation and larger sample sets for the purposes of statistical analysis.

3.8.5 Future chemical etching experiments

Based on initial experiments it was decided to pursue KOH etching. The aim was to both correlate surface material removal to strength, as well as investigate effects of etching on crystal and random orientated samples. Preliminary work was conducted to explore these aims but was cut short due to the COVID-19 pandemic. Details on the work carried out and future chemical etching work proposals are detailed in appendix A.

3.9 Conclusion

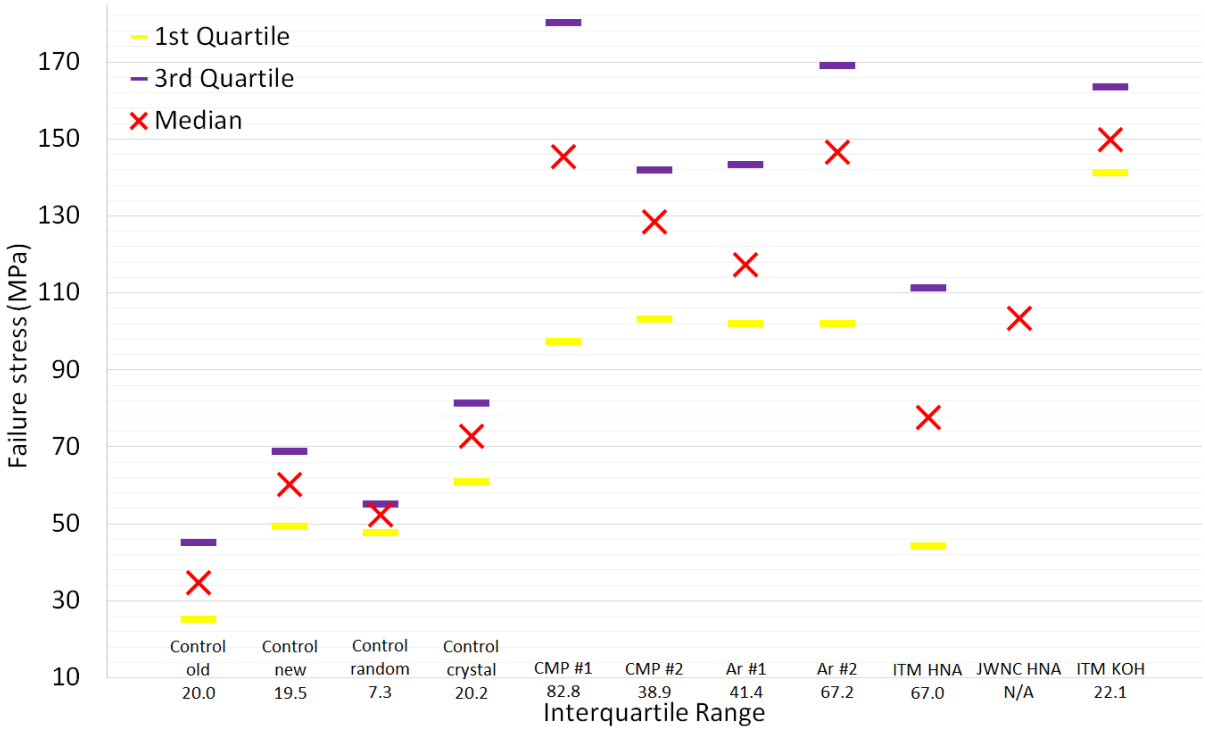

Figure 3.73: Median, 1st quartile, 3rd quartile and IQR values for controls and respective surface treatments. Discounted data from experimental error has been removed and justified in the respective results discussion section for each surface treatment.

This final section looks at the median, 1st quartile, 3rd quartile and IQR values together for all surface treatments compared to the control samples in order to draw comparison between each. The results are shown in figure 3.73.

Figure 3.73 clearly shows all surface treatments improved the median strength values compared to the 4 control sets.

The results of the control samples clearly show that laser cutting of silicon samples induces substantial surface damage, even with improved cutting techniques. The geometries available from laser cutting are also limited and restricted to wafers of certain thickness.

While CMP treatments are promising in terms of overall strength, one notes the large IQR in CMP treatment #1. This is likely due to the asymmetrical lapping and polishing on one side of the samples and appears to be improved in CMP treatment #2, with the 8 misaligned samples removed from the data. It is notable the weakest samples of both subsets are at or around the highest median stress values for the control sample sets.

This surface treatment will rely on substantial physical contact with suspension fibres or ribbons which may not be practical or scalable for circular fibre geometries.

Argon ion etching, particularly argon ion etch #2, achieves promising median strength results with reasonable IQR by comparison. It is notable the weakest samples of both subsets are at or

around the highest median stress values for the control sample sets.

This surface treatment has the advantages of being non-contact, with the potential for scalability, the treatment of circular fibre geometries and shaping of suspension element geometries. However, the surface finish may result in undesirable surface loss, although this will have to be investigated. Furthermore, inherent stress in the treated fibre or ribbon, from heating and subsequent cooling, may cause issues.

Chemical etching shows promising results for silicon surface and strength enhancement. The most promising result is the ITM KOH etch. This is due to a combination of having the highest median stress, but equally important, having the smallest IQR of any treated sample set. The small IQR means fewer strength outliers, increasing confidence in the ability to run 3G detector suspensions at higher stress or lower safety factors. While this etching left the roughest surface finish of all of the treatments, the weakest sample exceeded the median of 3 surface treated sample sets. As with all of the chemical etching data, it should be noted this is from a smaller sample set and more data is required to draw firmer conclusions.

As shown in section 3.8.1.1, specular, highly smooth surfaces are possible with chemical etching so these poor results should not exclude the use of this surface treatment if surface finish is a concern.

The ITM HNA etch is the least promising result due to the low median stress and larger IQR arising from a large spread in results. It is difficult to comment on the one data point from the JWNC HNA etch, however due to the small associated error, it is notable that the sample matches the strongest samples of all the control sets.

Fundamentally, chemical etching surface treatment are the most practical and scalable treatment for silicon ribbons or fibres. Wet etching allows all surfaces of a ribbon or fibre to be accessed and treated at once, evenly across the sample length and regardless of geometry provided etchant can sufficiently access and flow. Also, chemical baths can be and are built to dimensions scaling to 10s of metres [323].

From the work conducted in this chapter, it is recommended chemical etching be pursued as a surface treatment method of silicon ribbons or fibres. Section 3.8.5 details on how to proceed specifically with KOH etching. It is also recommended that further investigations be made into practicality and scalability of argon ion etching, or similar ion source experiments to find ways of treating suspension-scale ribbons or fibres.

Finally, all strength testing conducted in this chapter was carried out at room temperature. Future work should investigate strength testing within cryogenic regimes to investigate any potential effects temperature has on ultimate tensile strength.

This work has demonstrated an increase in ultimate tensile strength in silicon ribbons. These higher stresses will enable the use of small cross-sectional ribbons or fibres that are less stiff which will ultimately decrease detector noise through a decrease in the thermoelastic noise and increase in the dilution. This will also brings violin frequencies further up, or the vertical bounce

frequency down, out of the detection band frequencies. Higher stresses within suspension elements will also allow for larger engineering and alignment tolerances which may otherwise be restrictive in untreated silicon.

Chapter 4

Cryogenic silicon suspension prototype

4.1 Introduction

To the best of the author's knowledge, no silicon suspension has been successfully built and cryogenically cooled. In order to realise 3G cryogenic silicon detectors, prototype detectors must first be designed and built in order to withstand room temperature, non-vacuum environments as well as be robust enough to handle vibrational and mechanical loads that may impinge on the suspension system during assembly and installation.

As shown in chapter 3, silicon is a brittle material, making it difficult to work with for mechanical handling and assembly. While most materials will also exhibit brittle properties at cryogenic temperatures, the detector will have to be assembled at room temperature, where silicon still retains its brittleness [248] [249]. As such, it is necessary to design and build prototype suspensions from silicon in order to investigate known, and discover unknown, challenges for using silicon as a mechanical material at macroscopic scales.

Furthermore, the silicon suspension prototype will then need to be cryogenically cooled. The silicon suspension will be subject to further mechanical vibration from the cryogenic cooling system and will experience thermo-mechanical stresses, arising from silicon's unique coefficient of thermal expansion properties, when undergoing cooling from room temperature down to cryogenic temperatures.

This chapter covers these challenges by discussing the iterative design and build processes for assembling the first single ribbon silicon suspension loaded with a 1 kg mass. Assembly of the suspension and results from cryogenic cooling are covered after mechanical designs and assembly of the suspension support structure. Before this, an overview of the cryogen-free pulsed tube cryostat is given along with interface design between the suspension support structure and the cooling system.

4.2 Cryostat operation and interfacing

Significant cryogenic infrastructure will be necessary to cool large silicon test masses. There are a number of practical solutions available for cryogenically cooling gravitational wave detectors, each with advantages and disadvantages. Fundamental cryogenic considerations must be; the ability to cool the test mass to the required temperature, the time taken to cool to the required temperature and minimisation of any mechanical or thermal noise impinged onto the detector via cooling.

This has led to a variety of cryogenic method proposals which can involve radiative cooling, conductive cooling or a combination of both. Some proposals utilise cryogenic shields that are cooled via heat-links attached to standard liquid nitrogen and/or liquid helium reservoirs [198] [324] [325]. These applications are only viable for the high temperature 123 K operating regime as the test mass is predominantly cooled radiatively by the cryogenic shields but has the advantages of simpler mechanical design with little or no system vibration injection and does not require conductive primary heat extraction through the silicon suspension elements [122]. However, compared to conductive cooling, this method of cooling is inefficient, increasing the detector downtime during cooling [198] [326]. This downtime may be reduced by the use of helium exchange gas which can then be pumped out before laser injection [126]. One also has to consider the Newtonian noise arising from the heat shields in close proximity to the test mass [325].

For the low temperature 18 K operating regime, conduction is the only method available for cooling 3G detector temperatures arising from the T^4 relationship for radiative cooling which becomes much less effective at low temperatures [200]. All methods for low temperature regime 3G detector designs rely on direct conduction through suspension elements to reach the ultimate operating temperature [119] [127]. This can however still be complimented by radiative shielding to minimise heat load from surroundings onto the conductively cooled test mass [119].

As the only operational detector-scale cryogenic suspension, KAGRA offers excellent examples of low temperature cooling systems for cryogenic suspensions. Conductive cooling of the test masses is achieved using low-noise pulse tube refrigerators (PTRs) connected to heat links that are attached to the marionette, as shown in figure 1.16(b) of chapter 1 [327]. While PTRs are already low-noise it was found that too much mechanical vibration still coupled into the detector through existing heat links [222]. This problem was addressed by using soft, highly thermally conductive, pure aluminium heat links in order to minimise transmitted vibration from the PTR and transfer heat efficiently [328] combined with a heat link vibration isolation system (HLVIS) [329].

The work in this chapter aims to demonstrate conductive cooling to the lowest possible temperature available from the experimental equipment rather than specifically aim for radiative or conductive cooling regimes for the primary purpose of prototype demonstration.

4.2.1 Leiden Cryogenics cryostat

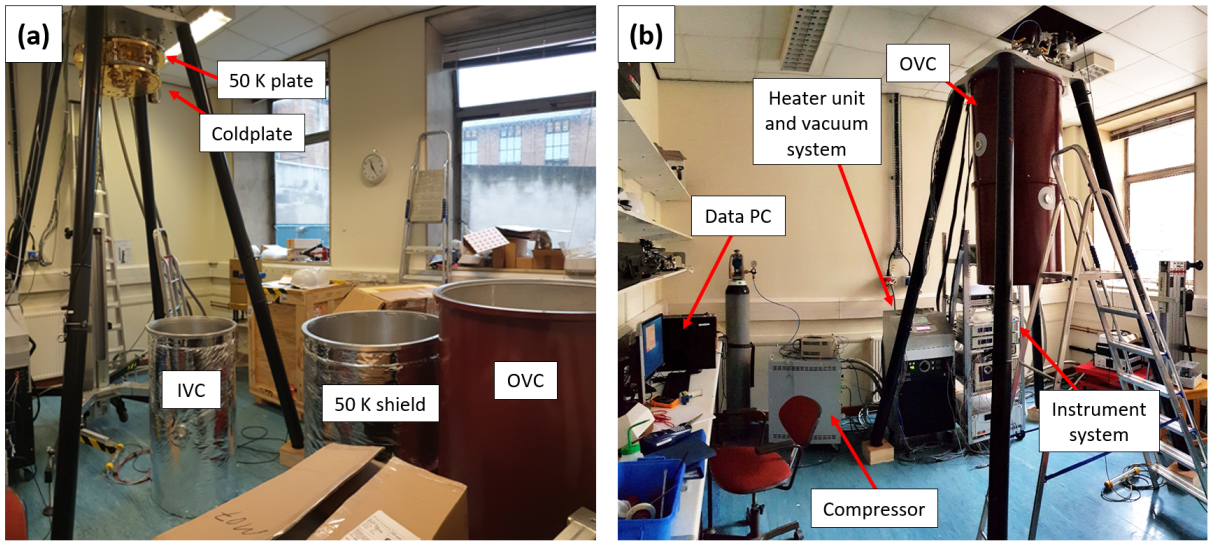

Figure 4.1: (a) Leiden Cryogenics cryostat with IVC, 50 K shield and OVC removed. (b) Leiden Cryogenics cryostat closed up and operational.

As shown in figure 4.1, a custom built cryogenic system (CF-CS81-3K-Glasgow) was designed by Leiden Cryogenics for the purposes of cryogenic suspension experiments. The cryostat has the advantage of a large experimental volume with $\approx 0.2\,\mathrm{m}^3$ of experimental space contained within the inner vacuum chamber (IVC), shown in figure 4.1(a), which can cool the coldplate down to 4 K [330] [331].

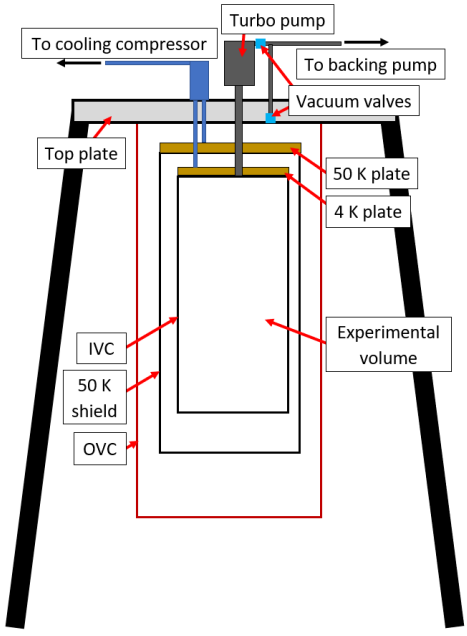

Figure 4.2: Schematic of Leiden cryostat vacuum and cooling system.

4.2.1.1 Vacuum system

As shown in figure 4.2, the IVC is bolted directly onto the rim of the cold plate, with a Kapton o-ring seal, and is evacuated using a scroll pump and a turbo pump, reaching pressures of $\approx 10^{-5}$mbar at room temperature. A 50 K shield is then bolted onto the 50 K plate. This is not a vacuum chamber and is purely used as a radiation shield to minimise external heat load on the IVC. Finally, the outer vacuum chamber (OVC) is bolted onto the top plate of the cryostat structure. This acts as an outer vacuum chamber, sealed with a rubber o-ring, and is evacuated using the same scroll pump as the OVC but is not evacuated by the turbo pump, reaching pressures of $\approx 10^{-3}$mbar at room temperature. The open IVC, 50 K shield and the OVC are shown in figure 4.1(a). The IVC and OVC have Pirani vacuum gauges fitted that read down to $\approx 5 \times 10^{-4}$mbar [332]. The IVC also has an Edwards wide range gauge fitted, capable of reading down to $\approx 10^{-9}$mbar [333].

4.2.1.2 Cooling system

This system benefits from a similar PTR mechanism to KAGRA which cools the coldplate to 4 K using a closed-loop compressor charged with helium gas. The PTR operates by compression and expansion of helium gas with few moving parts throughout and no moving parts in the cold head, which significantly reduces vibrations imparted onto the cryostat [327]. The helium gas is compressed in a compressor unit, shown in figure 4.1(b). This gas is then transported up to the top of the cryostat to the PTR head via a high pressure flexible metal line. This gas then expands adiabatically, in turn cooling, via a two-stage heat exchanger mechanism, the 50 K plate and coldplate to 50 K and 4 K respectively. This expanded helium gas is then transported back via a low pressure flexible metal line to the compressor to be re-compressed and the process is repeated in a closed-loop cycle continuously. This high and low pressure helium gas is cycled at the top of the cold head through a rotating remote motor mechanism. This system provides 1.5 W of cooling at 4 K [330].

Two 100 W heaters are attached on the bottom of the 50 K plate and the top of the 4 K plate. These are designed to bring the temperature of the cryostat back to room temperature quickly by applying current across the resistors using a simple variac on the heater unit, shown in figure 4.1(b). These 100 W heaters are not designed for temperatures stabilisation of cryogenic experiments.

4.2.2 Cryostat instrumentation, electronics and sensing

As shown in figure 4.3(a), the 4 K plate of the cryostat comes with an array of M3 to M10 tapped threads ideal for affixing sensors, experimental apparatus or mechanical structures. This plate also comes with a series of feedthrough channels for cold-insertable probes or other in-vacuum feedthrough requirements, such as electronics.

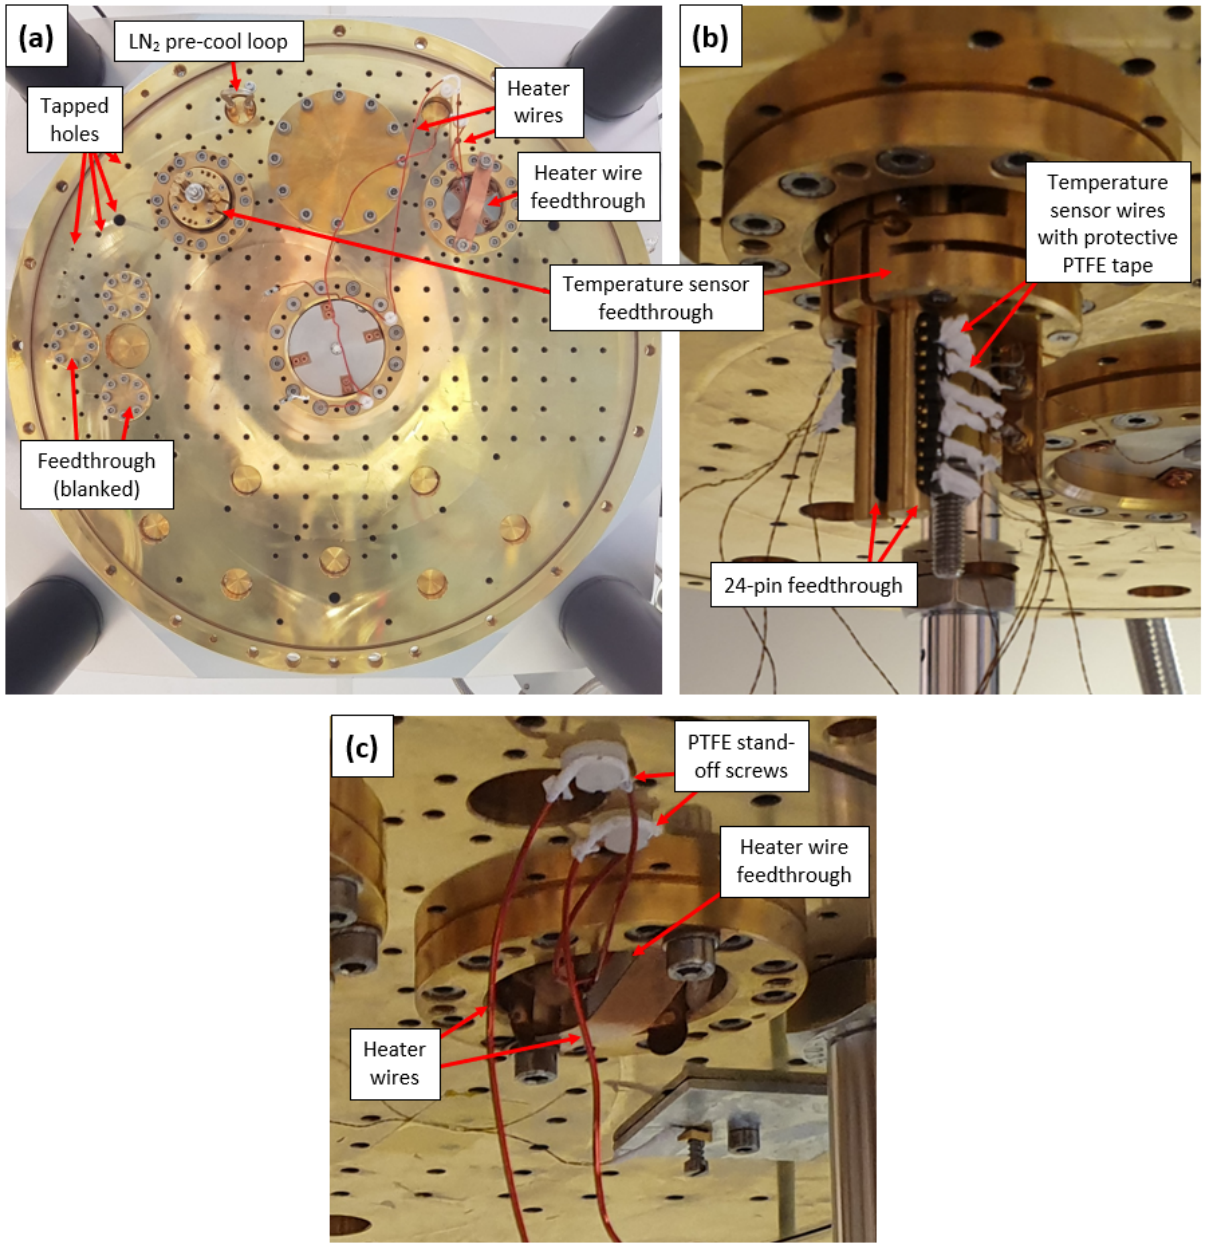

Figure 4.3: (a) 4 K coldplate layout. (b) EVF feedthrough. (c) Heater wires and feedthrough.

For this cryostat, most of these feedthrough sections are blanked off or unused. Only two of the feedthroughs were required for use. One feedthrough supplies two 24-pin electronic vacuum feedthrough (EVF) posts, shown in figure 4.3(b). These pins are reserved for low current and low voltage cryogenic temperature sensors.

The second feedthrough was used for two stiff copper insulated wires. These wires pass from the room through a vacuum feedthrough plate attached to the external vacuum system on the top plate, before passing through a series of thermal baffles within the feedthrough channel. They are then mechanically anchored to the 4 K plate. These wires were required for providing larger voltages or currents than could be handled by the two 24-pin EVF posts, primarily for the purposes of heating any experiment within the IVC, as such these wires are referred to as the

"heater wires", as shown in figures 4.3(a) and (c).

Figures 4.3(b) and (c) highlight three key considerations that must be balanced when wiring cryogenic experiments. Mechanically, anchoring of wiring is required, primarily to secure and create wire routing within the cryostat. Sufficient slack must be left to account for the small linear thermal contraction of the wires that will occur due to cryogenic cooling in order to avoid wire breakage. Electrically, sound connections must be made, taking all precautions against electrical shorting. While a wire may be electrically sound at room temperature, the aforementioned mechanical contraction due to cooling may cause the wires to shift position leading to a shorting scenario. Furthermore, mechanical abrading due to movement of the wires may strip the thin insulation layers. Mechanical anchoring helps avoid this but further protective insulation in the form of PTFE tape and stand-offs or thicker annealed wire coatings, for example, can help to mitigate this further. Thermally, all wires connected to a cryogenic experiment should be designed in such a way as to minimise any heat load from room temperature being transmitted along the wire into the experiment. Primarily, this is complimented by mechanical anchoring to the coldest surfaces, such as the coldplate but can conflict with electrical insulation requirements and so a balance must be struck between risk of heat load against electrical shorting onto the cold structure. Figure 4.3(c) shows a scenario where the risk of electrical shorting onto the coldplate was deemed too high compared to the transmission of heat load into the cryogenic experiment. However, as will be discussed later, this thermal anchoring issue was addressed by an adapted connection to specific cryogenic wires that were then thermally anchored elsewhere before being contacted to the experiment. Many cryogenic wire materials are specifically designed with a low coefficient of thermal expansion, cryogenic-compatible insulation and very low thermal conductivity in order to tackle these issues. Furthermore, low electrical resistivity is desirable to minimise resistive heating of the wires [334].

With the exception of the copper insulated heater wires, all wiring used within the cryostat IVC was Lake Shore Quad-Twist WQT-36 phosphor bronze cryogenic wiring [335]. This wiring highlights two further considerations for cryogenic wiring. Firstly, twisting wires is a common method for minimising pickup of electromagnetic noise [335]. Secondly, the wire consists of 4 smaller wires to enable 4-wire sensing.

4-wire sensing is required for accurate temperature sensor measurement and is superior to 2-wire sensing [334]. This is because measurement of resistance or diode temperature sensors requires passing a current through the sensor to produce a sensor voltage that can then be measured. Referring back to thermal design considerations, one ideally requires wiring of small cross-section and high resistance per unit length in order to minimise heat flow down the wires into the cryogenic experiment. This high resistance per unit length is the dominant source of error in 2-wire sensing, creating a systematic error on the voltage, and hence temperature reading. 4-wire sensing overcomes this issue with 4 leads, with current leads, I+ and I-, supplying current to the sensor and voltage leads, V+ and V-, eliminating the effect of lead resistance by

measuring the voltage at the sensor. This works since the voltage drop across the current leads is not measured and the voltage drop in the voltage leads requires only a very small current by the voltmeter (pA or less) to make the measurement. A schematic of 2-wire and 4-wire sensing is shown in figure 4.4 [336]. Further protection is offered by using properly shielded cables for lead wiring coming from the top of the cryostat to the measuring instrument. Belden 4-core shielded cabling was used for all instrument lines [337]. Each of the cores can carry the I+, I-, V+ and V- lines and are shielded by shielding tape and a drain wire along their length.

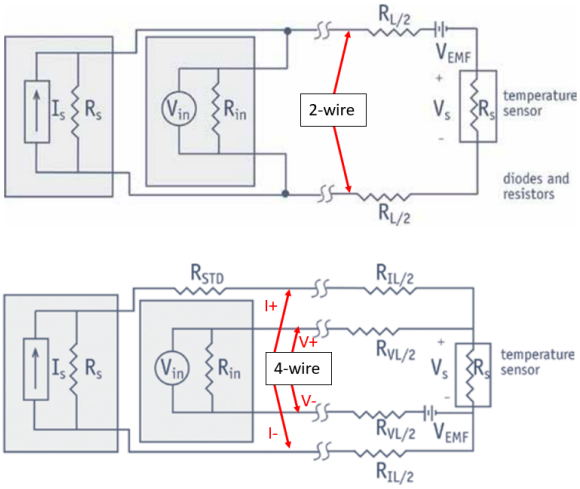

Figure 4.4: 2-wire resistance measurement and 4-wire resistance measurement. Images modified from [336].

To further enhance sensing capability it is necessary to correctly ground sensors for measurement stability and repeatability. Incorrect grounding can induce non-constant current flow via ground loops which can result in fluctuating voltages and hence erroneous or drifting temperature measurements [336]. Furthermore, care must be taken to avoid multiple grounding points to both minimise ground loops and ensure no negation of any enhanced protection measures offered by temperature instrumentation. An example of this, is the shielding potential supplied by a Lake Shore instrument controller. Primary temperature measurements within the cryostat are made on two Lake Shore 340 temperature controllers [338]. A Lake Shore 336 temperature controller was later added for auxiliary temperature measurements with the controllers contained within the instrument rack shown in figure 4.1(b) [339]. Both units offer shielding potentials for temperature sensor lead wiring. Sensors are isolated from earth ground to reduce the amount of earth ground noise in the measurement leads. This isolation is defeated by grounding any sensor leads to any earth ground, effectively bypassing instrument noise reduction measures. A shield is most effective when it is near the measurement potential, so the Lake Shore Model 340 and 336 offer a driven shield that stays close to the measurement potential [338] [339]. The drain wire, which is electrically connected to the shielding tape, is wired to the shield pin on the back of the instrument which then drives the potential of the shield to closely match the measurement

potential. The shielding tape and drain wire should *not* be connected at the opposite end of the cable where it enters the cryostat as other earth ground sources, such as that of the turbo or scroll pump, will defeat the driven shield and introduce ground loop fluctuations into the measurement system. A schematic of this set up is shown in figure 4.5(a).

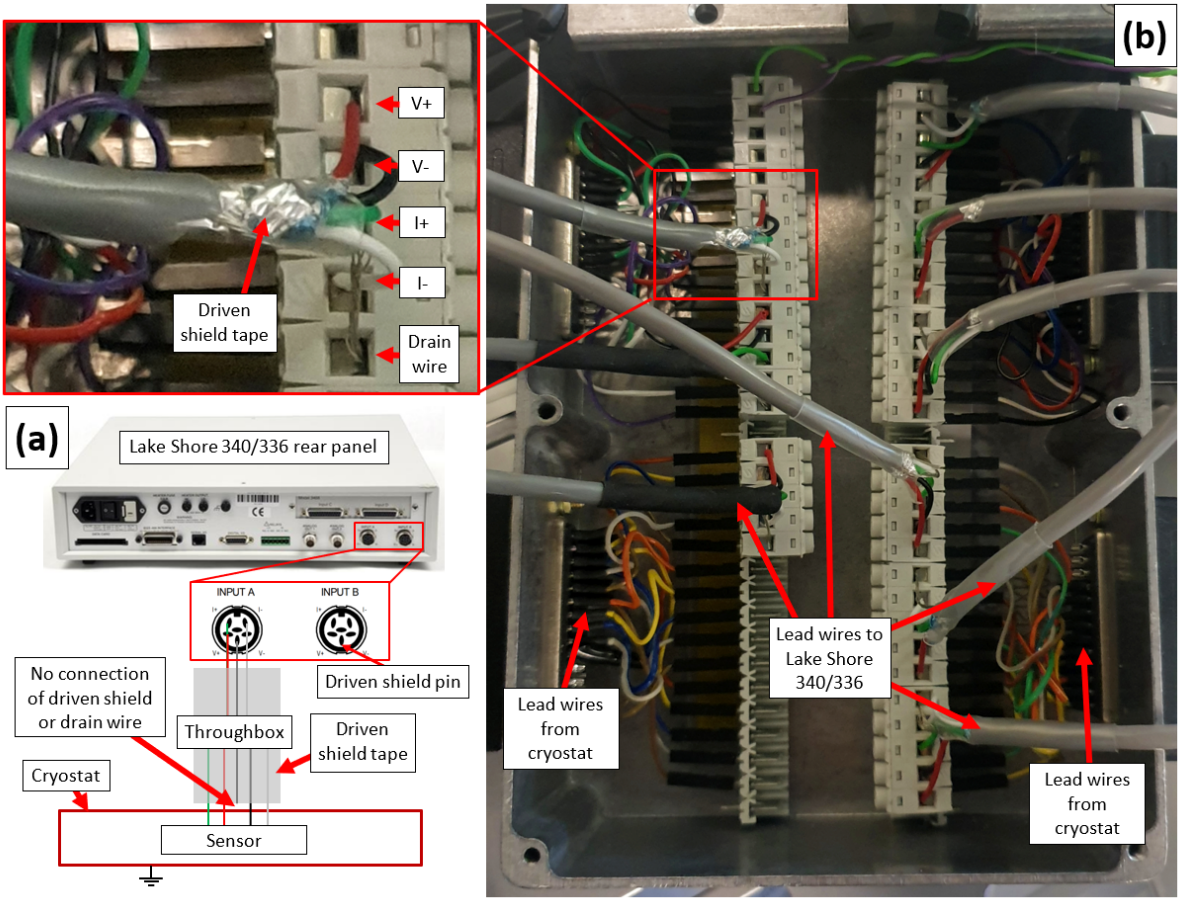

Figure 4.5: (a) Schematic of sensing wire set up between Lake Shore temperature controllers and sensors inside cryostat. (b) Custom built throughbox for temperature sensors. Images in (a) obtained from [338] [340].

To make changing of sensors and temperature channels easier an EVF throughbox was designed and built, as shown in figure 4.5(b). This allowed "plug and play" capability for temperature sensors fitted inside the cryostat. The throughbox was designed to permit 4-wire shielded sensing with a driven shield, as shown in the inset of figure 4.5(b).

4.2.2.1 Temperature sensors

There were two types of temperature sensors used inside the cryostat, both of which utilised 4-wire sensing. Both sensors types were used due to their immediate availability, compact size, accuracy and ease for mounting and dismounting in various locations throughout the cryostat and suspension structure.

The first type of sensors were Lake Shore DT-670 SD silicon diode sensor packages [341], shown in figure 4.6. This sensor was sometimes used in combination with the CO mounting screw also shown in figure 4.6. These sensors provide accurate temperature sensing over temperature ranges of interest. While these sensors come with only two lead wires, 4-wire measurement can still be implemented. All DT-670 sensors used were wired in the 4-wire measurement configuration. These sensors come in either generic uncalibrated form or calibrated form, provided with sensor-specific calibration curve data. Both calibrated and uncalibrated DT-670 sensors were used in these cryogenic suspension experiments.

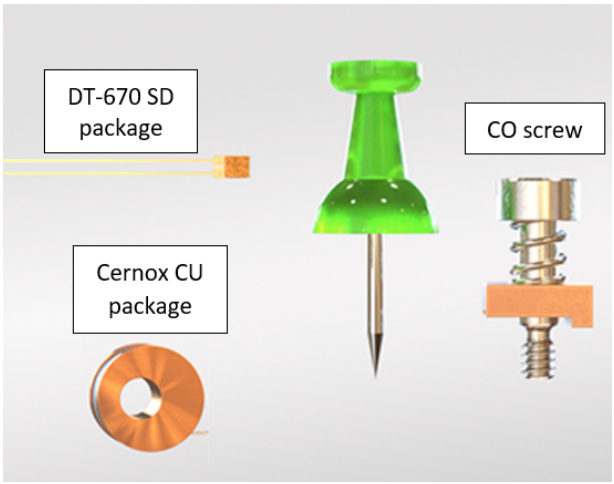

Figure 4.6: Lake Shore DT-670 SD package silicon diode temperature sensor and Cernox CU package temperature sensor. The CO screw package is sometimes used in conjunction with the DT-670 SD package. Green thumb tack for scale. Image modified from [342].

The second type of sensors used were Cernox-1050-HT CU thin film resistor sensor packages, shown in figure 4.6. These sensors are more accurate than the DT-670 sensors over the full temperature range of interest although due to their bulkier size can be more difficult to mount. These sensors come with 4 lead wires, designed to easily enable 4-wire sensing. Each Cernox sensor is individually calibrated and is provided with its own calibration curve.

(a)

	Typical sensor accuracy[2]
1.4 K	±12 mK
4.2 K	±12 mK
10 K	±12 mK
77 K	±22 mK
300 K	±32 mK
500 K	±50 mK

(b)

	Typical sensor accuracy[5]	Long-term stability[6]
1.4 K	±5 mK	±3 mK
4.2 K	±5 mK	±3 mK
10 K	±6 mK	±6 mK
20 K	±9 mK	±12 mK
30 K	±10 mK	±18 mK
50 K	±13 mK	±30 mK
77 K	±16 mK	±46 mK
300 K	±40 mK	±180 mK
400 K	±65 mK	—

Figure 4.7: (a) Lake Shore DT-670 typical calibrated sensor accuracy [343]. (b) Lake Shore Cernox typical calibrated sensor accuracy [344].

Typical calibrated sensor accuracies are shown for DT-670 and Cernox sensors in figure 4.7(a) and (b) respectively. These values are typically taken as the absolute temperature errors for the sensors however this is incorrect, as this error only arises from the sensor itself and does not include errors, such as the measurement resolution or electronic accuracy, of the measurement instrument itself [345].

Tables 4.1 and 4.2 show the true temperature accuracy including that of the measurement instrument, in this case a calibrated DT-670 silicon diode connected to a Lake Shore 340 temperature controller. All data has been linearly interpolated between given temperature points. The absolute temperature error data in table 4.2, calculated using [345], was found to exactly match the typical sensor performance expected as given by Lake Shore [346]. This same method was used to determine the absolute error for Cernox temperature sensors. For uncalibrated DT-670s, used occasionally as part of this work, the same method is used where the calibrated accuracy is taken from a Lake Shore tolerance table [343]. For the DT-670A uncalibrated sensors used, from 2–100 K the calibrated accuracy is ± 25 mK. The absolute temperature error for uncalibrated sensor is calculated and shown in appendix C.4.

Temperature (K)	Typical sensor accuracy (mK)	Interpolated accuracy (mK)	Typical sensor sensitivity ($\mathrm{mV\,K^{-1}}$) [347]	Instrument measurement resolution (μV) [348]
4.2	12.0	12.0	-31.6	
10	12.0	12.0	-26.8	
20	12.0	12.5	-15.6	
30	12.0	13.5	-2.0	10.0
50	12.0	15.5	-1.6	
77	22.0	22.0	-1.7	
300	32.0	32.0	-2.3	

Table 4.1: Sensor accuracy and sensitivity for a calibrated DT-670 silicon diode sensor and Lake Shore 340 temperature controller measurement resolution.

Temperature (K)	Measurement resolution converted (mK)	Electronic accuracy (V) [343]	Electronic accuracy converted (mK)	Absolute temperature error (mK) [345]
4.2	0.3	1.6	5.0	17.0
10	0.4	1.4	5.6	17.6
20	0.6	1.2	9.0	21.5
30	5.1	1.1	69.0	82.5
50	6.2	1.1	83.0	98.5
77	5.8	1.0	75.9	97.9
300	4.4	0.6	47.0	79.0

Table 4.2: Electronic accuracy of Lake Shore 340 temperature controller showing absolute combined temperature error calculated from [345].

4.2.3 Suspension interfacing

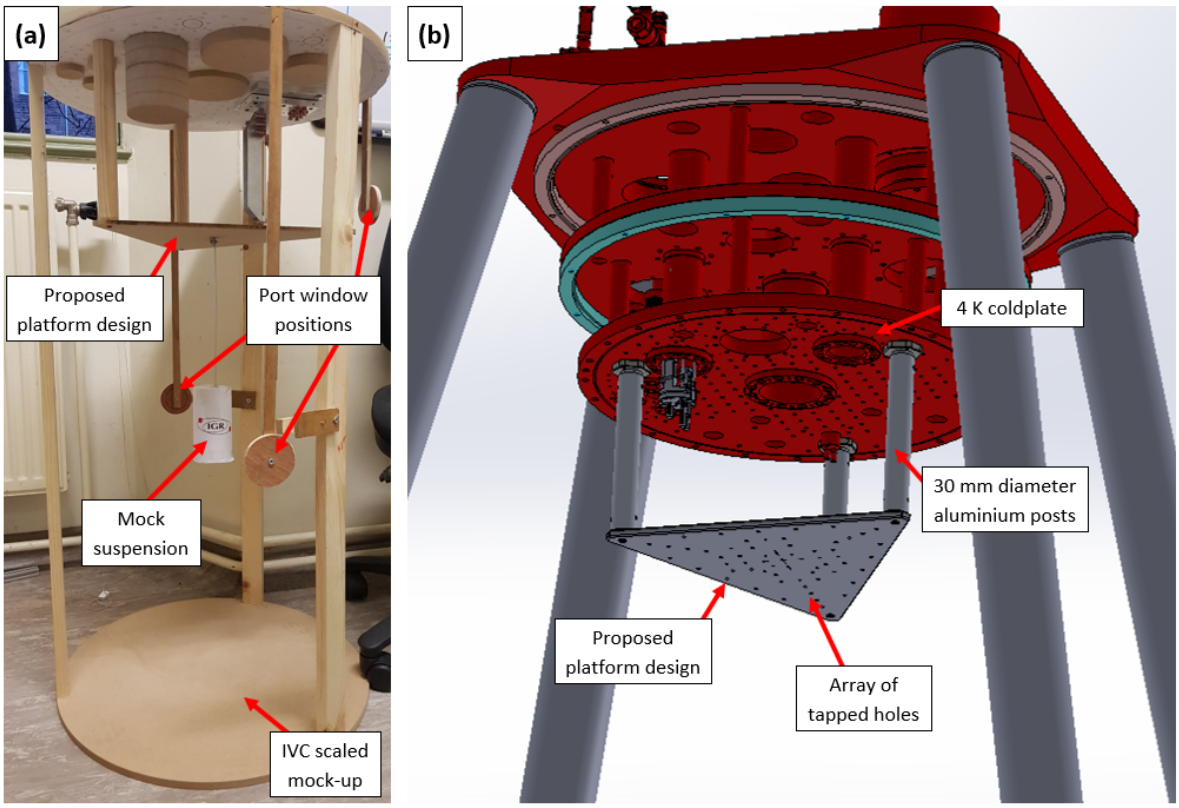

Figure 4.8: (a) Wooden mock-up of IVC with port window positions and proposed platform design. (b) Solidworks drawing of proposed aluminium platform design. CAD drawing supplied by Mr Russell Jones.

In order to affix the suspension into the cryostat, one requires both sound mechanical and thermal connections. The cryostat 4 K coldplate does not allow for easy affixing of large structures due to the uneven surface area produced by the many feedthroughs and blanks shown in figure 4.3(a).

Furthermore, the cryostat is fitted with 3 viewing ports, two of which are located half-way down the OVC, opposite one another, with the other located perpendicular and closer to the coldplate. Due to the position of the two opposite ports half-way down, experiments fitted directly onto the coldplate are unlikely to be observed visually unless they are approximately 60 cm in length. The viewing port windows align through the OVC, 50 K shield and IVC and are made of 5 mm thick, uncoated sapphire, sealed with rubber in the OVC, and indium in the IVC. The 50 K shield windows do not require sealing since this is a not a vacuum tank. For initial suspension testing, it would be beneficial to observe the suspension under cooling to visually confirm any mechanical failure of the suspension.

For these reasons, a wooden mock-up of the IVC complete with port window positions was designed and built, as shown in figure 4.8(a). This allowed for the design of an aluminium triangular platform to be bolted onto the three tapped M10 screws on the coldplate, as shown in

figure 4.8(b). This plate is held in place by three 30 mm aluminium posts to ensure secure mechanical fixing and a sufficient conduction path from the coldplate to the platform. The platform is patterned with an array of tapped holes to allow for the mechanical fixing of experiments.

4.3 Suspension support structure

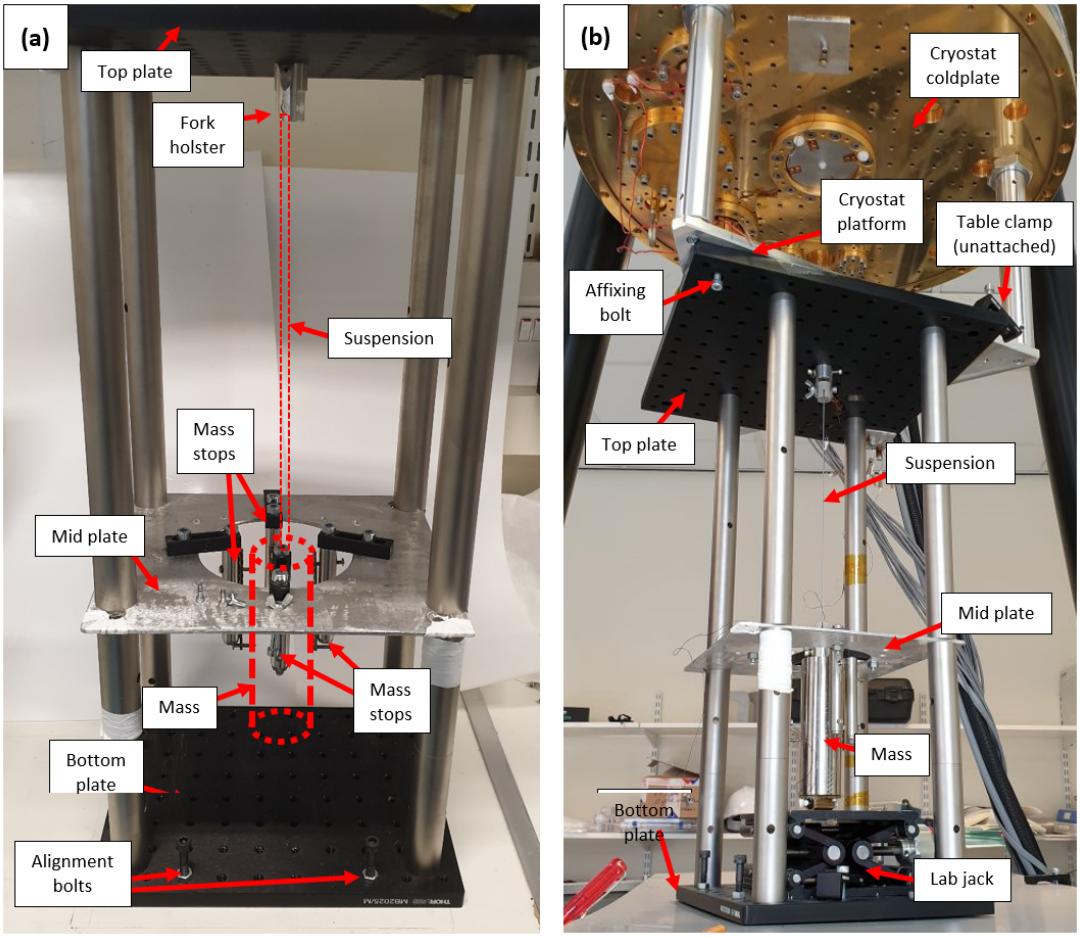

Figure 4.9: (a) Suspension support structure showing location of a proposed silicon suspension with dashed red line. (b) Suspension support structure interfacing with the Leiden cryostat.

A suspension support structure is required to both mechanically support the silicon suspension during assembly as well as protect the suspension in-situ from external mechanical shocks during installation.

The suspension structure design, and how it interfaces with the cryostat and platform, is shown in figures 4.9(a) and (b). The structure consists of an aluminium top plate, an aluminium mid plate and an aluminium bottom plate. The top plate is attached to the mid plate by four 1" diameter, 300 mm long, stainless steel, optical posts designed to allow the insertion of the silicon suspension length [349]. The mid plate is attached to the bottom plate by eight 1" diameter,

100 mm long optical posts, designed to accommodate the suspension mass with clearance below for a lab jack and structure to catch the mass should the suspension fail.

The top plate has a flat, flush surface to allow for easy interfacing with the cryostat platform. The top of the suspension structure is secured into the cryostat platform by the use of two table clamps and one bolt that screws into a tapped hole on the platform stage which allows for both good mechanical and thermal contact.

The mid plate contains a central hole, large enough to accommodate the diameter of the mass as well as four mass stops. These stops consist of posts containing a vertical array of tapped holes. These tapped holes allow small screws to be driven towards the sides of the top and bottom of the mass. By adjusting these screws inward, one can lock the mass in place to stop any horizontal motion. When the suspension is to be hung, one can release these stops enough to free the mass but limit its horizontal motion, or swinging, by keeping the tips of the screws very close to the mass edges.

The bottom plate has three bolts placed in an L pattern, screwed through the plate in order to contact the table below. This allows for the structure to be vertically aligned in both horizontal axes by adjusting the bolts. This minimises any induced stress on the silicon suspension by reducing off-axis hanging. These bolts are visible on the bottom plate of figure 4.9(a).

Fork holsters are attached to the bottom of the top plate, as shown in figure 4.9(a), and top of the mass (not shown) that allow for the suspension fuse ends to be inserted and held in place by passing a retaining bolt through the clearance holes, forming a half-universal joint.

4.4 Suspension design, cooling and results

This section discusses the design, cooling and results from each iteration of the silicon suspension prototype.

The general suspension design was to hang a 1 kg mass on a single silicon ribbon. This was deemed an achievable goal in the timescale for this body of work. As discussed in chapters 2 and 3, silicon's brittle properties make tolerancing of a multi-ribbon or fibre suspension extremely difficult. The mass is not designed to be an interferometer mirror but to purely apply a load to the suspension while undergoing cryogenic cooling. For practical and cost reasons, a metal mass was used throughout.

In total, 10 suspensions were constructed, hung and cooled in the cryostat with varying levels of success. Each iteration sought to improve on the last design, while taking steps to realise the ultimate aim, to demonstrate a cryogenically cooled HC bonded silicon suspension.

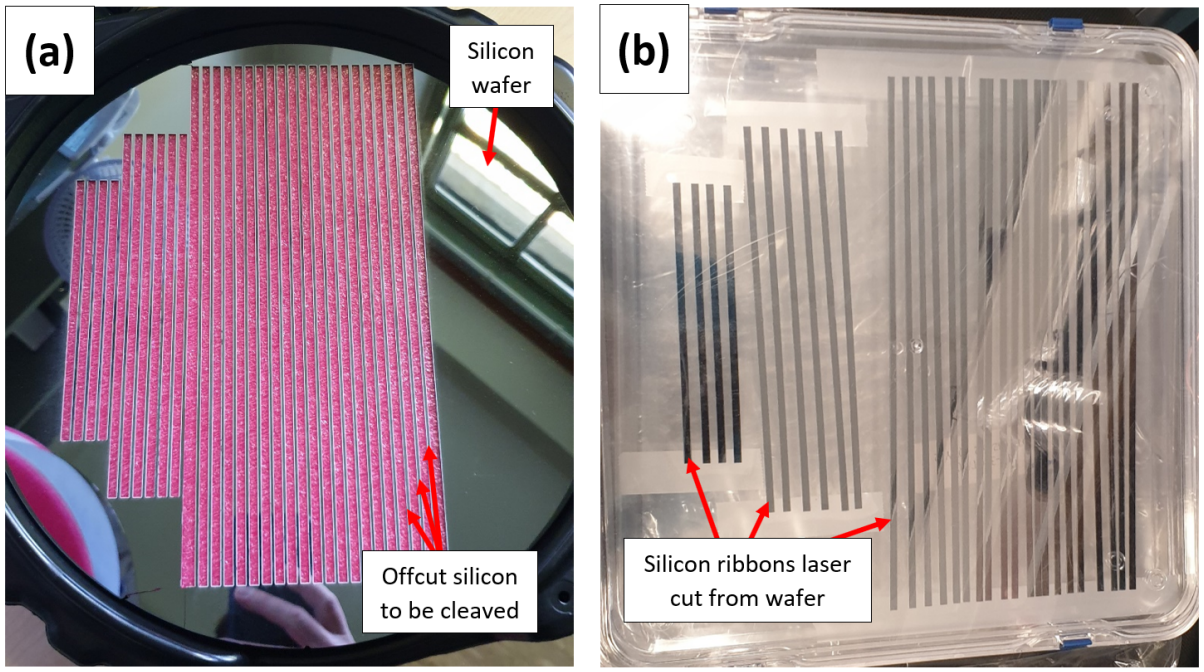

Figure 4.10: (a) 300 mm diameter (100) silicon wafer after laser cutting of ribbons. (b) Laser cut ribbons, 125 mm, 175 mm and 250 mm long.

All silicon suspension elements consisted of 250 mm long silicon ribbons cut from 300 mm diameter, 775 µm thick (100) silicon wafers by LML using the same new laser cutting technique discussed in section 3.4 of chapter 3, as shown in figure 4.10(a) and (b). Some shorter ribbons, 125 mm and 175 mm long, were cut in order to utilise the full wafer although they were not used during any part of this work. Offcuts were cleaved from the wafer shown in figure 4.10(a), significantly increasing the number of available ribbons for testing.

4.4.1 Dummy aluminium suspension

Before the first silicon ribbon suspension was built, a dummy silicon ribbon was used made from shimmed aluminium. This was used for development of hanging procedures involving the suspension structure as well as movement and integration of the structure into the cryostat without the risk of breaking a silicon ribbon. This also allowed for testing of proper mechanical and thermal attachment and routing of temperature sensors and the cryogenic wiring.

Figure 4.11 shows the dummy aluminium ribbon suspension. The left inset image shows the attachment mechanism proposed for the silicon suspension. Strength testing conducted in chapter 3 utilised the fuse end design which allowed for gluing of short silicon ribbons into aluminium fuse ends. This allowed for the insertion into the strength testing apparatus. This same technique was proposed for the initial silicon suspension. The right inset shows the 1 kg mass made of stainless steel hanging freely with the mass stops released, but close enough to minimise swinging in both horizontal axes. There are two more important features shown. The

first is the mass catcher, which replaces the lab jack after suspension hanging and is designed to stop the mass dropping into the bottom of the IVC should the suspension fail. The second is a small 12 Ω resistive heater. This is designed to apply a heat load to the suspension to simulate absorption of heat into the test mass that would result from an incident interferometer beam.

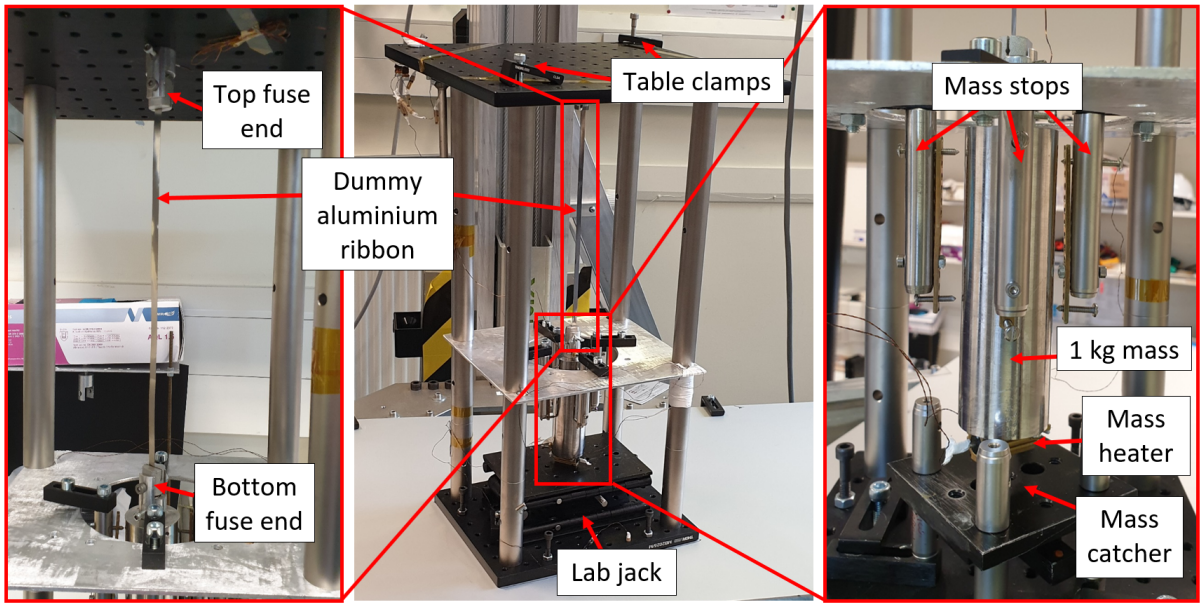

Figure 4.11: Dummy aluminium suspension. Left inset shows fuse end attachment design for suspension element. Right inset shows hanging 1 kg stainless steel mass alongside mass stops, mass heater and mass catcher.

4.4.2 Suspension 1

The first silicon suspension comprised of an offcut silicon ribbon. This offcut was 1.65 mm wide, which means the ribbon will be stressed to 7.7 MPa, 23 % of the breaking stress for the weakest new laser cut sample, or 13 % of the median *control new* sample set from chapter 3.

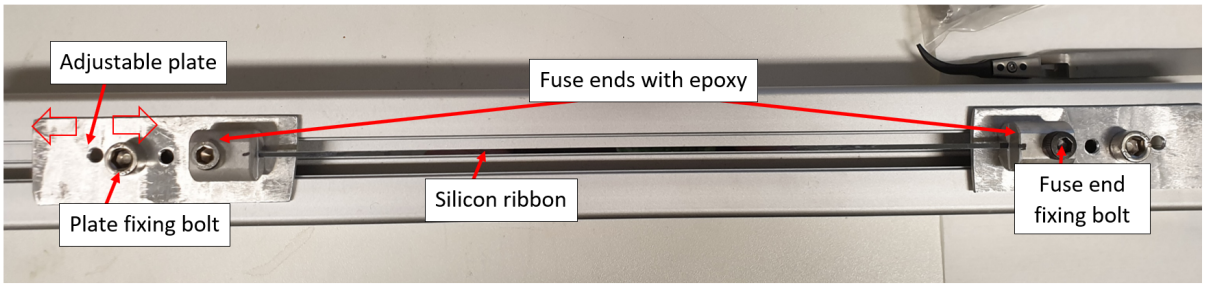

Figure 4.12: Silicon suspension epoxy jig.

The ribbon was epoxied into two fuse ends using a silicon suspension jig to align the suspension, shown in figure 4.12, using the same two-part Araldite 2012 epoxy used for strength testing

in chapter 3 [271]. While this epoxy has no data on cryogenic compatibility it was deemed suitable due to its known high strength properties. The jig has adjustable plates for different lengths of silicon ribbons. The ribbon was then installed into the suspension structure and loaded with the 1 kg mass.

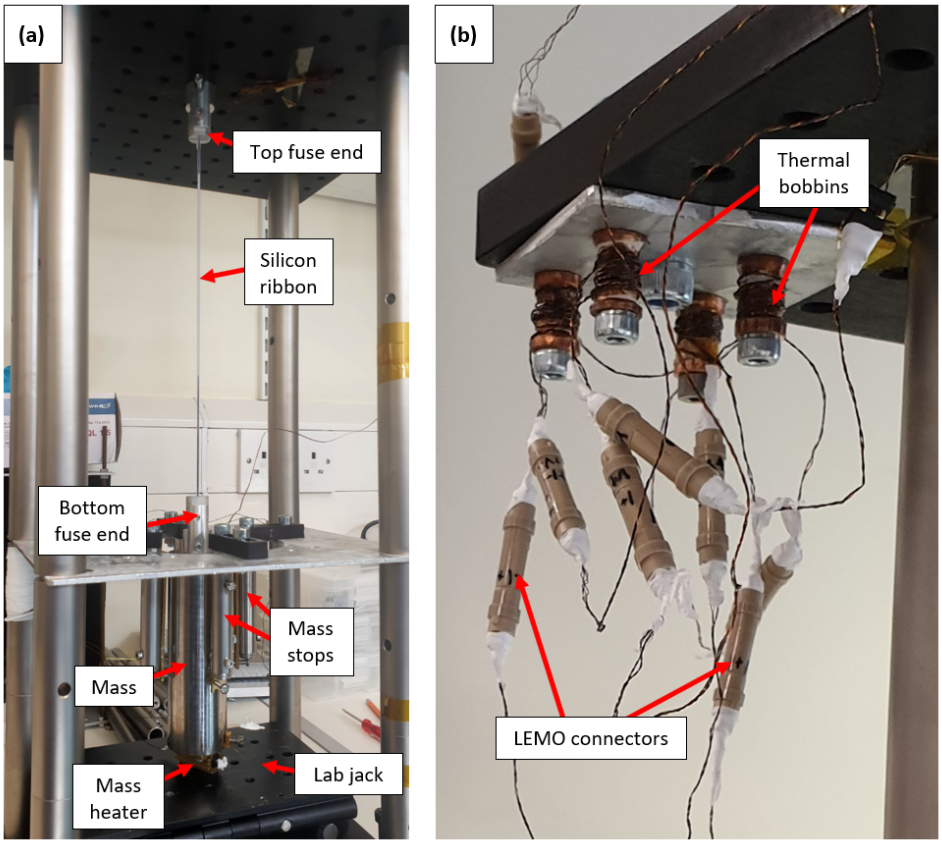

Figure 4.13: (a) First silicon ribbon suspension. (b) Thermal bobbins used for thermal and mechanical anchoring of temperature sensor lead wires.

Every silicon suspension in this chapter was hung in air for a minimum of 24 hours to check for epoxy or bond failure, or failure in the silicon ribbon suspension material.

Figure 4.13(a) shows the first silicon ribbon suspension before being installed into the cryostat for cooling. Cernox and calibrated DT-670 sensors were fitted to the suspension structure in order to measure the temperature a various points. Every temperature sensor affixed to the suspension structure was thermally and mechanically anchored to the top plate before going to the EVF. Thermal bobbins made of copper were designed and built with plug-and-play LEMO connectors on either end to allow for easy repositioning of 4-wire sensor anchoring, as shown in figure 4.13(b). Each bobbin had a cigarette paper attached using cryogenic varnish around its diameter in order to prevent any potential electrical shorting onto the cryostat structure [350]. Each sensor lead wire was then wrapped around the bobbin for a minimum of 10 turns and then varnished in place. Each bobbin was attached to an aluminium plate which in turn was attached to the top plate of the suspension structure by bolts. Each bolt and aluminium plate were coated

in a thin layer of Apiezon N thermal grease to ensure optimal thermal contact [351].

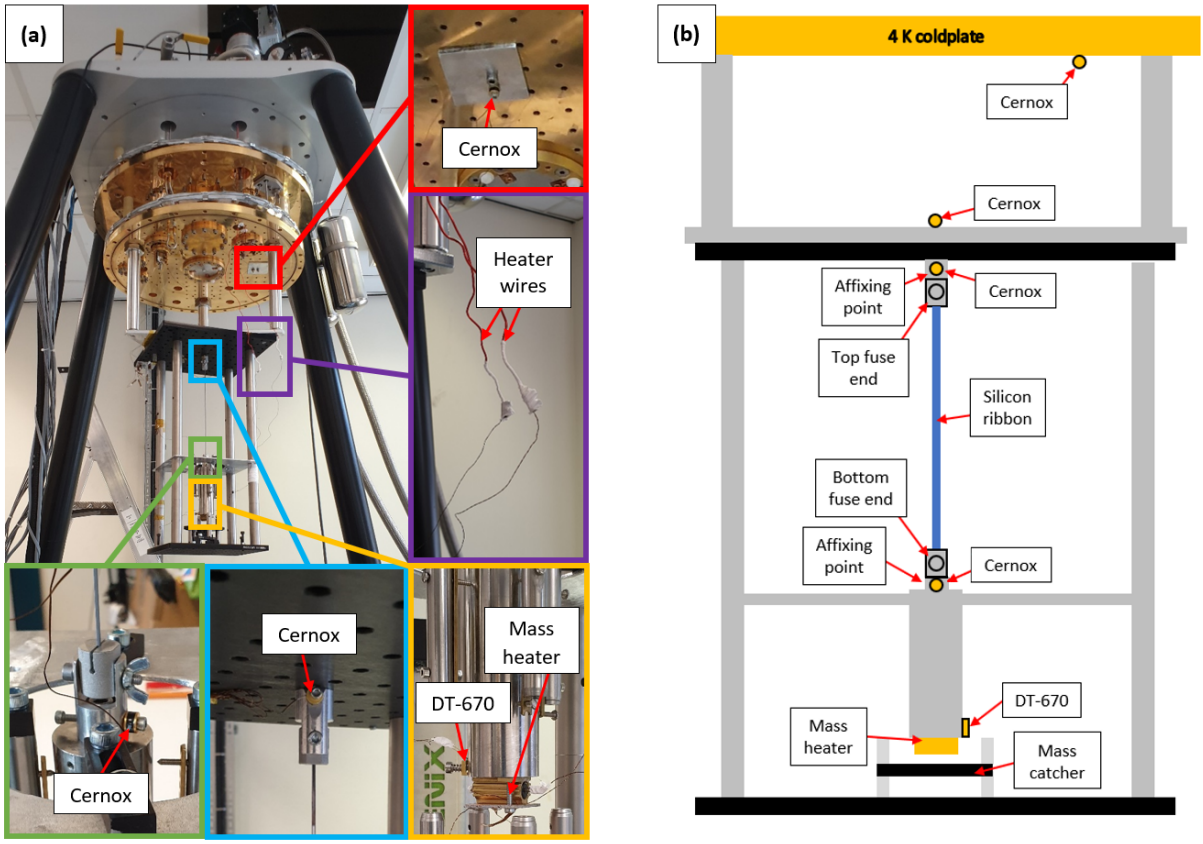

Figure 4.14: (a) Silicon ribbon suspension installed into cryostat with coloured boxes highlighting temperature sensor placement and heater wire connections. The Cernox sensor on the top of the cryostat platform is not shown. (b) Schematic of the temperature sensor type and location within the cryostat.

Figure 4.14(a) shows the hanging silicon suspension affixed into the cryostat with all of the temperature sensors fitted. The heater wire connection, discussed in section 4.2.2, is connected to cryogenic wiring from the stiff copper insulated wires. This wiring is then thermally and mechanically anchored just below the mid plate to one of the suspension structure posts before connecting to the mass heater. Figure 4.14(b) shows the schematic of where and what type of temperature sensors are affixed onto the suspension. It is worth noting that the top and bottom fuse end Cernox sensors are actually attached to the fork holsters.

The cryostat was closed up, evacuated and cooled. Recording of the temperatures of each sensor was taken using a LabView program developed by Dr Karen Haughian, which has been modified by Dr Alan Cumming and Dr Mariela Masso Herrera before being further modified by the author [202]. This program recorded the temperature from the Lake Shore controllers for each sensor at a sampling rate of 1 Hz.

The aim of this experiment was to cool the suspension to as low a temperature as possible before using the heater on the bottom of the test mass to inject a heat load.

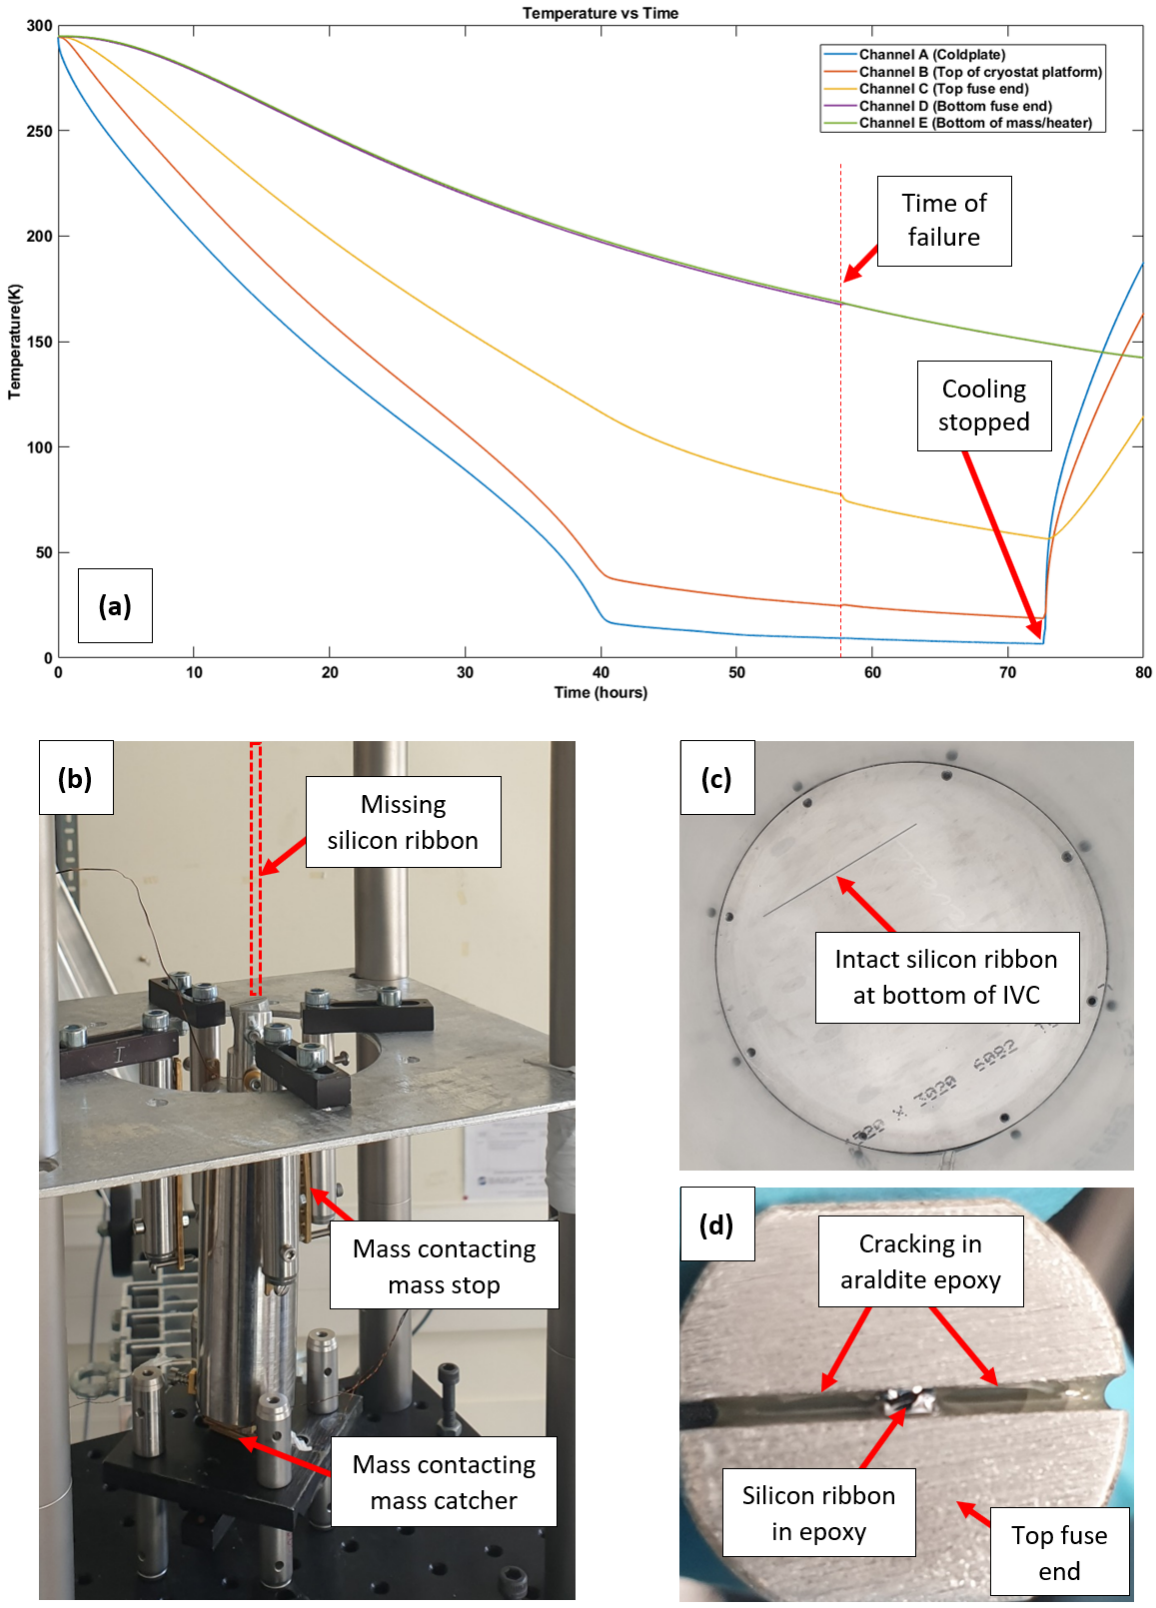

Figure 4.15: (a) Cooling curve of first silicon suspension, highlighting point of failure. (b) Status of mass after failure. (c) Silicon ribbon at bottom of IVC after failure. (d) Evidence of cracking in the Araldite epoxy.

The cooling curve for cooling suspension 1 is shown in figure 4.15(a). After approximately 58 hours of cooling, the suspension appeared to fail, as indicated in the small temperature change seen across all suspension structure sensors. The primary evidence for this is the fact the bottom fuse end on the top of the mass and the bottom of the mass/heater sensor temperature merge, where there had previously been a small thermal lag. This is due to the fact the bottom of the mass would have contacted the mass catcher, while the top of the mass would have contacted one of the mass stops. This means the conductive cooling path, which was previously only through the silicon ribbon would now be through two cooling points at the top and bottom of the mass. It is reasonable to assume the mid plate and bottom plate temperatures are of similar temperature due to the four, 1" poles connecting the two plates, providing a large conduction path, and so the temperature of the mass becomes more uniform as this new conduction path is formed.

This failure of the suspension was confirmed visually by looking through one of the cryostat port windows. When the cryostat was opened, the assumption about the mass touching the mass catcher and one of the mass stops was also confirmed, as shown in figure 4.15(b).

The suspension ribbon fell into the bottom of the IVC but landed intact as shown in figure 4.15(c). It is not known whether the suspension initially failed at the top or bottom fuse end but the suspension ribbon failed across both fuse end interfaces. It is possible the suspension failed at the top fuse end first, causing the ribbon to rotate around the bottom fuse end's half-universal joint. As the bottom fuse end rotated and contacted its fork holster this may have then jolted the ribbon causing a secondary failure at the bottom fuse end interface, allowing the ribbon to freely fall out of the suspension structure and into the bottom of the IVC. Further evidence to support this failure location is shown in figure 4.15(d). This shows cracked Araldite epoxy within the top fuse end slot. By contrast, the bottom fuse end showed no visible evidence of cracking in the clear epoxy.

There are a number of possible reasons that could have caused the suspension to fail. The obvious visible cracking of the Araldite epoxy was of immediate concern as this has potentially mechanically shocked the ribbon, causing it to fail. This motivated a move away from Araldite 2012 which is discussed below.

Further to this, the top fuse end temperature sensor read 77.4 K at the time of failure from figure 4.15(a). The fact this temperature is so close to the boiling point of nitrogen was interesting. It was initially suspected that the rough laser cut edges of the silicon may have trapped air, and hence nitrogen, between the epoxy-ribbon interface. If this nitrogen then liquefied this may have caused some form of unwanted interaction between the epoxy-ribbon interface.

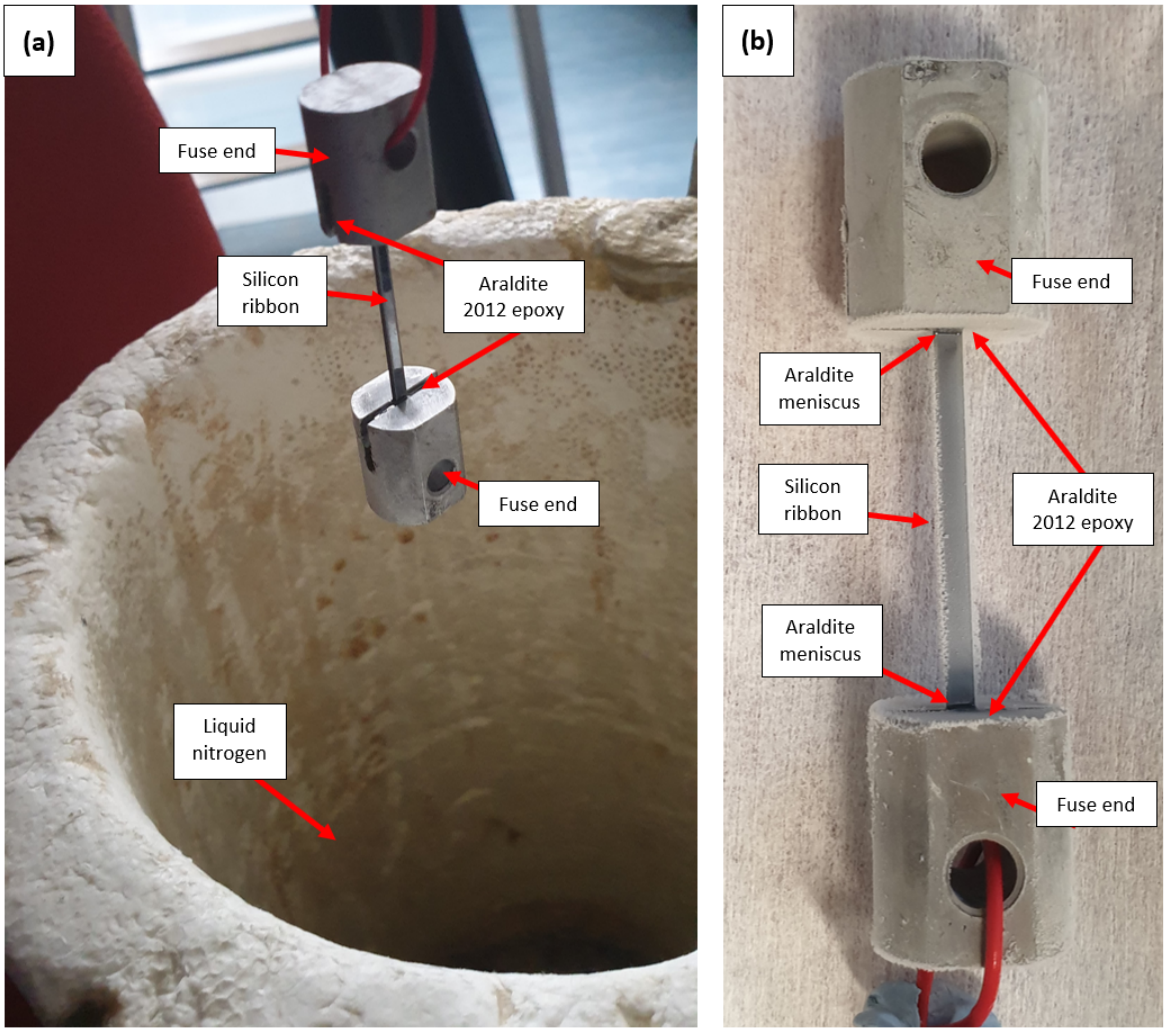

Figure 4.16: (a) Silicon fuse end assembly before being placed directly into liquid nitrogen below. (b) An example of an assembly that appeared to survive liquid nitrogen cooling, during warm up. It was later discovered the sample had failed internally and was held together only by the epoxy meniscus.

To investigate this potential failure mechanisms, a number of 45 mm long silicon ribbons were taken from the same laser cut wafer batches used in chapter 3. These were epoxied into fuse ends, as shown in figure 4.16(a), using Araldite 2012. These assemblies were then placed directly into liquid nitrogen. Some of the samples resulted in the top or bottom fuse end falling off once the assembly had fully cooled in the liquid nitrogen. Other samples appeared to survive cooling down but once left to warm up, as shown in figure 4.16(b) were manipulated by hand where it was discovered the silicon has broken internally in the epoxy and it was purely the epoxy meniscus around the silicon that was holding the assembly together. It was considered that the violent boiling of the liquid nitrogen may have unduly stressed the assembly so the experiment was repeated by housing the assembly in a small electrical box and slowly lowering it into the liquid nitrogen. This did not change the outcome with all assemblies resulting in loose

or broken fuse ends. The araldite in all samples appeared to be cracked.

In order to remove the issue of rough edges, one sample was taken from the CMP treatment #2 batch in chapter 3. Once epoxied into the fuse end assembly it was placed in the liquid nitrogen, inside the small electrical box to protect it from any forces exerted from the liquid nitrogen boiling. Once removed and returned to room temperature the sample was gently manipulated by hand and appeared to be fully intact. The sample was strength tested in section 3.6.3 of chapter 3, resulting in the second weakest sample of the batch at 63.3 MPa. This appeared to demonstrate that cooling the sample with no rough laser cut edges helps to mitigate for this failure mechanism.

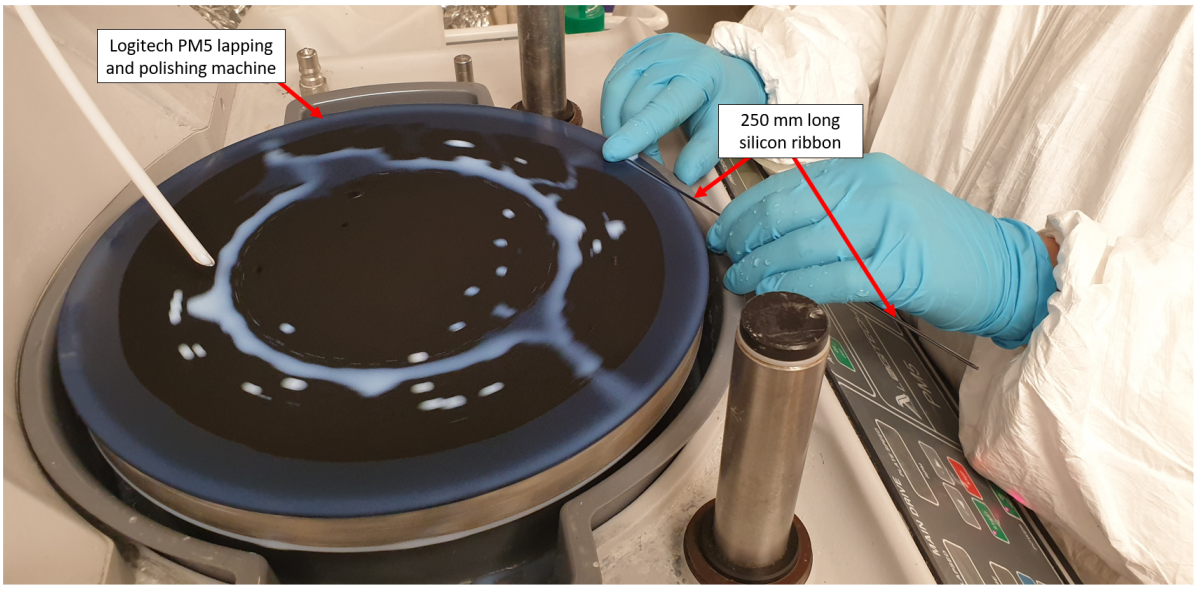

Figure 4.17: Hand polishing 250 mm long silicon suspension ribbons on the Logitech PM5 polishing machine. Photograph taken by Dr Gregoire Lacaille.

It was therefore necessary to lap and polish suspension length 250 mm silicon ribbons at the ends where they interface with the epoxy and fuse end. Due to their length, this had to be done by hand using the Logitech PM5 polishing machine, discussed in section 3.6.1 of chapter 3, and is shown in figure 4.17. This was initially attempted on the suspension 1 ribbon offcut that had fallen into the bottom of the IVC since the mechanical shock from the fall may have weakened the ribbon. After a successful result, this was repeated on a new ribbon offcut, to be used for suspension 2.

4.4.3 Suspension 2

Suspension 2 differs from suspension 1 in a few areas. Firstly, a new silicon offcut ribbon was used, this time 1.79 mm wide, reducing the breaking stress to 7.1 MPa, 12 % of the median *control new* sample set from chapter 3. This ribbon was lapped and polished at both ends in order to remove the potential for any trapped air in the epoxy-ribbon interface.

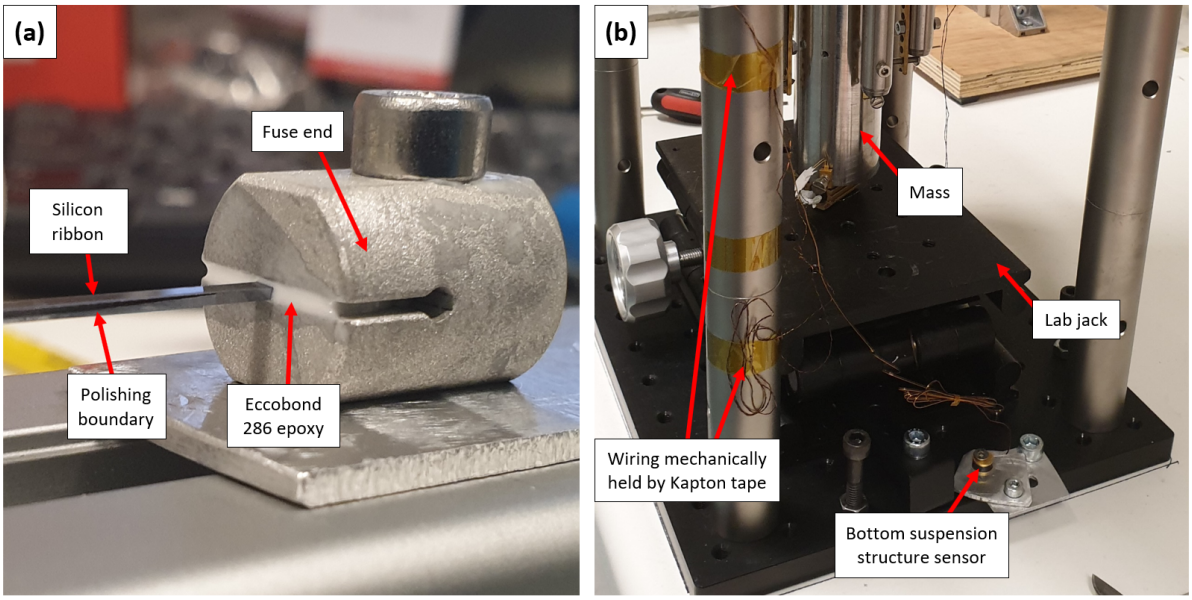

Figure 4.18: (a) New Eccobond 286 epoxy with polished silicon ribbon where the polishing boundary is visible. (b) New position of Cernox temperature sensor, moved from the top of the cryostat platform to the bottom of the suspension structure.

To better increase the chances of the suspension survival, it was decided to replace the Araldite 2012 epoxy entirely. As such, a new two-part epoxy, specifically designed for cryogenic operation, was used. This new epoxy was Eccobond 286 also known as Loctite Ablestik 286 [352]. This epoxy contained a ceramic-like filler and behaved similar to Araldite 2012 at room temperature but appeared to be much harder when set. The epoxy was not transparent and cured opaque white, as shown in figure 4.18(a).

The final change between suspensions was the repositioning of the cryostat platform Cernox sensor. It was decided to move this to the bottom plate of the suspension structure in order to measure the difference between the top and bottom of the suspension structure, as shown in figure 4.18(b). Due to the sensor size, aluminium adapter plates were used to attach the sensor onto the bottom plate. Each plate and bolt used was coated in a layer of Apiezon N thermal grease to ensure optimal thermal contact. Sensor wiring was mechanically held in place by Kapton tape before being further mechanically and thermally anchored by thermal bobbins attached to the top plate. Kapton tape is regularly used throughout the suspension structure due to its vacuum and cryogenic compatibility [298]. This new sensor placement would give a better indication of the conduction through the four, 1" diameter posts that separate the top, mid and bottom plates of the suspension structure.

The suspension was cooled for 150 hours before cooling had to be switched off due to the Christmas break. The suspension survived this cooling down period, with the top fuse end sensor reaching a minimum of 20.5 K and the bottom of the mass reaching 54.0 K. When cooling was stopped, the suspension warmed and failed at the bottom fuse end interface. The temperature

at this interface was recorded to be 268.4 K at the time of failure. The cause of this failure was unable to be determined but it was clear that suspension 2 was a significant improvement on suspension 1.

4.4.4 Suspension 3

Suspension 3 used the same 1.79 mm wide silicon ribbon as suspension 2, where the broken bottom fuse end section was re-polished and a new fuse end attached with Eccobond 286 epoxy, reducing the suspension length to 242 mm.

A new ambient room temperature sensor was fitted. This sensor was a simple PT1000 attached to the top of the cryostat externally. This was to monitor the temperature of the lab during cryostat operation.

The most significant change for this suspension was the design and build of a custom made cryogenic switch, herein referred to as the "cryoswitch". This cryoswitch removes the ambiguity around suspension failure and means one does not have to rely on temperature readings to establish exactly when the suspension has broken. The design principle worked on measuring the voltage across a circuit which was kept open when the mass was hanging. If the mass fell, the mass would contact a plate, completing the circuit which would be recorded in the LabView program.

The external wiring for the cryoswitch is visible at the top of figure 4.5(b) in the form of the green and purple twisted wiring. A 5V signal was placed across these wires using a NI USB 6211 device.

The cryoswitch is shown in figure 4.19. A copper plate was bolted onto the bottom of the mass heater. Two copper wipers were soldered onto the bottom of this plate. These copper wipers ensured a good chance of contacting the contact plate if the mass was to fall. The copper contact plate is attached to the top of the mass catcher. Both plates have cryogenic wiring soldered to them. These wires were connected to the USB-6211 box, with the 5 V signal across them. If the mass failed, the copper wipers would contact the contact plate, short circuiting this 5 V potential which is then recorded in the LabView program in order to determine when the suspension failed. Since one only had to discriminate between 0–5 V no shielding or special wiring was required externally. The internal wiring consisted of the Lake Shore Quad-Twist WQT-36 phosphor bronze cryogenic wire, which was mechanically and thermally anchored both at the suspension structure poles, as shown in figure 4.19, and at the thermal bobbins on the top plate. This design required the mass catcher posts to be shortened to accommodate the cryoswitch.

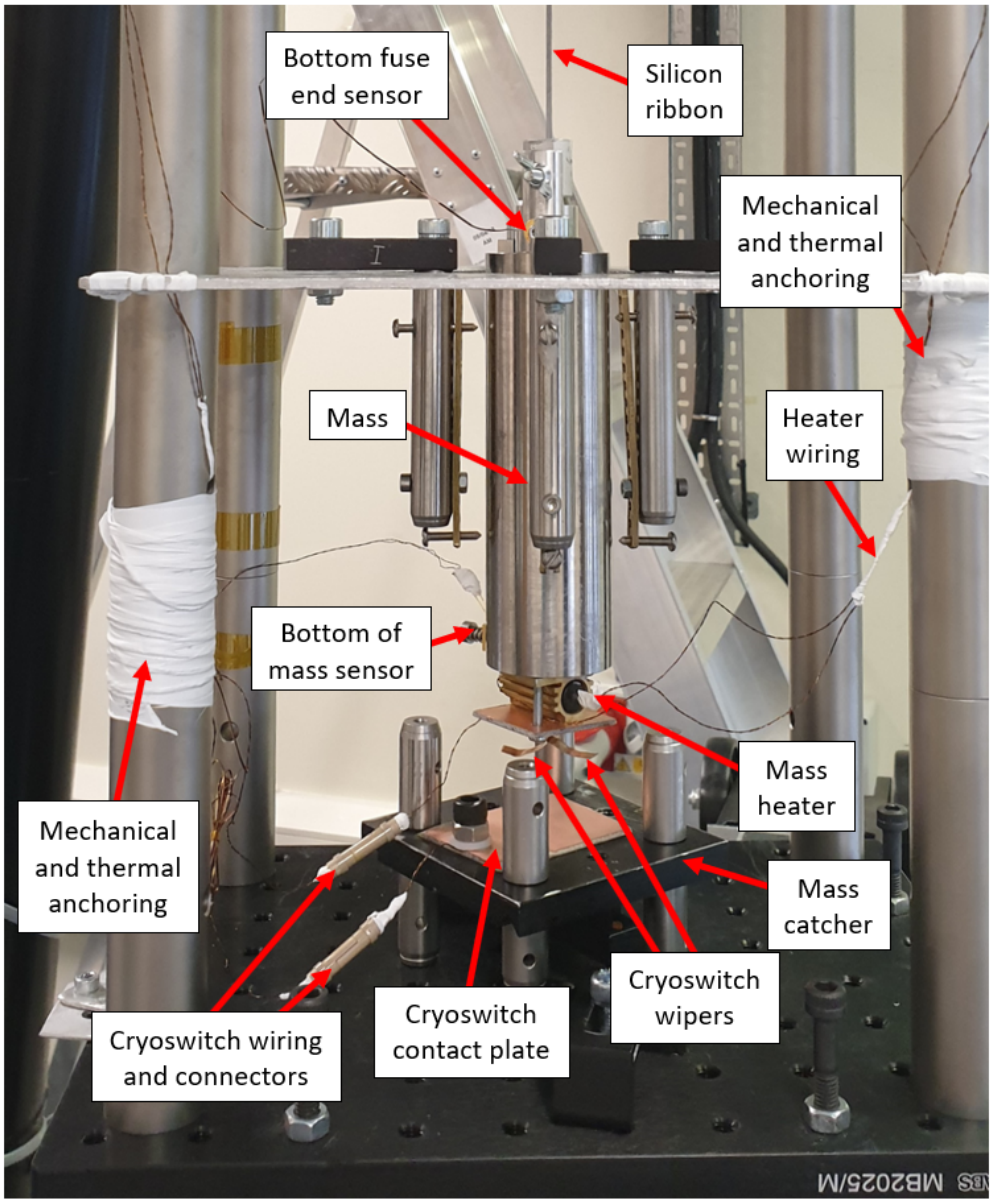

Figure 4.19: Cryoswitch designed detailing components and wiring.

The cryostat was evacuated and the suspension was cooled. This suspension failed after only 37 hours of cooling, with the top fuse end at 117.1 K and the bottom fuse end at 190.4 K. The failure caused the ribbon to fall to the bottom of the IVC where this time it shattered into multiple pieces. After detailed inspection of the Eccobond, fuse ends, broken ribbon and temperature data, the location and cause of this failure was unable to be determined. Due to this failure, it was not possible to attempt the aim of injecting a heat load into the bottom of the suspension.

4.4.5 Suspension 4

Suspension 4 was identical to suspension 3, except the same 1.65 mm wide, silicon ribbon used for suspension 1 was re-used for suspension 4. As discussed at the end of section 4.4.2, this ribbon was lapped and polished as part of the hand polishing trial on a long silicon ribbon. This ribbon was now only 234 mm long due to it being removed from the 8 mm deep, epoxied fuse ends slots. The ribbon was inserted into two new fused ends with the Eccobond 286 epoxy.

Furthermore, this ribbon was mechanically shocked by dropping to the bottom of the IVC. As such, hanging the ribbon in air for 24 hours in order to test the structural integrity of the ribbon was vital. This hang showed no structural issues with the silicon ribbon, so the suspension was installed into the cryostat, as shown in the inset of figure 4.20, and cooled.

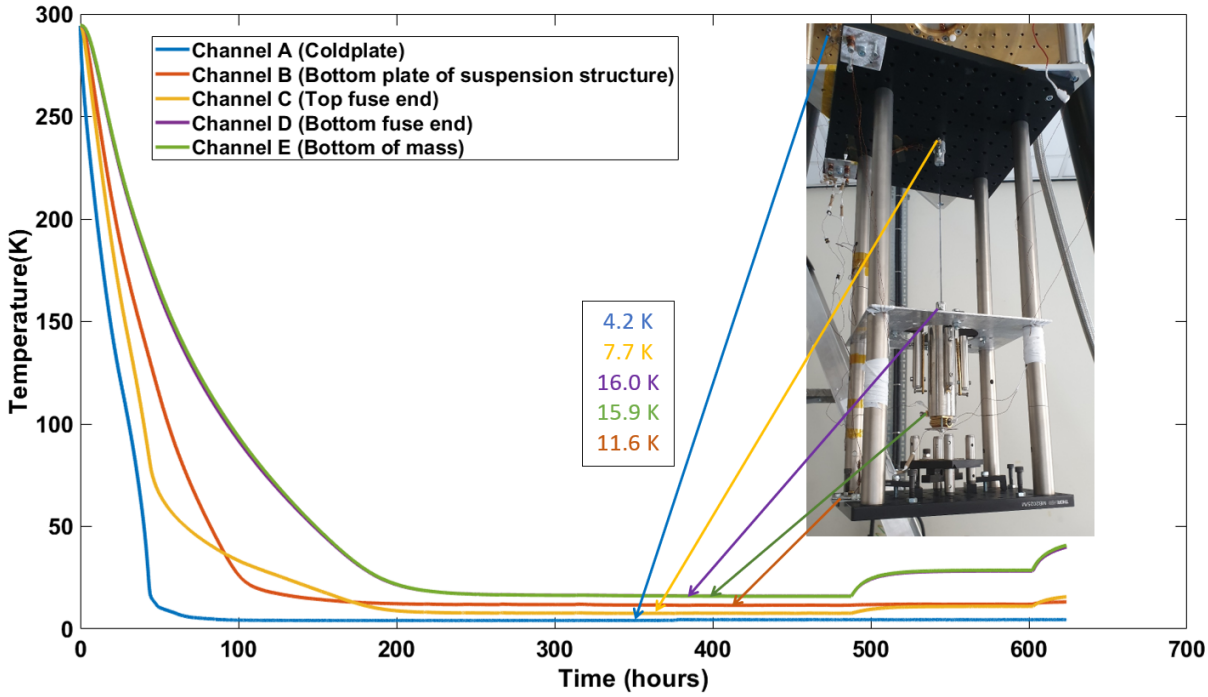

Figure 4.20: First successful cryogenic silicon suspension. Inset shows the suspension with coloured lines corresponding to temperature readings from sensor locations on the suspension structure.

As shown in figure 4.20, the first cryogenically cooled silicon suspension has been demonstrated.

The suspension reached its thermal equilibrium after approximately 250 hours of cooling, with the top fuse end sensor reading 7.7 K, the bottom fuse end sensor reading 16.0 K and the bottom of the mass sensor reading 15.9 K. The bottom suspension structure temperature sensor stabilised at 11.6 K with the coldplate sensor held at a constant 4.2 K.

Initial results of the temperature profile of the suspension structure appear sensible. Firstly, the suspension temperatures increase the further away one gets from the coldplate. This is expected when an ambient heat load impinges on the suspension structure. This could be via

conductive heating, although this is minimised due to the fact any and all wiring touching the suspension structure is thermally anchored to the coldplate or thermal bobbins at the top of the suspension structure. Thermal convection is not possible due to the UHV environment in the IVC which is enhanced by cryogenic cooling, ensuring free gas molecules within the system freeze to cold surfaces, effectively creating a cryotrap. It was found the IVC pressure dropped from $\approx 10^{-5}$mbar to $\approx 10^{-7}$mbar when transitioning from room to cryogenic temperatures. This leaves radiation as the potential remaining heat load on the suspension.

Logically, the bottom fuse end and bottom of the mass temperature sensors closely matched each other during cool down and thermal equilibrium. This is due to the fact they are effectively attached to the same piece of stainless steel, and separated by the length of mass. It is notable however, that the bottom of the mass sensor reads 0.1 K lower, than the bottom fuse end sensor located on the top of the mass. Referring to tables 4.1 and 4.2, this 100 mK difference is well outside the absolute temperature error of both the calibrated DT-670 and Cernox temperature sensors and so is believed to be a real temperature difference. The lower temperature recorded on the bottom of the mass sensor is likely due to the fact it is closest to the 3 wires touching the bottom of the mass, the resistive heater wires and one of the cryoswitch wires. Since the top fuse end sensor, located extremely close to the top plate of the suspension structure, stabilised at 7.7 K and the bottom plate of the suspension structure stabilised at 11.6 K, it is reasonable to assume the four, 1" diameter structure poles are at a temperature between 7.7–11.6 K. Therefore, the suspension structure itself is colder than the mass at all times, as reflected in the cooling curve of figure 4.20. Since these 3 wires are all thermally anchored to these 1" diameter structure poles, they will only aid in heat extraction from the mass. While this is not ideal, these wires, at least, do not contribute to the heat load on the mass, providing further evidence that radiation, likely from the IVC is the cause of the ambient heat load. However, since no temperature sensor is fitted to the IVC for this experiment, it was not possible to definitively identify this as the source of the ambient heat. Experimental troubleshooting of potential heat sources was conducted for subsequent suspensions, as detailed in the relevant section. More detailed analysis of potential heat sources or losses is conducted in section 4.5.

With a successfully cooled and thermally stabilised silicon suspension, one could now investigate applying a further heat load into the suspension system, to simulate that of absorption in the test mass from the interferometer laser.

The simple heater circuit, shown in figure 4.21, consists of a set of resistors in series; the 12 Ω resistive heater on the bottom of the mass, inside the cryostat, and an external 110 Ω resistor contained within a box in the lab. An Aim-TTi 1908P digital multimeter (DMM) was able to measure both current and voltage simultaneously, in order to calculate the power put into the bottom of the mass [353]. The power to the circuit is supplied by one of the Lake Shore temperature controller unit's analogue outputs in the form of a 10 V DC voltage. One can select increments of 0.01 %, as a percentage of this 10 V to be supplied to the circuit. The LabView

program discussed in section 4.4.2 was modified to record the voltage across and current in the circuit every time the temperature sensor readings were taken.

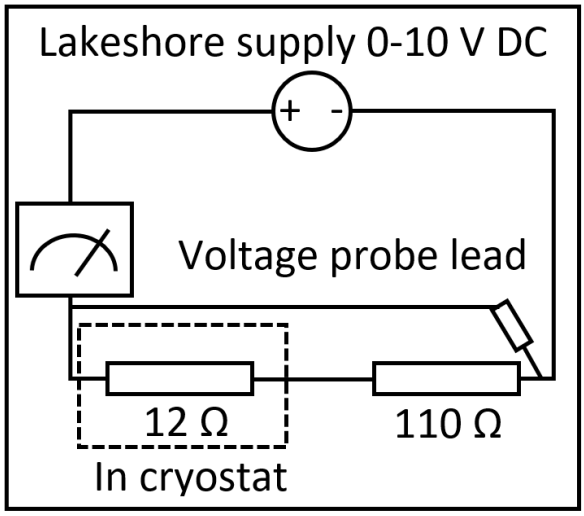

Figure 4.21: Schematic of heater circuit.

Two heating steps are visible around the 500 and 600 hour mark on figure 4.20. This data will not be analysed or discussed here as investigations, triggered by a mismatch in expected voltage and current readings, indicated a short in the heater circuit. This short circuit was traced to a legacy connection made inside the vacuum system, behind a high voltage 4-point vacuum feedthrough.

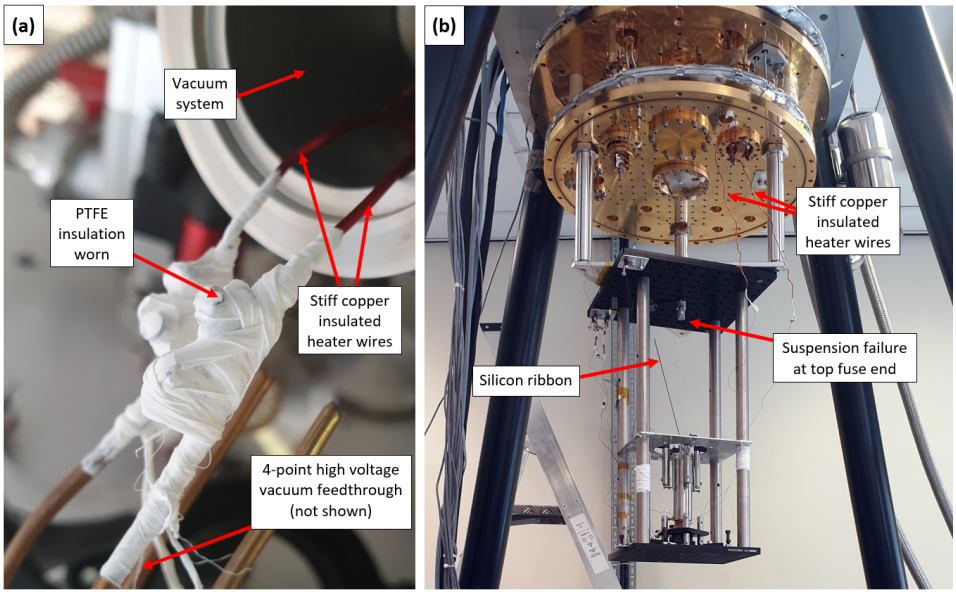

Figure 4.22: (a) Identification of short circuit in heater wire circuit where worn insulation allowed for electrical contact onto the grounded vacuum system tubing. (b) Suspension failure at top fuse end.

As shown in figure 4.22(a), protective PTFE insulating tape had worn away which shorted the heater wires onto the inside tubing of the vacuum system which is grounded by both the turbo pump and the scroll pump safety earthing connections. This connection was properly insulated to avoid future short circuiting scenarios.

During this investigation, either through mechanical shocking, due to researchers working around the cryostat, or thermal shocking, switching the compressor on and off during the short circuit investigation, the suspension failed. The failed suspension is shown in figure 4.22(b). If the suspension were to break at the bottom fuse end, one may be able to give more weight to mechanical shocking over thermal shocking, as thermal shocking would be more likely to affect the top fuse end first due to its proximity to the coldplate. As such, the exact cause of this suspension failure was unable to be determined.

The bottom fuse end was deliberately broken off to allow for re-polishing and to avoid over-stressing or breaking the ribbon by removing this extra mass on the end. When the ribbon was to be broken from the bottom fuse end, it was noted how easily the parts separated. This implied that the ribbon had also failed at the bottom fuse end during the cooling of suspension 4 and only the hard epoxy meniscus was holding it together.

After the suspension failure, all lab activity immediately ceased and the author could not access experimental labs for 6 months due to the COVID-19 pandemic.

4.4.6 Suspension 5

Suspension 5 had some substantial changes compared to suspension 4. The successful hang and cool of suspension 4 implied that the experiment was repeatable so no mechanical changes were made to the experiment, only sensing and vacuum improvements.

Although the 1.65 mm wide silicon ribbon in suspension 4 broke, it could be used again. Both ends of this ribbon were re-polished, with a further reduction in length to 218 mm. This was due to the removal of the fuse ends from suspension 4 and along with the 8 mm of silicon ribbon encased in them.

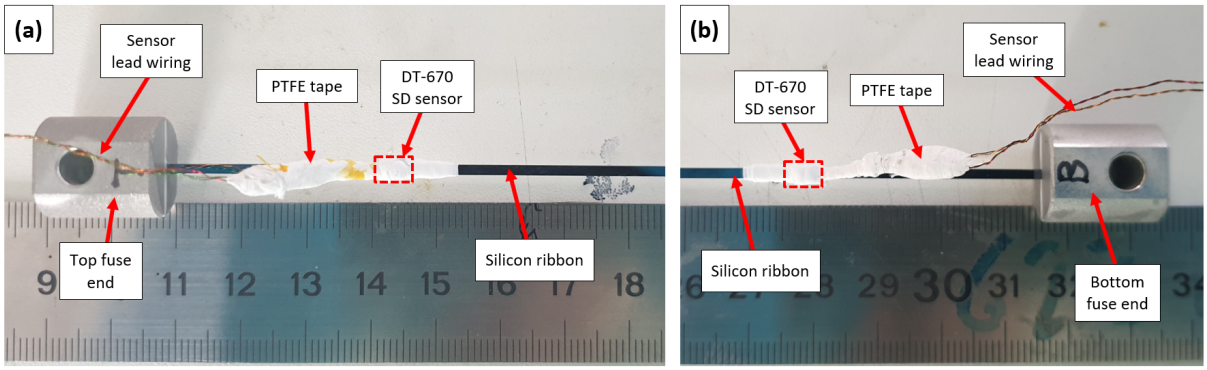

Figure 4.23: (a) Top ribbon sensor and (b) bottom ribbon sensor attached to silicon ribbon for suspension 5.

Previous suspensions only had temperature sensors near the top and bottom fuse ends that were mounted on the fork holsters for each fuse end. This means the thermal path from the mass up to the top of the suspension structure involved passing through the bottom fork holster, the bottom fuse end retaining bolt, the bottom fuse end, the epoxy and then the silicon ribbon, following the reverse path at the top fuse end area. After successfully hanging the ribbon in air, the suspension was disassembled to allow for the fitting of two temperature sensors directly onto the ribbon. These sensors would give a measurement of the temperature gradient across the ribbon which would aid understanding of heat flow throughout the system.

Due to the width of the ribbon and temperature sensor package options available, it was decided that the DT-670 SD package sensors were the only option as Cernox sensors were too wide to be mounted securely. Mounting sensors directly onto very delicate silicon suspension ribbons was not trivial and required much practice before finding an acceptable solution that satisfied good thermal and mechanical contact. The solution was to use a thin layer of Apiezon N thermal grease to mount the sensor thermally onto the ribbon and mechanically secure the sensor in place using self-amalgamating PTFE tape. The sensor attached near the top fuse end, referred to as the top ribbon sensor, is shown in figure 4.23(a). The sensor attached near the bottom fuse end, referred to as the bottom ribbon sensor, is shown in figure 4.23(b). Both sensors are calibrated DT-670 SD sensors.

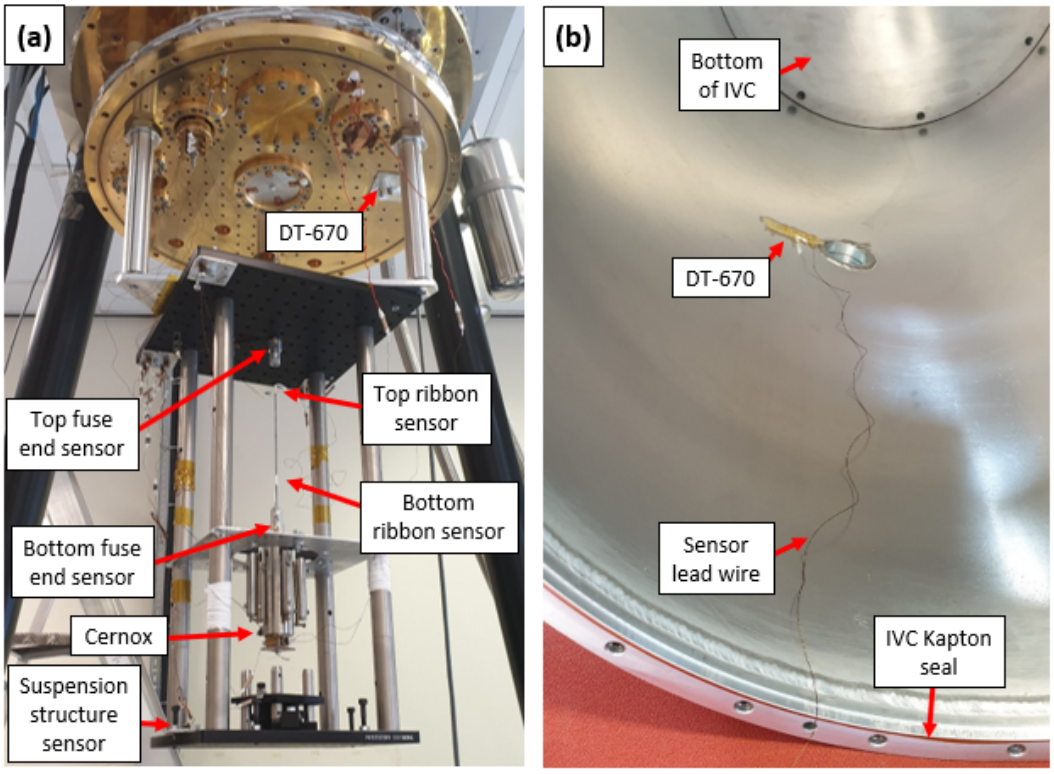

Figure 4.24: (a) New sensor layout for suspension 5. (b) IVC wall sensor fitted next to port window.

The calibrated DT-670 SD sensor affixed to the bottom of the mass was used for one of the ribbon sensors. A spare uncalibrated DT-670 SD sensor was acquired and it was decided that it should be switched with the Cernox sensor on the coldplate. This then allowed the more accurate Cernox sensor to replace the previous calibrated DT-670 SD sensor on the bottom of the mass. This was to increase the accuracy of all the sensors that are part of the suspension structure and thermal path leading from the heater. This change is shown in figure 4.24(a).

In order to investigate ambient heat loads, a new uncalibrated DT-670 SD sensor was fitted to the IVC wall next to the port windows, as shown in figure 4.24(b). This sensor was affixed to the IVC wall with thermal grease, followed by cigarette papers that were then coated in cryogenic varnish. This varnish was then allowed to set hard, holding the sensor in place. The sensor lead wire was connected to a LEMO connected on the EVF once the IVC was lifted into place, just before the IVC was sealed.

At the time of this experiment, only two Lake Shore 340 temperature controllers were available. Both can only take 4 temperature sensor inputs, and so the ambient temperature sensor previously fitted was disconnected for this experiment in order to accommodate the 8 primary sensors; coldplate, top fuse end, top ribbon, bottom ribbon, bottom fuse end, bottom of mass, bottom of suspension structure and IVC wall. A schematic of this new sensor set up is shown in figure 4.25(b).

Finally, a leak valve was fitted onto the vacuum system primarily to allow a slow re-pressurisation of the vacuum chamber. The Virgo detector has previously suffered from contaminants, originating from scroll pumps, damaging and breaking the fused silica suspensions during venting [354]. If a silicon suspension survived cooling and warming back up to room temperature it could be used again. One can then isolate the vacuum system from both the scroll and turbo pumps by closing the nitrogen gas-operated valves. By opening the leak valve, one can then slowly re-pressurise the IVC back to room pressure, minimising the risk of injecting contaminants into the suspension system at speed. Fitting this leak valve also allowed for the potential to leak helium contact gas into the system for rapid and complete cooling of the suspension structure and IVC. This was never performed during the course of this work but should be pursued in future work in order to investigate ambient heat load effects.

The suspension was reassembled with top and bottom ribbon sensors attached and installed into the cryostat. When the cryostat was evacuated, it was discovered the pressures were higher than previous runs indicating a potential vacuum leak issue. After cleaning all seals, this appeared to fix the pressure issue although evacuating took longer than expected, by a few hours, compared to suspension 4. However, the IVC reached 6.2×10^{-5} mbar, similar to suspension 4, which was sufficient for cooling. The experiment was cooled but failed after only 48 hours of cooling, as shown in figure 4.25(a).

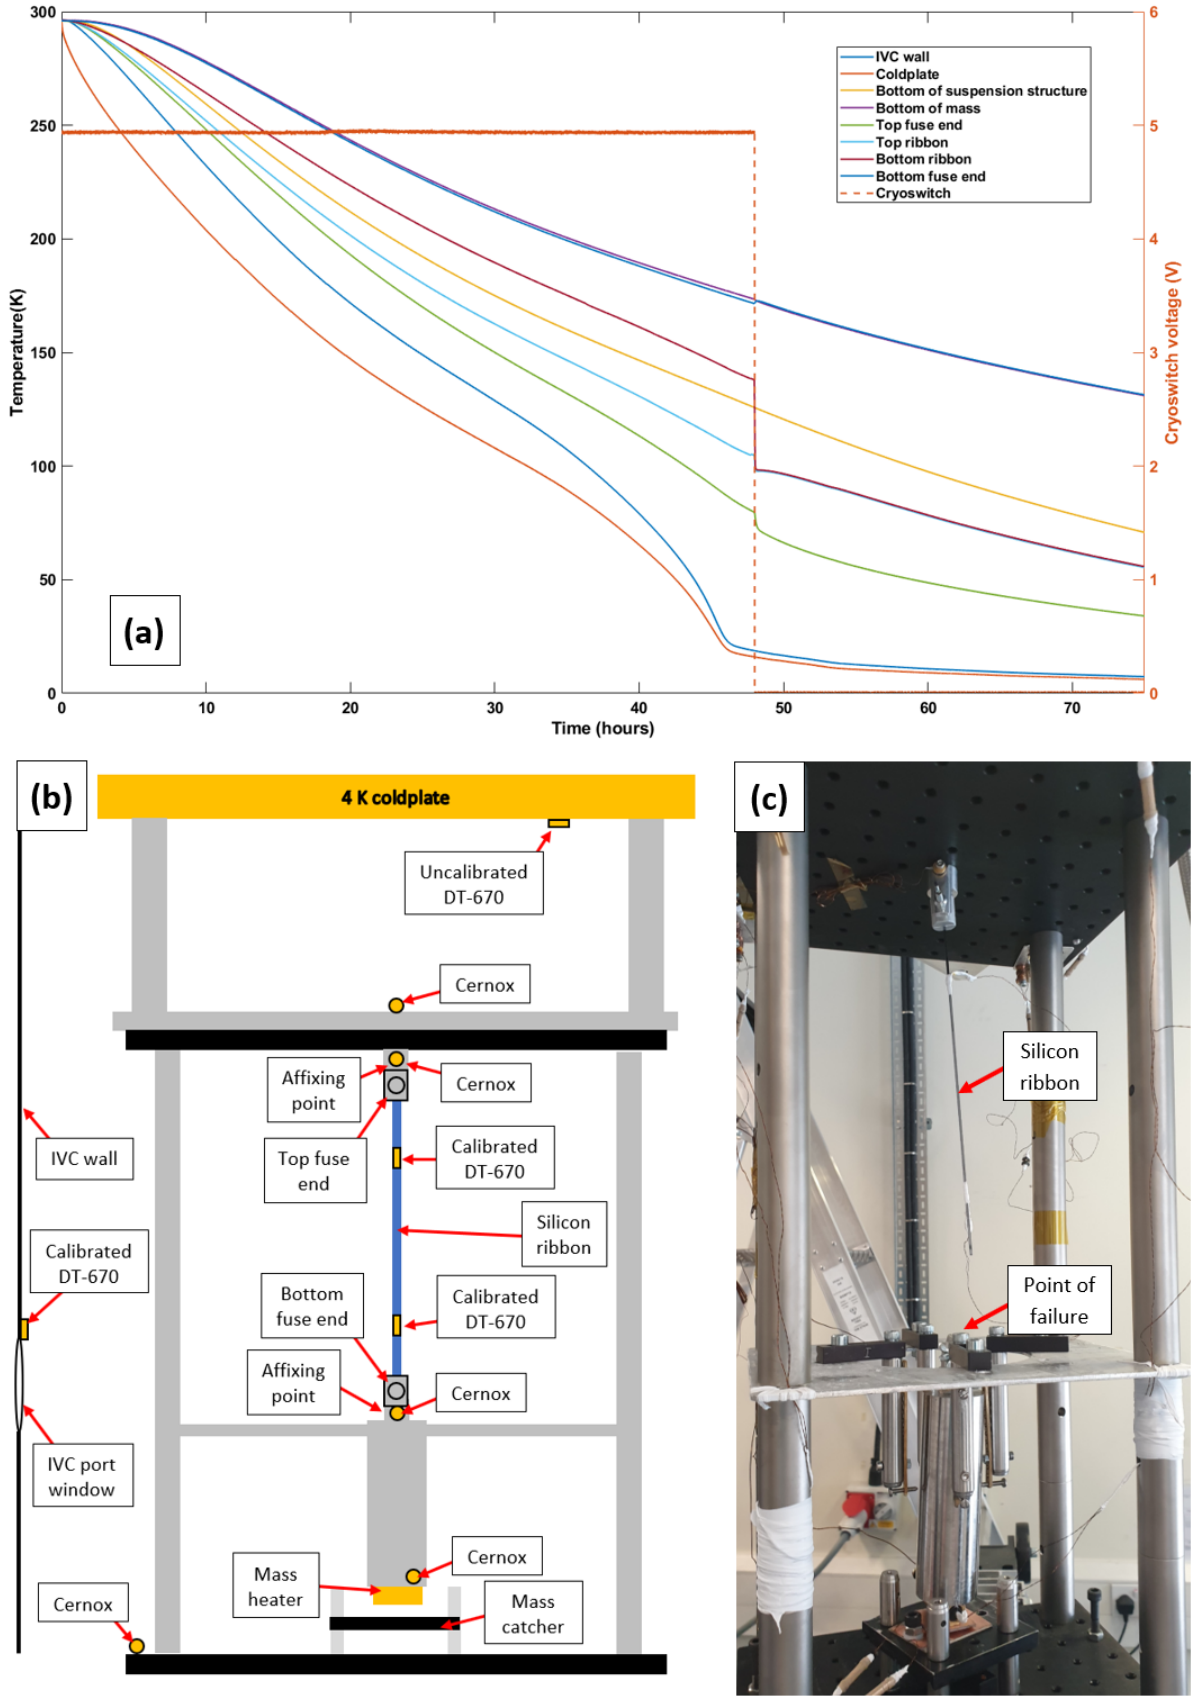

Figure 4.25: (a) Suspension 5 cooling curve, highlighting the cryoswitch's successful operation. (b) Schematic for sensor layout in suspension 5. (c) Failure of suspension 5.

The cryoswitch trigger correlated exactly with the steep change in cooling rate observed in figure 4.25(a). The ribbon sensors show a fast thermal response, consistent with a relatively high thermally conductive element like silicon. Furthermore, the temperature drops towards the top fuse end temperature, which itself shows an immediate drop in temperature. From this data alone, this would seem to imply the ribbon has failed at the bottom fuse end interface, since the large thermal load of the mass has been removed. Upon inspection of the suspension after opening, this is exactly what was observed. The ribbon was intact and hanging from the top fuse end, while the mass had fallen into the mass catcher, as evidenced in figure 4.25(c).

There was no obvious reason for this failure at the bottom fuse end and the cause was unable to be determined. As such, this failure did not allow for further analysis of the suspension since thermal equilibrium was not reached.

4.4.7 Suspension 6

Suspension 6 saw only two minor changes made from suspension 5. The offcut silicon ribbons were replaced with a full width laser cut ribbon, 3.5 mm wide. This reduces the stress to 3.6 MPa, 6 % of the breaking stress measured for the median *control new* samples set discussed in chapter 3. This ribbon was lapped and polished at both ends in order to remove the potential for any trapped air in the epoxy-ribbon interface. This ribbon was lapped and polished at both ends before being epoxied with Eccobond 286 into fuse ends.

Upon evacuating the cryostat, similar pressure issues experienced with suspension 5 arose. This time, it was decided the OVC rubber o-ring be replaced to see if this was causing the issue. While this appeared to improve the vacuum, the OVC pressure was only at 1 mbar, approximately 2 orders of magnitude above the ideal scroll pressure. Similarly, the IVC only reached 3.4×10^{-4} mbar. Significant vacuum checks were conducted to investigate any sources of vacuum leaks but none were found. It was decided to cool the experiment with these higher pressures in the system.

Suspension 6 failed at the top fuse end after 58 hours of cooling. It was noted at this point, the cooling rate of the coldplate was significantly reduced indicating either a degradation of the cooling system or the high vacuum pressures allowing greater thermal conduction to impinge on the 50 K shield and IVC meaning the cooling system has to work harder in order to remove the thermal energy. Initial assessment of the cooling system appeared to show sufficient performance, implying the vacuum system was the cause of the issue. This was investigated in more detail for suspension 7.

After six suspension hangs, with one success, it was clear failures were still occurring at the fuse end interfaces. With the epoxy changed and ribbons polished, the only remaining mechanical changes to be made were the fuse ends themselves. It was not suspected that the coefficient of thermal expansions (CTE) of the 1 mm wide aluminium fuse end slot would be a significant problem as this would lead a 225 µm gap between the ribbon and the silicon to be

filled with the epoxy.

Some MATLAB code was written to take in CTE data for different materials from room temperature down to 4 K. By fitting to this data and integrating over the temperature range of cooling, one can determine the integrated coefficient of thermal expansion (or contraction). Given the original length of the material, with the initial and final temperatures, one can then calculate the overall contraction over this cooling regime.

By integrating from 4–293 K over the CTE, it was found that the overall contraction of this aluminium fuse end slot would be 4 µm. As shown in chapter 2, silicon has a more complicated CTE due to the fact it undergoes expansion between 18–123 K, meaning the silicon contracts, then expands, then contracts again during cooling. However, integrating silicon's CTE from 4–293 K leaves a total contraction of 0.18 µm for a 775 µm thick ribbon. The region of expansion from 18–123 K of the CTE of silicon is $\approx$ 6 times smaller than the overall contraction from 4–293 K.

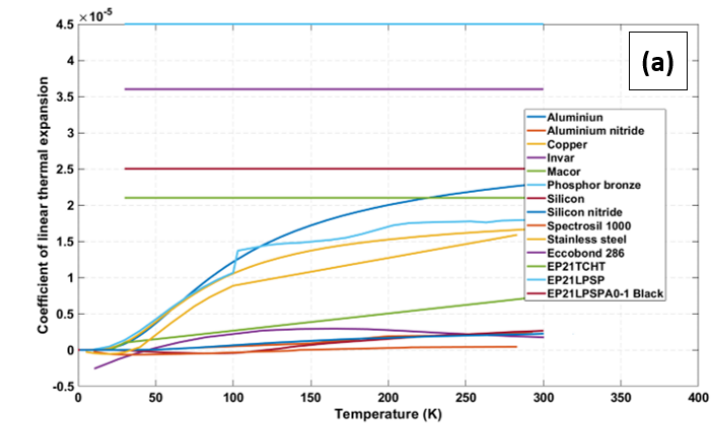

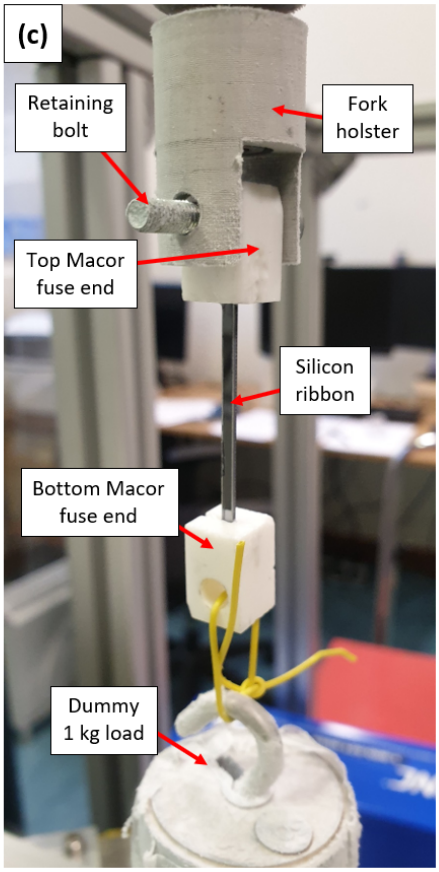

Figure 4.26: (a) Coefficient of thermal expansion for a variety of materials and epoxies. (b) Corresponding total contraction from 4–293 K for the materials and epoxies listed in (a). (c) New Macor fuse end design showing survival after loading with 1 kg mass and cooling in liquid nitrogen.

Finding CTE data, or indeed thermal conductivity data, for Eccobond 286 for cryogenic temperatures was not possible and according to the manufacturer is not available below 220 K

[355] [356]. This was unusual, considering the epoxy is manufactured and sold for use in cryogenic windows at 4 K [352] as well as being used successfully in previous cryogenic works [357] [358]. The stated CTE for Eccobond 286 is 36 ppm/°C and had to be assumed this was the case at all temperatures [356]. Similarly, the thermal conductivity was stated as $1\,\mathrm{W\,m^{-1}\,K}$ and had to be assumed constant over all temperature ranges [356]. For similar epoxies, data shows both the CTE and thermal conductivity reduce by approximately 1 % and 10 % respectively, from 300 K down to 10 K, providing conservative estimates for Eccobond 286 from these assumptions [359] [360]. From 4–293 K, this provides a contraction length of 2.4 µm.

While these length changes seem small in comparison to the 225 µm gap between the silicon ribbon and aluminium fuse end slot, this was filled by the epoxy. Eccobond did appear to set particularly hard and the cracking in the softer set Araldite 2012 indicated thermal stresses were occurring within the fuse end structure. However, suspension 4 managed to hang which pointed to the fact that these stresses were not always enough to fully break the silicon-fuse end interface, with the hard epoxy meniscus being able to adhere to silicon ribbon. The fact silicon contracts, expands and then contracts again will however exacerbate the stresses applied across the silicon-fuse end interface which is also a likely contributor to these failures along with material CTE mismatch.

To best increase the chance of suspension survival, a series of other cryogenic epoxies and fuse end materials were investigated. The results of this investigation are shown in figures 4.26(a) and (b). It was decided that work would continue with Eccobond 286 but the ceramic material, Macor, would replace the aluminium fuse ends. This was selected over similar CTE matching materials to silicon, such as more silicon or silicon nitride fuse ends due to the ability to easily manufacture and shape into a fuse end design.

The fuse ends were cut from 10 mm × 10 mm Macor bar with the same 1 mm wide slot cut to insert the ribbon. A series of 45 mm long silicon ribbons were hung using the Macor fuse ends with Eccobond 286 epoxy, as shown in figure 4.26(c). These were loaded with a 1 kg mass and submerged in liquid nitrogen. These particular ribbons were not lapped or polished. All of these samples survived both loading and cooling which proved promising.

4.4.8 Suspension 7

Suspension 7 employed these new Macor fuse ends with the same Eccobond 286 epoxy. This time a new silicon ribbon offcut was used that was 1.76 mm wide. This increases the stress in the suspension back to 7.2 MPa, 12 % of the breaking stress of median *control new* sample set from chapter 3. This ribbon was not polished at either end due to the Macor fuse end investigation showing this was not necessary.

As expected, upon initial evacuation of the cryostat, similar pressure issues experienced with suspension 5 and 6 arose. Cleaning the OVC rubber o-ring this time did not appear to improve the vacuum and further vacuum leak checks did not yield any conclusive results. It was

decided to cool the experiment again with the OVC pressure at 1.1 mbar and the IVC pressure at 1.6×10^{-4} mbar in order to compare the performance to suspension 6.

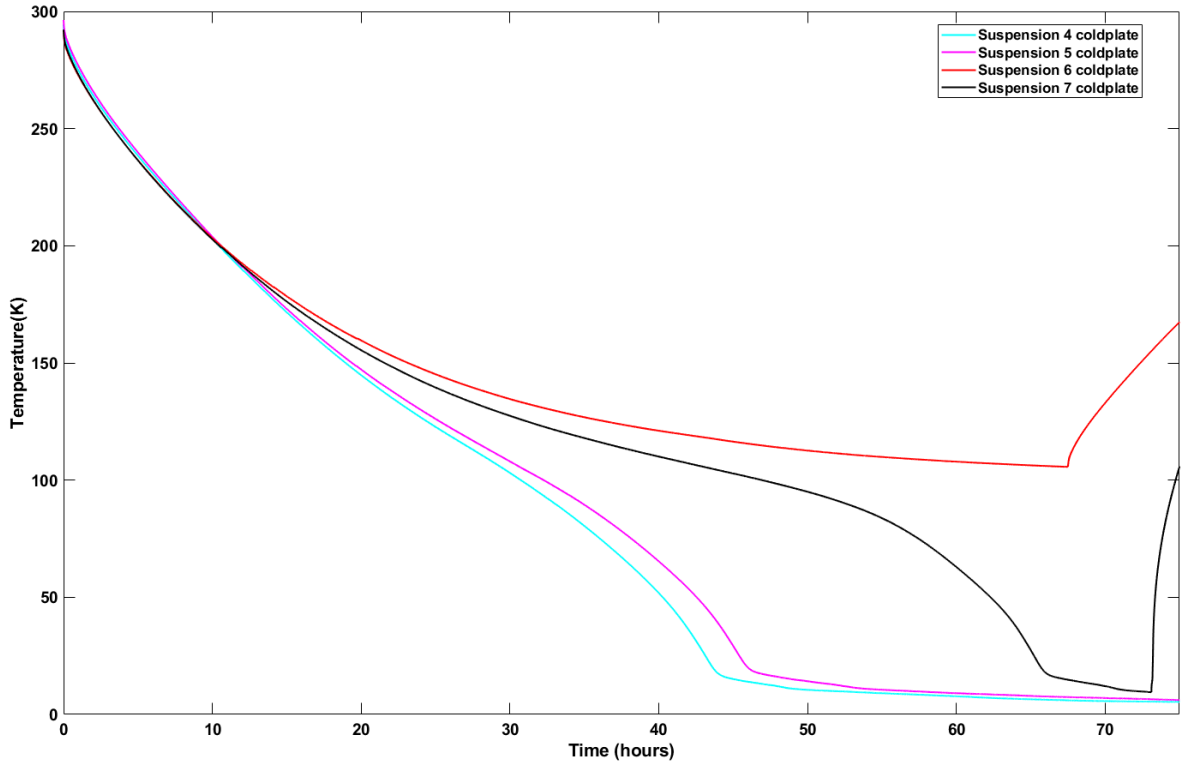

Figure 4.27: Suspensions 4 to 7 coldplate temperatures compared during cryostat cooling. Increases in temperatures observed for suspensions 6 and 7 are when cooling was stopped.

Figure 4.27 shows a comparison of the coldplate temperatures between suspensions 4 to 7, where a clear degradation in cooling is seen.

Suspension 5, after cleaning the OVC o-ring, had very similar OVC and IVC pressures to suspension 4 but appeared to have a lower cooling rate. This was potentially the first sign of cooling issues arising within the cryostat. Furthermore, the cooling rate of suspension 7 appeared to improve compared to suspension 6 with no changes made to the vacuum or cooling system.

While there was clearly a vacuum issue, this also provided strong evidence towards the fact there may be a cooling issue, related to the cryogenic compressor or the heat-exchanger.

With a hanging suspension inside and cooling, it was decided to continue the experiment. After approximately 63 hours, the suspension failed at the top fuse end. After investigation, the exact cause of failure at the top fuse end interface was unable to be determined.

The primary focus now shifted to fixing the vacuum system. This would allow for the investigation of the cooling system once this was fixed. Further vacuum investigations finally found the tip seals on the scroll pump were worn resulting in a poor vacuum provided to both the OVC and IVC. These tip seals were replaced for suspension 8.

4.4.9 Suspension 8 and compressor faults

Suspension 8 used a new, full width, silicon ribbon, that was 3.51 mm wide, reducing the stress in the ribbon back to 3.6 MPa, 6 % of the breaking stress of the median *control new* sample set from chapter 3. Like suspension 7, the silicon was not lapped and polished at either end.

During routine electrical testing of the sensors, conducted before cooling each suspension, it was discovered that the top ribbon sensor, a calibrated DT-670 SD, had failed. It was decided to move the uncalibrated DT-670 SD sensor from the coldplate to become the top ribbon sensor. With the lab jack raised to unhang the mass and the mass stops locked in place, this sensor was replaced in situ. A spare Cernox sensor was placed onto the coldplate to replace the uncalibrated DT-670. While the uncalibrated DT-670 SD sensor is the least accurate sensor, it was the only sensor available at the time that could be mounted on the ribbon. Furthermore, this experiment would firstly focus on investigating the cryostat cooling issues and, if successful, would then aim to look at heating the suspension if it survived cooling.

With the new sensor arrangement, the cryostat was evacuated reaching an OVC pressure of 7.4×10^{-2} mbar and an IVC pressure of 1.4×10^{-4} mbar before cooling commenced. After 45 minutes of cooling the cryostat compressor stalled which causes the remote motor to stop operating and cooling to cease. The previous assumption of both vacuum *and* cooling problems was proven to be correct.

A series of troubleshooting was conducted with advice given by Leiden Cryogenics. No immediate solution was found. At this point, it is typical for a qualified engineer from Cryomech, who make the compressor, to come out and investigate the issue, however due to the COVID-19 pandemic this was not possible.

Over the course of the next 4 months the author conducted a series of investigations ranging from opening the remote motor assembly and re-wiring to re-charging the helium gas in the compressor. This line of work was supported by Cryomech remotely. After working backwards from the most likely source of failure, the cold head and remote motor assembly, the issue was eventually traced to a power distribution board that was intermittently failing causing the compressor to trip power and stall. Further faults were also uncovered during the course of this maintenance, including incorrect helium pressures within the compressor and worn wire insulation within the remote motor assembly. All of these issues were rectified.

Suspension 8 was kept suspended in the cryostat during this work and survived cooling to cryogenic thermal equilibrium, as shown in figure 4.28. It should be noted that the compressor issues were not fully resolved during this suspension experiment. After approximately 1416 hours, or 59 days, the compressor had to be switched off, highlighted in figure 4.28, and the experiment allowed to return to room temperature. This was to allow for the electrical investigations that later resolved all underlying compressor issues.

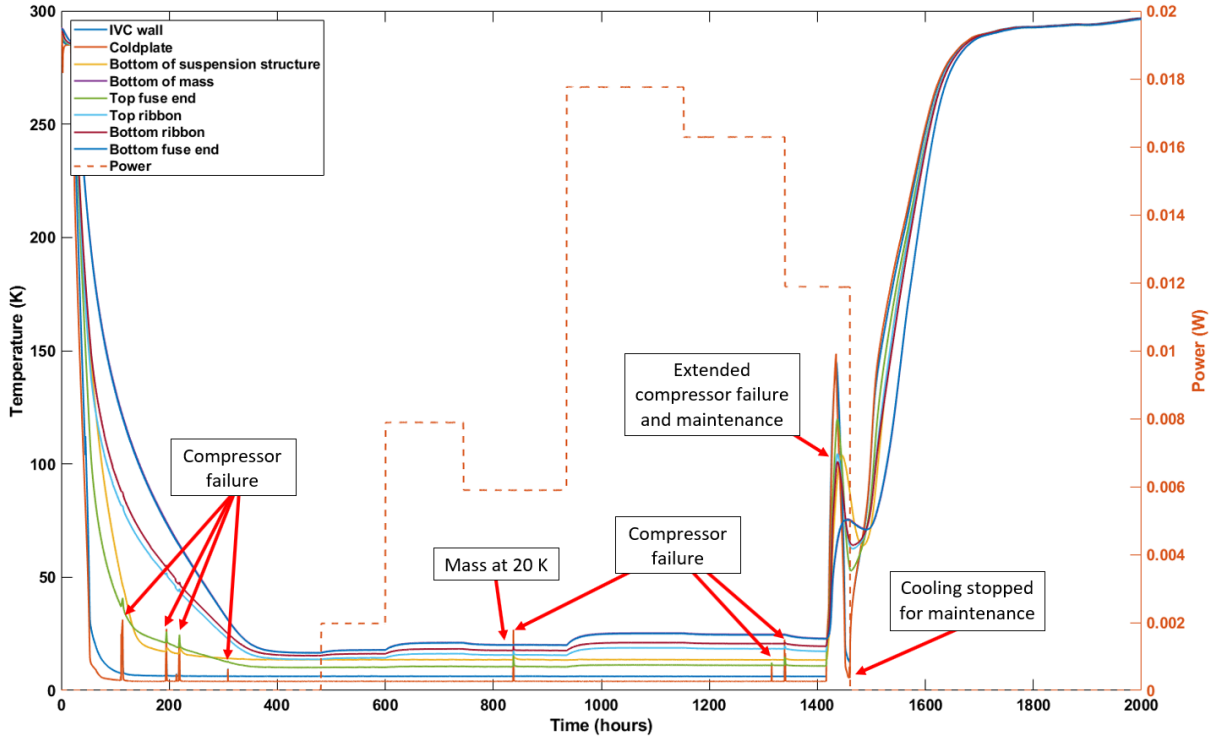

Figure 4.28: Full cooling curve for suspension 8 highlighting multiple points of compressor failure including extended compressor downtime for maintenance and investigations. Also highlighted is the mass stabilised at 20 K. The cryoswitch data is removed for clarity due to no failure of the suspension.

Compressor issues did not ultimately affect most of the results of suspension 8 allowing for thermal stabilisation of each heat step applied. In total, 6 heat steps were applied for this suspension, with the final heat step not stabilising before the final compressor failure and maintenance shown in figure 4.28. Heat step 3, highlighted on figure 4.28, shows the first demonstration of a loaded silicon suspension where the mass temperature is held at 20 K.

Recording and analysis of heat steps, ambient heat load and thermal conductivity for this suspension are analysed below in section 4.5.

As can be seen in figure 4.28, at cold thermal equilibrium a heat load still remains. It was suspected the 5 V channel as part of the cryoswitch sensing may have be inadvertently applying power, and hence heating, into the suspension structure through some path that was unable to be determined. To test this, the cryoswitch external wiring was disconnected from the NI box. This external wiring was left disconnected for 24 hours but no temperature change was observed throughout the suspension and so the cryoswitch was ruled out as a source of ambient heating.

4.4.10 Suspension 9

Due to the survival of suspension 8, the suspension 9 experiment is the same suspension undergoing cooling once the compressor was fixed. As such, no 24 hour hang in air was required.

A new Lake Shore 336 temperature controller was added to the existing 2 Lake Shore 340 temperature controllers. This allowed for the ambient temperature sensor to be reconnected since the capacity for temperature sensors was now increased from 8 to 12.

A new calibrated DT-670 SD replaced the uncalibrated DT-670 SD top ribbon sensor in order to increase accuracy. The uncalibrated DT-670 SD was moved onto the cold plate next to the Cernox sensor. This would allow for post-calibration of the sensor for suspension 8 data analysis, due to it being placed directly next to a more accurate Cernox sensor for this run.

The other significant change between suspensions 8 and 9 was the blanking of all port windows. All sapphire port windows were removed and replaced with blank metal inserts. The reliability of the cryogenic switch meant that visual confirmation of suspension failures was no longer required and it was hoped that this would help reduce any ambient heat load that may pass through the sapphire windows.

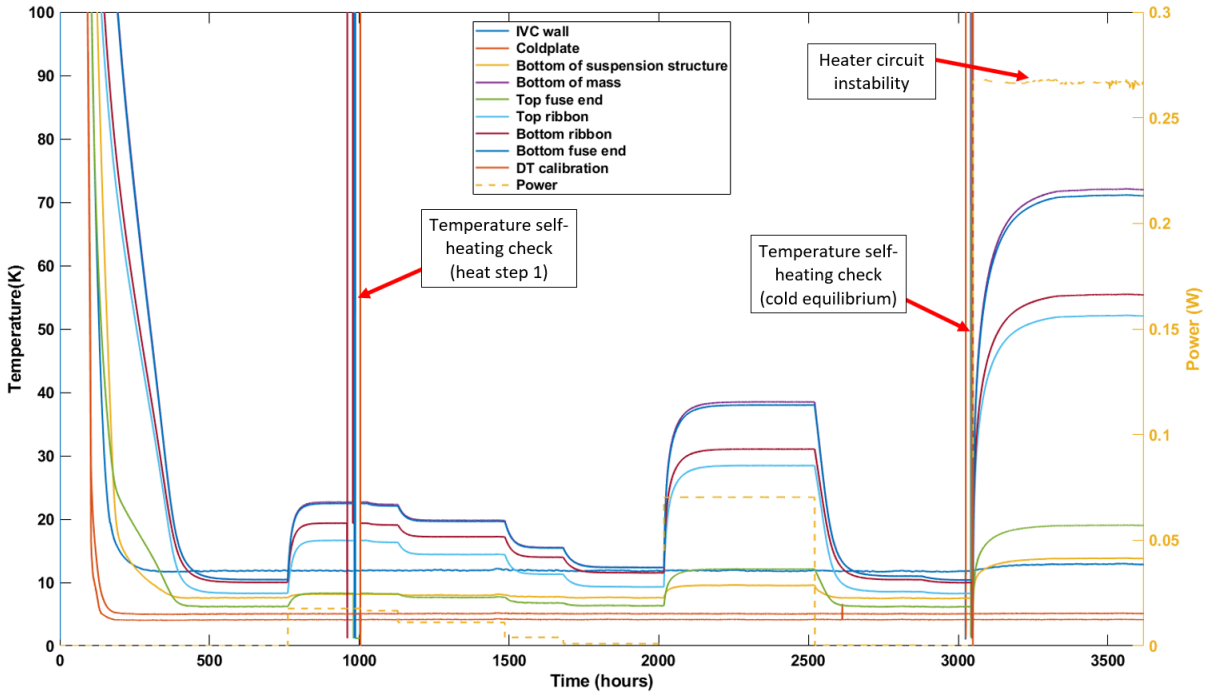

Figure 4.29: Full cooling curve for suspension 9 highlighting sensor self-heating checks and heater circuit instability. The room temperature sensor is not included due to the y-axis scaling and the cryoswitch data is removed for clarity due to no failure of the suspension.

Figure 4.29 shows the successful cooling curve of suspension 9. In total, 8 heat steps were applied for this suspension with one additional step returning the suspension back to cold equilibrium. Recording and analysis of heat steps, ambient heat load and thermal conductivity for this suspension are analysed below in section 4.5.

It is of note here that the IVC wall sensor stabilised at 11.9 K, whereas for suspension 8 it stabilised at 6.3 K. This was assumed, and later confirmed, to arise from the fact the sensor had detached from the IVC wall.

A series of sensor self-heating experiments were carried out for heat step 1 and later at cold thermal equilibrium. During heat step 1, the 5 temperature sensors attached directly to the suspension were unplugged; the top fuse end, top ribbon, bottom ribbon, bottom fuse end and bottom of mass sensors. These were left disconnected overnight. When the sensors were plugged back in individually to look for temperature responses, little change was seen on the Cernox sensors (top fuse end, bottom fused end and bottom of mass). However, when the top and bottom ribbon DT-670 sensors were plugged back in, a response of 14 mK over approximately 12 minutes was observed in each sensor. This is discussed further in section 4.5.4.

This experiment was repeated at cold thermal equilibrium where instead of unplugging all 5 sensors at once, the sensors were unplugged one by one and then re-connected. The results were consistent with the experiment conducted at heat step 1.

The final temperature step applied produced some power instability in the heater circuit, as shown in figure 4.29. The source of this instability was investigated and rectified as part of the next suspension experiment, as discussed in section 4.4.11.3.

After this temperature step, cooling was stopped and the system brought back to room temperature and pressure. With the cooling stopped by turning off the compressor, an unusual noise was heard from the Leybold Turbovac SL80 turbo pump that was not previously heard due to the compressor noise. It was discovered the turbo pump bearings had significantly degraded. Before the next suspension was installed, a new Leybold Turbovac 90i turbo pump was installed.

Suspension 9 was removed from the cryostat and the silicon ribbon-fuse end assembly was placed into storage for any future work required. This experiment allowed for the progression to installing a HC bonded silicon suspension into the cryostat as discussed in the next section.

4.4.11 Suspension 10 - a cryogenic, HC bonded, silicon suspension

With the success of suspension 8 and 9, work was conducted in parallel to prepare a HC bonded silicon suspension.

The fuse end designs were to be replaced with silicon blocks, where a silicon ribbon could then be HC bonded onto the sides of the blocks. The top block could be clamped to the suspension structure and the mass could then be clamped to the bottom block to load the suspension.

4.4.11.1 HC bonded suspension assembly and jig

Silicon blocks were acquired from Sil'tronix [361]. These blocks measured $5\times10\times20$ mm. The four largest faces were required to have a peak-to-valley flatness of <63 nm (equivalent to a $\frac{\lambda}{10}$ flatness at 633 nm) over 95 % of the surface. For the two largest surfaces this was to allow for HC bonding of the silicon ribbon onto the side of the block. For the two smaller areas, this was to allow for clamping the silicon block into a suspension support structure or to clamp the mass at the bottom but would also enable future designs to be HC bonded at these interfaces

rather than clamped. Furthermore no peaks could protrude from the edges so the blocks were chamfered on all vertices. Both of these are requirements for HC bonding [186] [191]. The silicon block samples were characterised on a Zygo GPI XP-D flatness interferometer upon arrival to ensure these standards were met before bonding [362]. These silicon blocks effectively become analogous to the aLIGO suspension ears described in section 1.5.4 of chapter 1.

The silicon ribbons used were laser cut 3.5 mm wide by 775 µm thick and 250 mm long acquired from LML using the same new laser cutting technique as in chapter 3. The ribbons were also characterised on the Zygo flatness interferometer beforehand to ensure no peaks from laser cutting would inhibit the HC bonding process.

Silicon can not be HC bonded without a sufficient oxide layer [190]. While these oxide layers should be thermally grown to a minimum of 50 nm for reliably strong bonds, they more commonly been previously grown to 100–200 nm thick [186] [191] [363]. Taking the bond thickness as 61 nm and a conservative oxide thickness of 150 nm for reliable strength this gives total bond thicknesses of 211 nm if only one silicon piece is oxidised, and 361 nm if both pieces are oxidised [150] [364].

Before bonding, the ribbon and silicon blocks were cleaned in Piranha solution. The 3 pairs of silicon blocks were then thermally oxidised in an oven for 5 hours at 1000 °C. This produces an oxide layer approximately 150 nm thick and was confirmed by post-analysis on a SE850 spectroscopic ellipsometer [365]. The silicon ribbons were not oxidised to minimise handling and maintain their ultimate tensile strength.

As shown in figure 4.30, a specialised jig was designed and constructed to allow for the HC bonding of the ribbon and blocks. The jig must not allow for contact on regions that are to be bonded in order to maintain the high cleanliness required for a successful HC bond. As shown from the work in chapter 3, the silicon ribbon must also not be contacted anywhere along its length to avoid surface damage in order maintain its tensile strength.

Figure 4.30(a) shows the original design drawing for the suspension assembly. Recesses were built to accommodate both silicon blocks into two moveable stages. When fully extended, these moveable stages, together with bond aligners, create a clearance of 0.5 mm at either end between the end of ribbon and the exposed edge of the silicon block. This is to avoid any unwanted contact, and induced stress, between the ends of the ribbon and clamping surfaces on the suspension support structure or mass. The blocks sit above the moveable stages to ensure the ribbon does not contact them. The bond aligners are referenced off the moveable stages, with one bond aligner controlling the longitudinal placement of the ribbon and both aligners ensuring central bonding alignment. The final built design is shown in figure 4.30(b) with a few modifications made to (a). In order to ensure best alignment of the moveable stages and minimise any off-axis stress on the ribbon, it was decided to reference and fix one stage to a breadboard and allow only the other stage to move. Furthermore, securing screws were threaded through the side of the stages to gently press the silicon blocks into the sides of the referenced

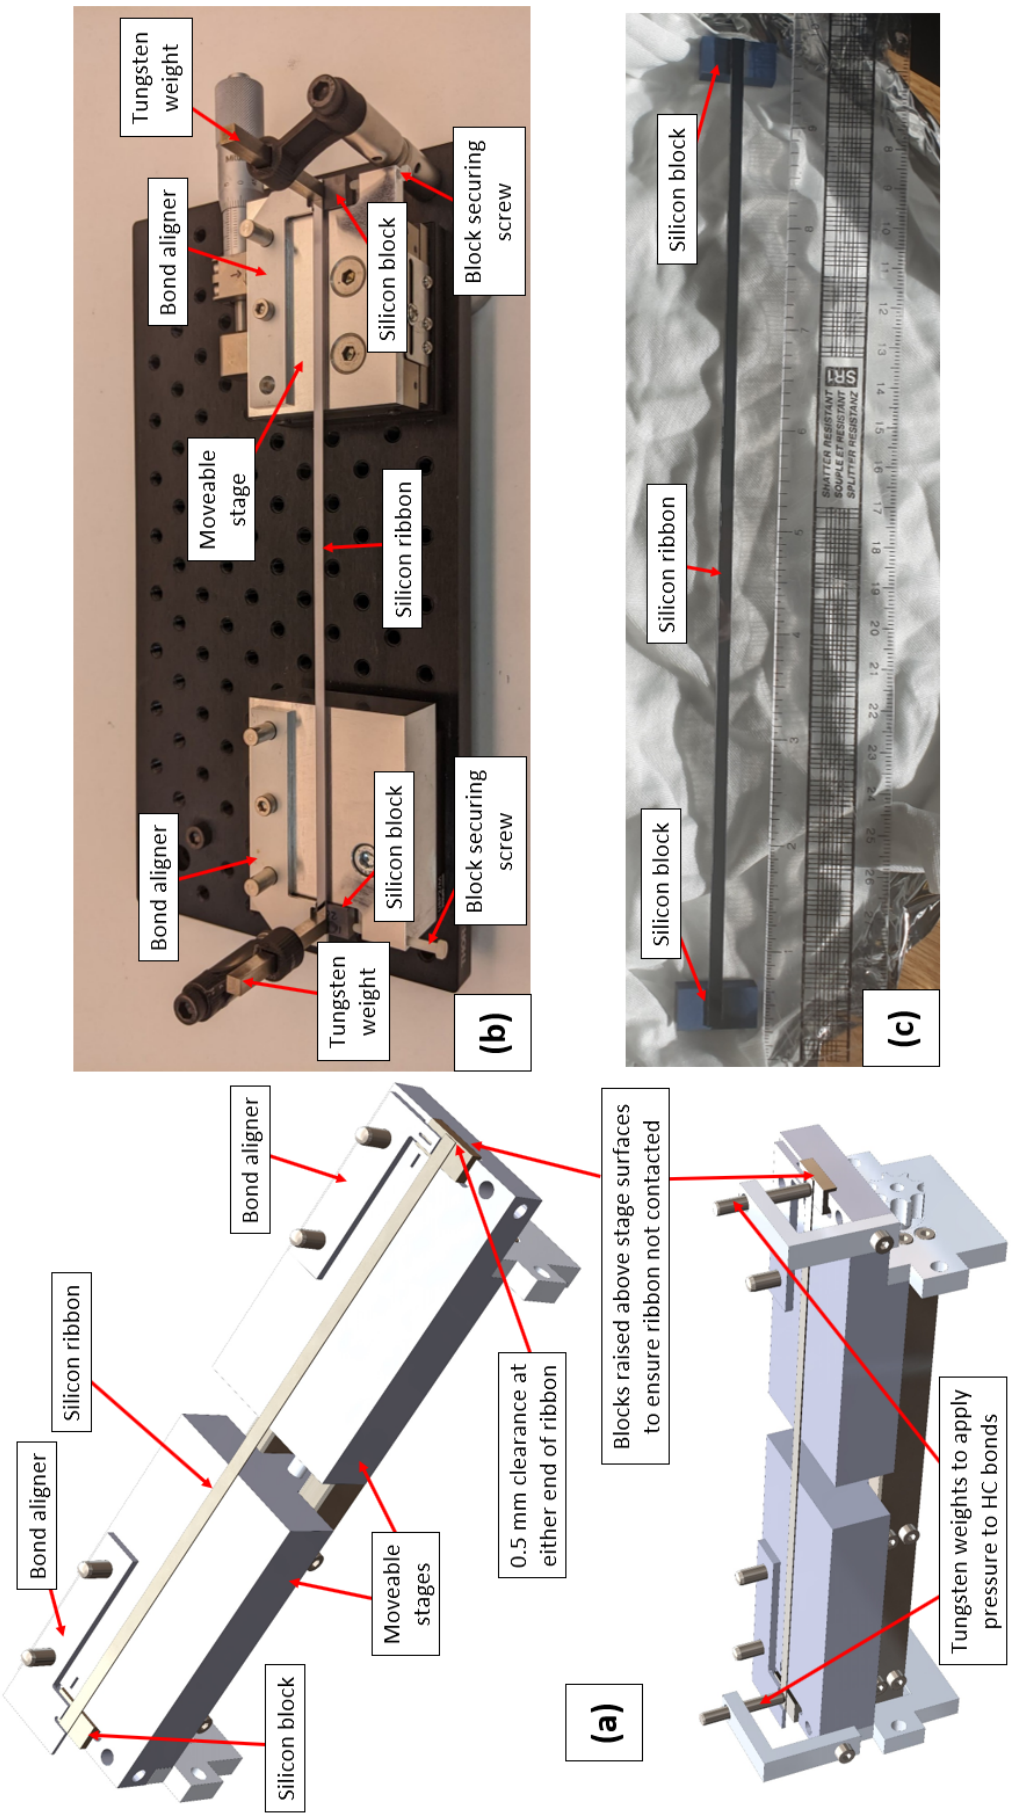

Figure 4.30: (a) Original design drawings for suspension HC bonding assembly jig. (b) Actual HC suspension in HC bonding assembly jig with minor modifications from (a). (c) Successfully HC bonded silicon suspension. CAD drawings supplied by Mr Russell Jones.

recesses.

The 3 pairs of blocks were then cleaned again in Piranha solution prior to bonding. With the silicon blocks and bond aligners in place and the moveable stages fully extended, $0.8\,\mu L\,cm^{-2}$ of sodium silicate bonding solution was applied to the silicon block surfaces. This solution is a 1:6 sodium silicate to DI water mixture as defined in the aLIGO bonding procedures [366]. The ribbon was then carefully placed on top of the bonding solution and pushed up against the bond aligners. Tungsten weights were then carefully inserted to apply a small pressure to the bond region to prevent the ribbon from drifting. The bond was left in the jig for 24 hours. After this, the tungsten weights were removed, the moveable stages were retracted slightly and the bond aligners were removed without disturbing the suspension. The suspension assembly was removed to ensure bonding had occurred, producing the final HC bonded silicon suspension, as shown in figure 4.30(c). The suspension assemblies were left to further cure, out of the jig, for >4 weeks, in line with standard HC bonding procedures [186].

Combined with a 61 nm HC bond, this provides a total bond thickness at the top and bottom of the ribbon of 211 nm, with a bond area of 9.5×3.5 mm. This results in an induced stress in the bond, due to loading from the 1 kg mass, of 0.3 MPa. The accepted breaking stress for HC bonds between silicon substrates is 30 ± 17 MPa, giving a safety factor of 51, taking the conservative silicon HC bond strength of 13 MPa [186]. For comparison, aLIGO's bond stress is 0.16 MPa with a safety factor of 94 [189].

A total of 4 suspensions were bonded in this way, with 3 of the suspensions bonded using oxidised blocks, providing the strongest bond and 1 suspension bonded using only the native oxides of the silicon block and ribbons. As will be discussed in chapter 5, FEA was conducted on these 4 bonded suspensions to determine the best-aligned suspension which would yield the lowest stress and hence best chance of survival.

4.4.11.2 Bonded suspension support structure and hanging in air

Due to the unique design of the HC bonded suspension, it was necessary to design a new support structure and interface. The overall support structure is based on a similar design to the original support structure shown in section 4.3, with top, mid and bottom plates as before.

Since the suspension already has top and bottom block attachment pieces as part of the overall structure, the suspension can not be simply inserted as with the previous fuse end design. A new design of interface between the suspension structure and silicon blocks must be created. This involved the use of a simple clamping bar with bolts either side, as shown in figure 4.31(a). This design has been used extensively for research involving clamping silicon cantilevers for cryogenic mechanical loss measurements and was adapted for this application [193]. In order to evenly and tightly clamp without damaging the silicon, it is necessary for all clamping surfaces to be mechanically polished to a specular standard, including the top of the aluminium mass.

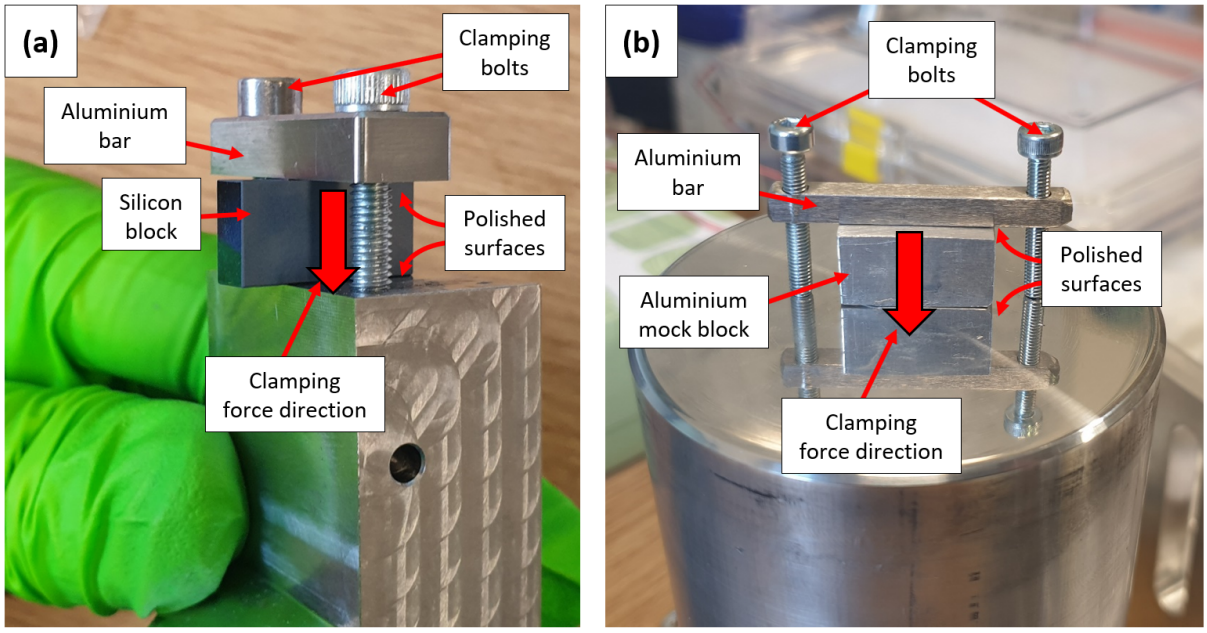

Figure 4.31: (a) Example of clamping technique applied to silicon suspension. (b) The same clamping technique applied to an aluminium mass, using a mock aluminium block.

This same concept was applied for the silicon suspension, as shown in figure 4.31(b). Due to the 20 mm width of the block, this required the bolts to be 30 mm apart, centre to centre. As such, the stainless steel mass previously used was replaced with an equivalent aluminium mass, measuring 65 mm in diameter and 120 mm long. The aluminium bar has a 4×4 mm cross-section. This allows for a 1 mm gap between the edge of the bar and the silicon ribbon bonded onto the side of the block, as the silicon block is 5 mm thick, avoiding any potential contact that could damage the silicon.

A bottom block alignment plate held on with a g-clamp, shown in figure 4.32(a), was created to ensure the silicon block was well aligned on the top of the aluminium mass to minimise any suspension stresses. This is removed once the suspension assembly is affixed to the top block.

Instead of having the mass placed inside the suspension support structure and then installing the suspension above it, it was deemed most practical to affix the suspension onto the mass first and then move the suspension and mass assembly into the suspension support structure. Furthermore, a new cryoswitch and mass catcher, working on the same principles as before were developed. The cryoswitch differs by having a copper trigger arm attached to the side of the mass instead of beneath it. The mass heater was bolted directly onto the bottom of the mass, as before, but this new procedure would require one to balance the suspension before hanging which is not possible. To mitigate this, a new mass catcher was designed with a recess that accommodated the heater and its wiring while being able to support the mass on a level plane. The bottom of mass Cernox sensor was also fitted to the bottom of the mass. All of the above modifications can be seen in figure 4.32(a).

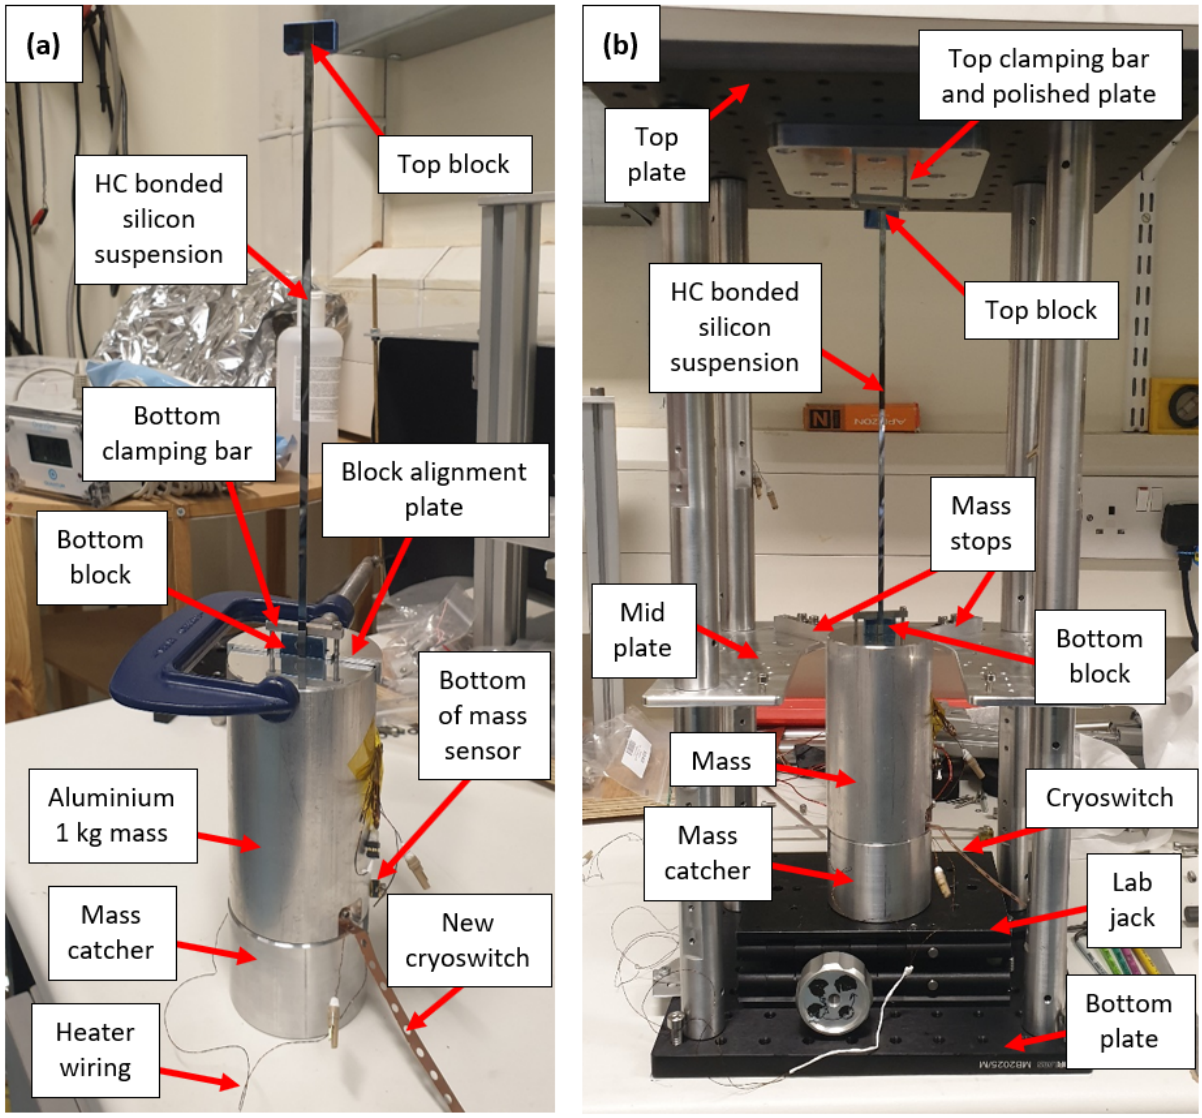

Figure 4.32: (a) Silicon suspension assembly affixed to the top of the mass, sitting on the mass catcher. (b) Suspension assembly from (a), with g-clamp and alignment plate removed, inserted into the suspension support structure.

The new suspension structure can be seen in figure 4.32(b). This new structure was almost identical to the original epoxied suspension's structure, except the supporting 1" diameter poles are made from 6061-T6 aluminium instead of stainless steel. This allowed for easy machining of flats in the poles to create new mounting points for thermal bobbins that provide thermal and mechanical anchoring of sensor wires to the suspension structure support poles. Further to this, a thin layer of Apiezon N thermal grease was applied between all poles and connecting plates, unlike the epoxied suspension's structure.

One notes this new structure is missing a front portion of the middle plate. This is to allow for the insertion of the suspension assembly in figure 4.32(a) into the suspension structure and onto the lab jack. One can see in figure 4.32(b) the next step is to elevate the suspension in order to clamp the top block. This is a particularly delicate procedure as over-elevation of the

suspension assembly will compress the suspension and break the bonds or ribbon.

For the top silicon block, this clamping technique can be applied in reverse, replacing the mass with a polished metal block with the clamping bar underneath instead of on the top. No alignment jig was required for the top silicon block.

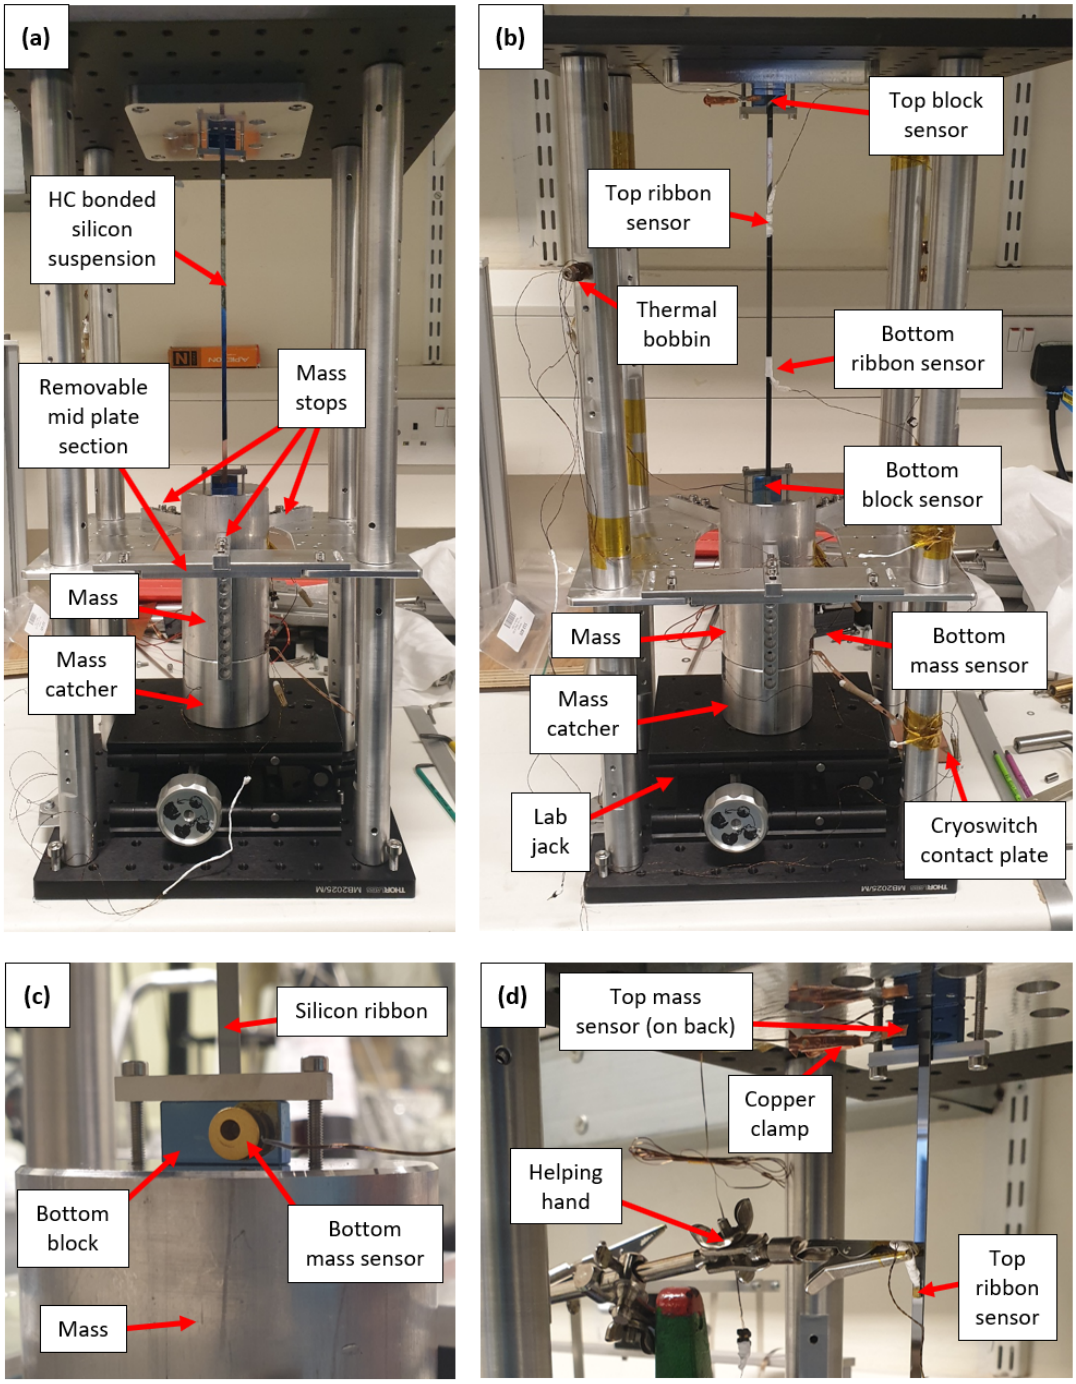

Figure 4.33: (a) Elevated suspension clamped at the top with mid plate inserted. (b) The last step of the suspension assembly, with temperature sensors attached ready to be hung. (c) Cernox sensor on bottom block. (d) Clamp for Cernox sensors and top ribbon sensor attached with cryogenic grease before wrapping in PTFE tape.

Figure 4.33(a) shows the suspension after elevation with the top block clamped to the top polished plate and the removable mid plate section inserted. This completes the 3 mass stop configuration. The mass stops have multiple threads to allow for the insertion of nylon-tipped grub screws that can be driven into the side of the mass to fix it in place.

As shown in figure 4.33(b), the remaining temperature sensors were attached; the top block, top ribbon, bottom ribbon and bottom block sensors. The top and bottom block sensors are the same Cernox sensors previously attached to the top and bottom fuse ends on the epoxied silicon suspensions. The sensors were affixed by a thin layer of cryogenic varnish directly onto the polished face opposite to the bonded ribbon, as shown in figure 4.33(c). Copper clamps were created to hold the sensors in place as the varnish cured, and one can be seen still attached to the top block sensor in figures 4.33(b) and (d) as it was curing. These are removed before suspension hanging and cooling. The top and bottom ribbon sensors were affixed to the ribbon as before, using cryogenic grease, where helping hands mounted on lab jacks were used to support the sensor in place before wrapping, as shown in figure 4.33(d). All lead wiring was affixed to the suspension structure poles via Kapton tape where necessary before being attached to thermal bobbins. These thermal bobbins are now directly screwed into the suspension structure poles, for both mechanical and thermal anchoring, an example of which is shown in figure 4.33(b). Finally, a copper plate, similar to the one attached to the original mass catcher in figure 4.19, was attached to the bottom plate of the suspension structure. This was connected to the same 5 V circuit through the same National Instruments USB-6211 multifunction device, which recorded the hanging state of the suspension in the same way as previous epoxied suspensions.

An initial attempt at hanging in air was conducted using the suspension bonded with only the native oxides of the silicon block and ribbons, since, from section 2.3.3.4 of chapter 2, this was likely to be the weakest bond. This suspension failed at the bond interfaces when attempting the first step shown in figure 4.32(a) as the suspension assembly was inserted into the bottom clamp. This was due to over-constraint of the suspension assembly when handling both the top and bottom blocks at the same time.

A second attempt was made using one of the remaining 3 suspension assemblies, where the blocks were oxidised, likely producing stronger bonds. This attempt was conducted before the cryoswitch, mass heater or any temperature sensors were attached in order to discover any assembly conflicts or unforeseen difficulties.

What is believed to be the world's first successful hang of a HC bonded silicon suspension is shown in figure 4.34. This suspension was left to hang in air for over a month with no failure occurring.

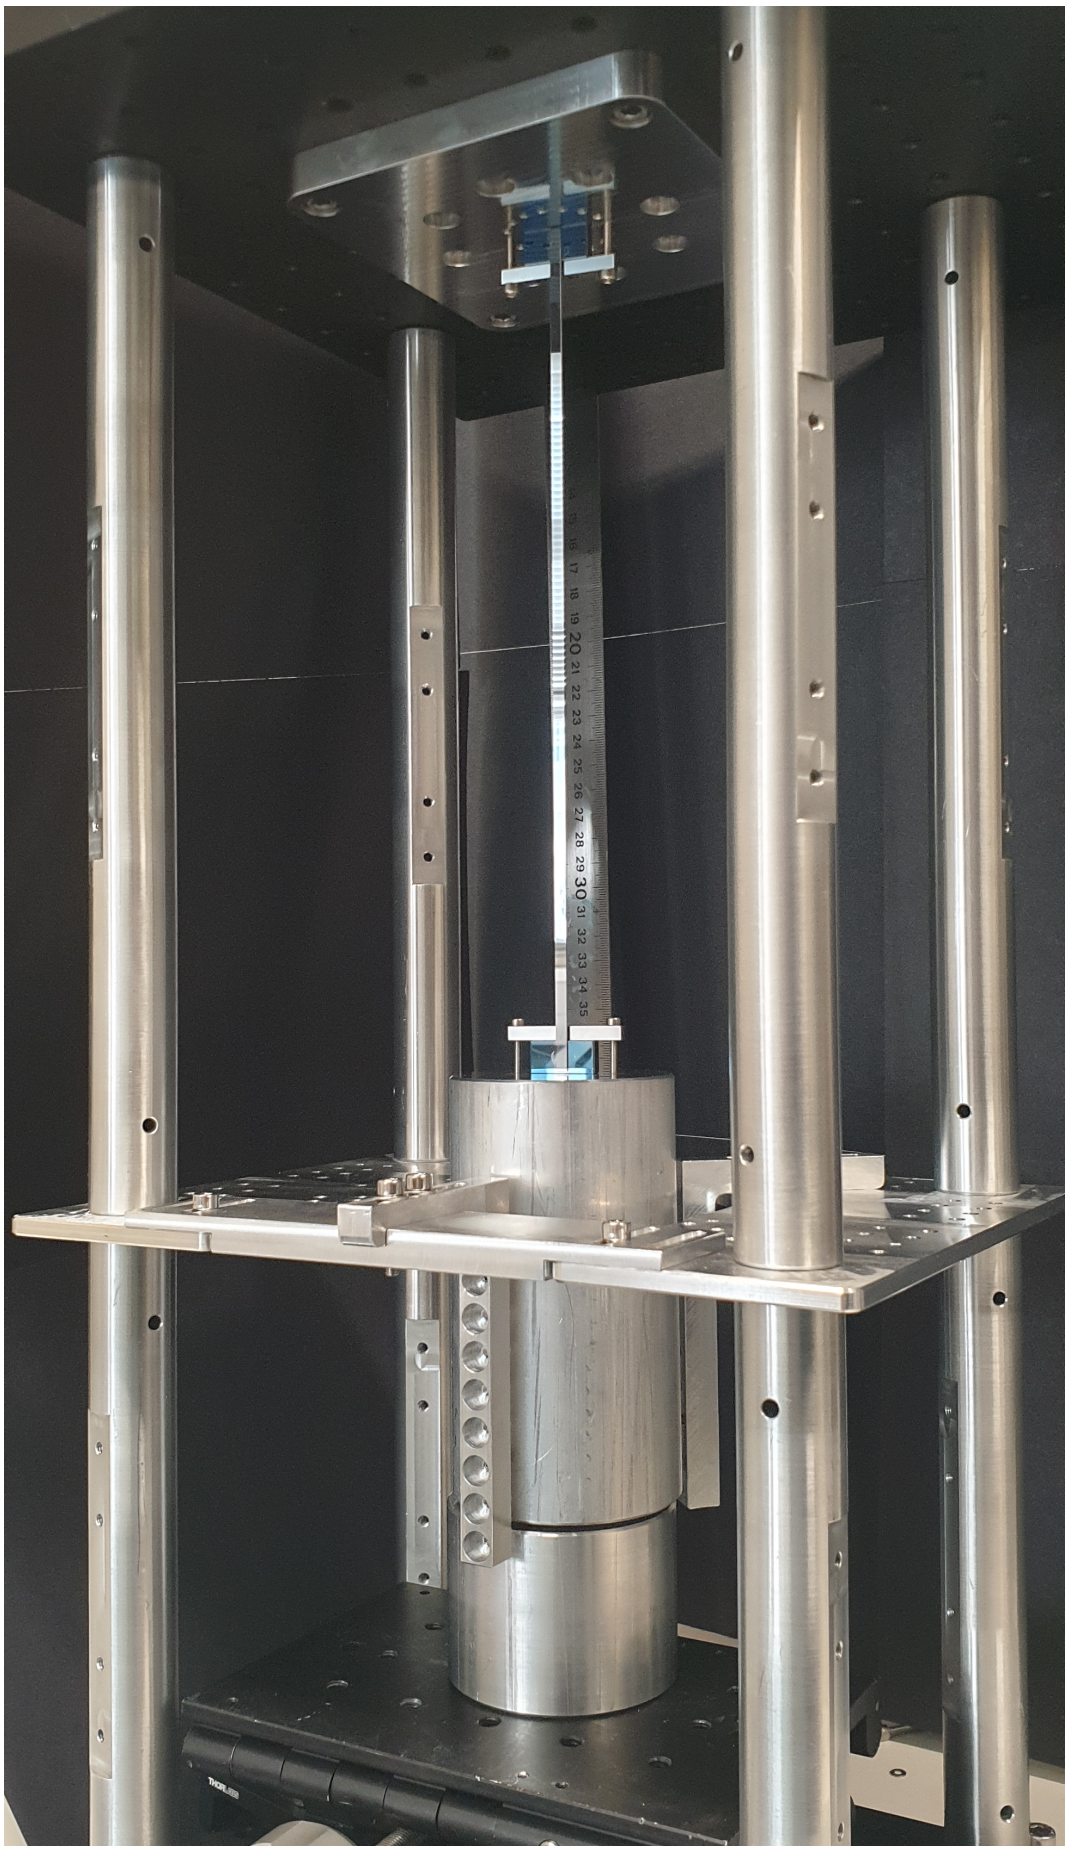

Figure 4.34: The world's first successful hang of a HC bonded silicon suspension.

4.4.11.3 Bonded suspension cooling

The significant changes between the epoxied and HC bonded suspension infrastructure were; the new aluminium mass, mass heater, new mass catcher, cryoswitch, suspension structure and turbo pump.

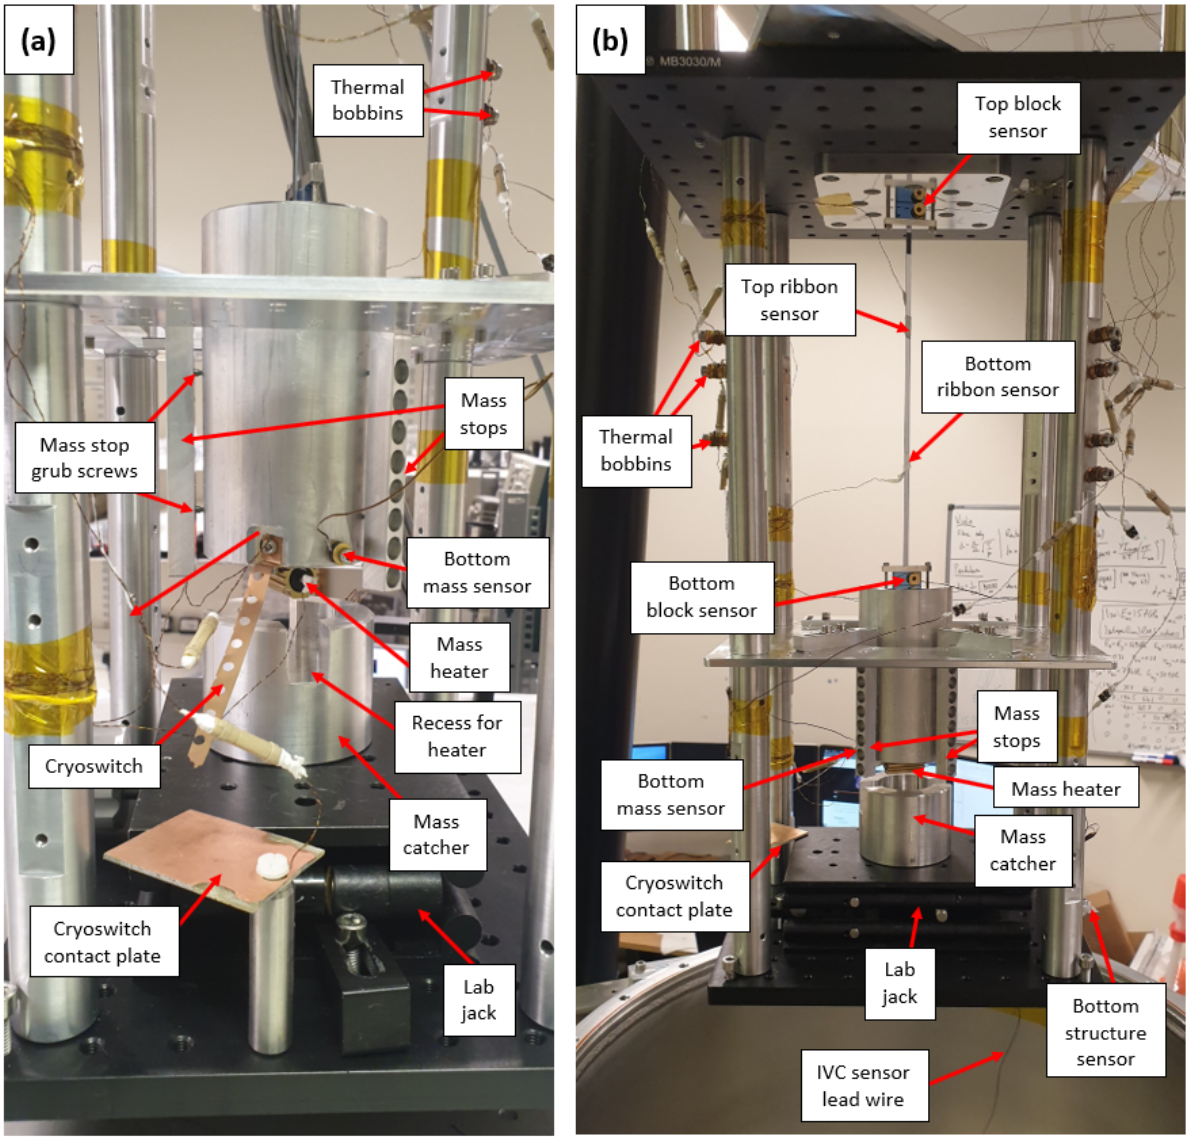

Figure 4.35: (a) Zoomed in image of mass area on HC bonded suspension. (b) HC bonded suspension installed in the cryostat, viewed from the back to show the positioning of temperature sensors on the suspension.

Figure 4.35(a) zooms in on the mass, mass stop and cryoswitch area to clearly show the new configuration with the mass hanging but locked in place by the nylon-tipped grub screws.

All temperature sensors outside the suspension structure were kept in the same place. The Cernox temperature sensor on the bottom of the old suspension structure was placed on the bottom plate of the new suspension structure. The top block, top ribbon, bottom ribbon, bottom block and bottom of mass temperature sensors were placed as discussed in section 4.4.11.2.

The temperature sensors attached to the suspension assembly and structure are shown in figure 4.35(b).

With the grub screws retracted, the mass is freely hanging. Figure 4.35(b) shows the final state of the HC bonded silicon suspension before the cryostat was closed up, evacuated and cooled.

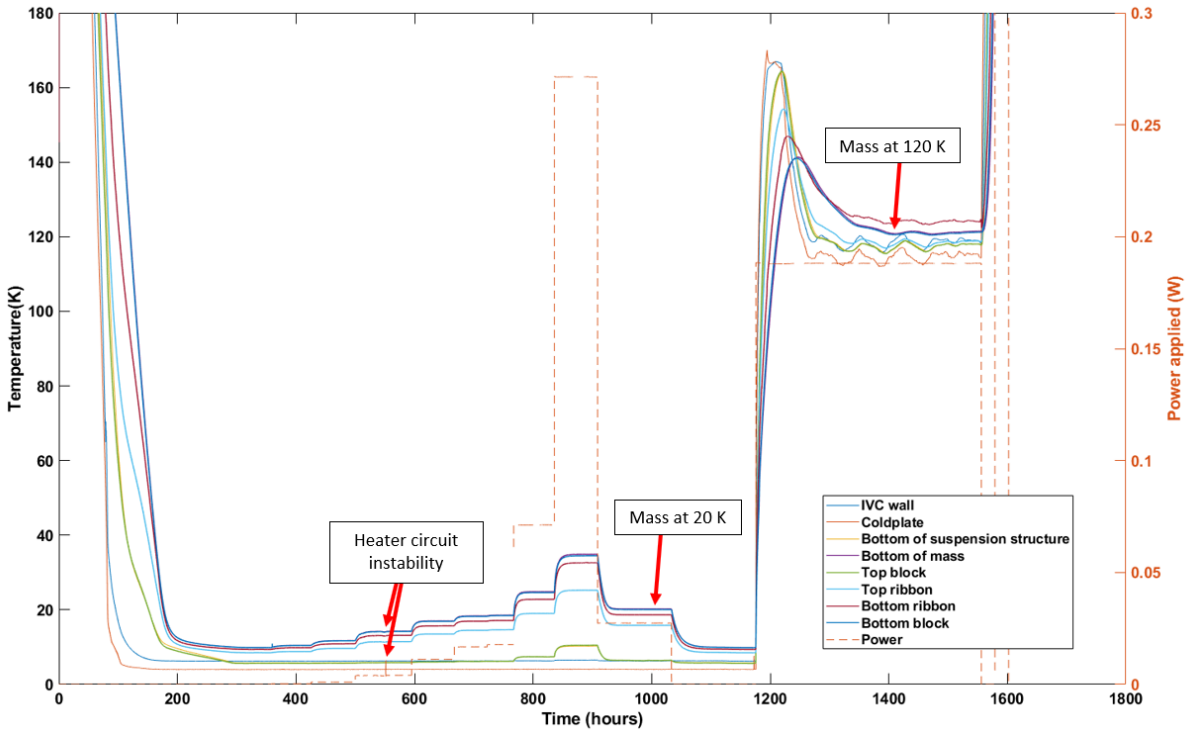

Figure 4.36: Full cooling curve for suspension 10 highlighting and heater circuit instability and the mass stabilised at 20 K. The mass is also run around 120 K although thermal stabilisation did not occur. The room temperature sensor is not included due to the y-axis scaling and the cryoswitch data is removed for clarity due to no failure of the suspension.

Figure 4.36 shows the successful cooling of suspension 10, a HC bonded silicon suspension. A total of 10 heat steps were applied, 9 of these applied the same power as the heat steps for suspension 9, with one additional step returning the suspension back to cold equilibrium.

Before applying the final heat step 10, heating was switched off to allow the suspension to reach cold thermal equilibrium again. This was to investigate the source of ambient heat that is still present on all suspensions. For suspension 8, the heater wire circuit was disconnected from the 5 V channel to ensure no excess current was flowing through the circuit or was being picked up from external electromagnetic fields. However, this investigation did not fundamentally disconnect the thick copper heater wires that travel from the high voltage vacuum feedthrough, shown in figure 4.22(a) down through baffling and into the cryostat IVC. It was suspected these were potentially introducing excess heat into the system. To investigate this, the wires from the heater circuit were disconnected externally and a heat gun was applied to the outside of the high voltage vacuum feedthrough with the intention of heating the thick copper wires. This was only

conducted for 30 minutes and no observable change in temperature was observed. This was likely not left long enough to observe a change. It should also be noted however, these thick copper wires terminate near the cryostat platform before connecting through LEMO connectors to cryogenic wiring that is then thermally anchored to the cold suspension structure poles before connecting to the mass heater, as can be seen in figure 4.22(b). For this reason, it is not believed that this is the main contributing source to the ambient heat load, however future work should rule this out by entirely removing the thick copper wires from the system.

The final heat step applied aimed to stabilise the temperature of the mass at 120 K, thus demonstrating the low and high temperature operating regimes for 3G cryogenic silicon detectors. To attempt this, the Lake Shore power output to the heater on the bottom of the mass was set to 50 % (5 V). This would allow for fine adjustment up or down of the temperature. For coarse temperature adjustment, the variac on the heater unit, discussed in section 4.2.1.2, was initially set to 90 % power applied to the 4 K plate 100 W heater. This led to the temperature spike around 1200 hours on figure 4.36. As the temperature rose, a further series of adjustments were made eventually setting the variac to 50 % power applied to the 100 W heater. This, combined with 50 % applied to the heater on the bottom of the mass achieved a mass temperature around 120 K.

One can see a clear thermal instability throughout the suspension at this power setting. Firstly, the 100 W heaters on both the 4 K and 50 K plates are not designed to heat cryogenic experiments to stabilised temperatures, they are purely to bring any cryogenic experiment quickly back to room temperature. There is a sinusoidal pattern in the temperature data, growing stronger towards the top of the suspension sensors and the cold plate sensor. It is not clear the source of this sinusoid as the difference between the peaks around 1400 hours equates to approximately 80 hours. The sinusoid does not correlate with a 24 hour daily cycle and investigating historical weather records shows little change in daily temperatures over this time period. The source of this fluctuation can not be explained however the electronics of the heater circuit are simplistic and unshielded, unlike the cryogenic electronics, so there is potential for either pickup from unshielded cabling or interference between the cryogenic and non-cryogenic electronics. Furthermore, radiative effects are expected to be a greater factor at these higher temperatures, however it is unclear from this plot how exactly this affects the suspension at these temperatures, particularly when the strongest fluctuations are seen at the coldplate temperature sensor. Due to the fact this plate is constantly cooled, it is unlikely this sinusoid affect is as strongly affected by radiation compared to lower down, where the suspension is more exposed to the IVC wall temperature.

After this was achieved, cooling was stopped and both 100 W heaters were set to 100 % power. Once the suspension was back to room temperature it was left hanging in vacuum from May 2022 to August 2022. The system was re-pressurised slowly using the leak valve attached to the IVC in order to minimise suspension failure from injection of particles into the vacuum

system during the process. The cryostat IVC, 50 K shield and OVC were removed and the suspension locked up and supported at the end of August 2022. This demonstrates that the HC bonded silicon suspension has hung for over 8 months and survived a full cryogenic cycle.

The bottom ribbon sensor showed random spiking on the temperature channel. The spikes were instantaneous, temporary and the reading returned to the same value before this spike. This was potentially due to poor soldering or loose wiring. As such, the data for the bottom ribbon in figure 4.36 has been put through median filtering in MATLAB to remove these spikes.

It was noted during heat step 3 that instability on the heater circuit, witnessed in suspension 9, was still present and was affecting the temperature stability of the heat step. Investigations involved shielding the custom built through box and 110 Ω heater circuit resistor box by wrapping them in UHV foil. Both of these boxes sat next to each other on the instrument rack and it was suspected there was some form of electrical interference in this area. As can be seen on figure 4.36, this fixed the heater instability issue as temperature steps 4 to 10 provided stable power to the heater on the bottom of the mass.

This experiment has demonstrated what is believed to be the first successful hang, cooling and warm up cycle of a HC bonded silicon suspension. The suspension was cooled to cold thermal equilibrium, heat steps were applied including steps that stabilised the mass to 20 K and around 120 K, the proposed high and low temperature operating regimes for 3G cryogenic silicon detectors. Recording and analysis of heat steps, ambient heat load and thermal conductivity for this suspension are analysed in section 4.5.

4.4.12 Summary of cryogenic silicon suspension experiments

This section summarises the 10 cryogenic silicon suspension prototype experiments in table 4.3.

Suspension	Ribbon stress (MPa)	Ribbon width (mm)	Connection type	Result (Failure location)	Comments
1	7.7	1.65	Aluminium fuse end and Araldite 2012	Failed (unknown)	Failed during cooling
2	7.1	1.79	Aluminium fuse end and Eccobond 286	Failed (bottom fuse end)	Ribbon polished at ends, survived initial cooling, failed on warm up
3	7.1	1.79	Aluminium fuse end and Eccobond 286	Failed (unknown)	Same ribbon from suspension 2, failed during cooling.
4	7.7	1.65	Aluminium fuse end and Eccobond 286	Survived	Same ribbon from suspension 1, polished at ends, survived full cooling, heater short-circuit problems, unknown failure mechanism.
5	7.7	1.65	Aluminium fuse end and Eccobond 286	Failed (bottom fuse end)	Same ribbon from suspension 4, failed during cooling.
6	3.6	3.50	Aluminium fuse end and Eccobond 286	Failed (top fuse end)	Ribbon polished at ends, failed during cooling, vacuum issues.
7	7.2	1.76	Macor fuse end and Eccobond 286	Failed (top fuse end)	Failed during cooling, vacuum and cooling issues.
8	3.6	3.51	Macor fuse end and Eccobond 286	Survived	Survived cooling and warm up, cryogenic compressor failure.
9	3.6	3.51	Macor fuse end and Eccobond 286	Survived	Same ribbon from suspension 9, port windows blanked, survived cooling and warm up, turbo failure.
10	3.6	3.50	HC bonded ribbon to silicon blocks	Survived	First demonstration of cryogenically cooled, HC bonded silicon suspension.

Table 4.3: Summary of cryogenic silicon suspension prototypes.

4.5 Analysis of suspension 8, 9 and 10

This section conducts analysis on suspensions 8, 9 and 10, with the primary aim of determining the thermal conductivity of the silicon suspension.

As shown in section 3.4.1 of chapter 3, these 775 µm thick ribbons were cut from 300 mm diameter silicon wafers with boron (p-type) dopant and a resistivity of 1–100 Ω cm.

4.5.1 Suspension cooling time

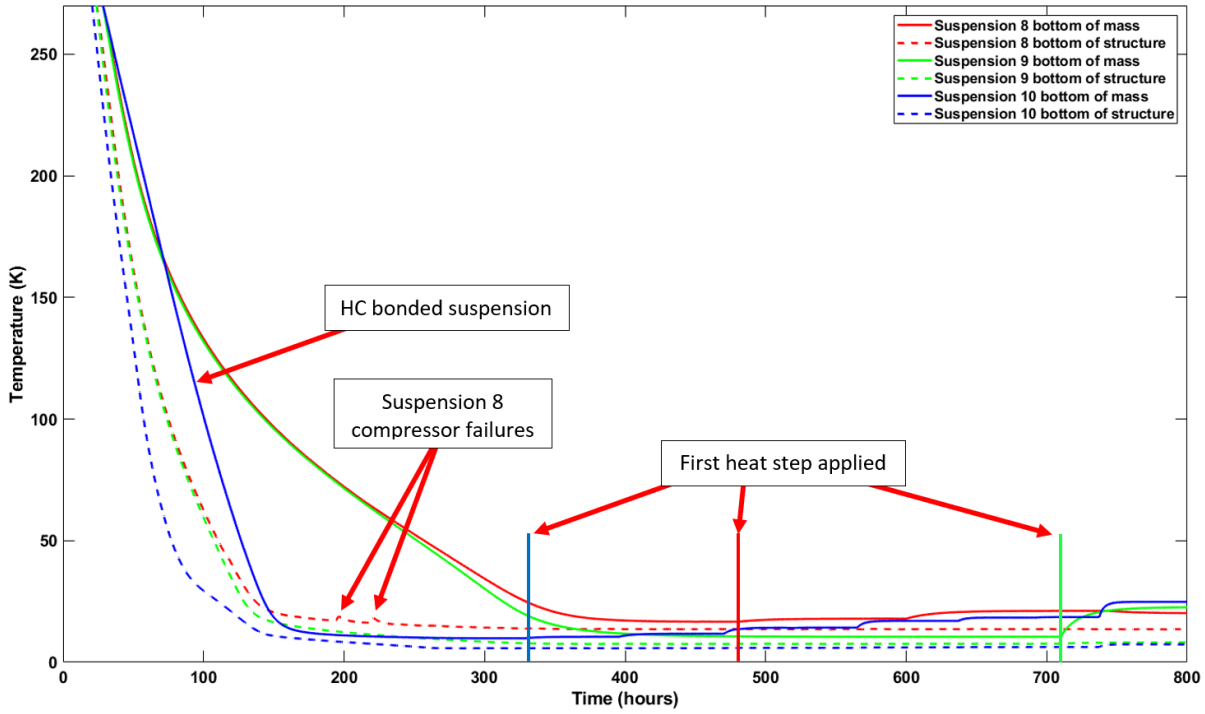

Figure 4.37: Cooling time for suspension 8, 9 and 10 showing the temperature measured at the bottom of the mass and bottom of the suspension structure. This data has been artificially time shifted to match the point of cooling from 270 K at the bottom of the mass sensor. The first heat step applied for each suspension is marked to show the cold thermal equilibrium temperature for each suspension.

Figure 4.37 shows the time taken for the bottom of the mass to reach steady state cold thermal equilibrium for suspensions 8, 9 and 10. The data shows the temperature measured by the Cernox sensors on the bottom of the mass and on the bottom of the suspension structure for each suspension. For comparative purposes, the data has been artificially time shifted to match the point of cooling at the bottom of the mass sensor from 270 K onwards. This is due to the fact that, for suspension 8, the compressor was switched on and off during the initial cooling from room temperature as part of the maintenance checks. Furthermore, compressor failures leading to loss of cooling around 200 hours show two warming peaks on the bottom of the suspension structure for suspension 8.

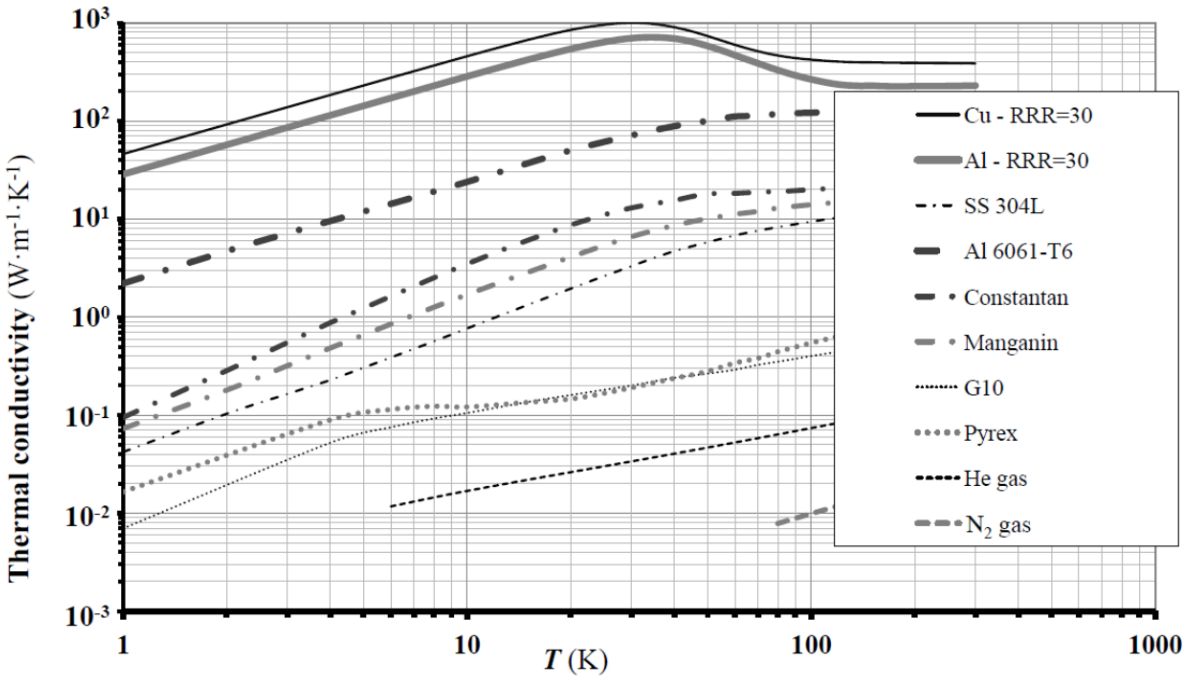

Figure 4.38: Thermal conductivity of aluminium compared to stainless steel [367].

One immediately notes that the cooling time for the HC bonded suspension to reach thermal equilibrium, compared to the epoxied suspension, takes significantly shorter. One can see that there is correlation between suspension structure cooling and the cooling of the mass. As the temperature of the structure drops, conduction dominates and so any change in conductivity of these 4 support poles will significantly change the cooling rate of the bottom suspension structure. This in turn will affect the temperature of the mass due to the contacting sensor wires thermally anchored to the suspension structure poles helping to cool the mass down. The epoxied suspension's structure poles are made from stainless steel, compared to the bonded suspension's aluminium poles. As can be seen in figure 4.38, these aluminium poles are approximately 2 orders of magnitude more conductive than the stainless steel poles from 50 K and below and also benefit from thermal grease between all points connecting poles to plates, unlike the epoxied suspension's stainless steel structure poles. As can be seen in figures 4.28, 4.29 and 4.36 the suspension structure poles are always colder than the mass, and so heat from the mass must flow through the sensor wires to the colder poles. While this may not be the sole reason for the faster cooling rate, it is a significant factor. Further work exploring heat loss mechanisms through these sensor wires is discussed in section 4.5.2.

One also notes the epoxied suspensions 8 and 9 cool down consistently with one another. This is expected considering this is the same suspension re-cooled.

For comparative purposes, the cold equilibrium temperatures are taken as the sensor reading just before applying the first heat step for each suspension, as marked on figure 4.37. The cold equilibrium temperatures between suspension 8 and 9 are quite different. Suspension 8 stabilised at 16.66 K, compared to suspension 9 at 10.46 K. The only significant difference be-

tween suspensions 8 and 9 was the blanking of all sapphire port windows with metal inserts. This 6.2 K difference is highly likely to be directly attributed to this change since uncoated sapphire windows do permit some infrared transmission below 6 µm wavelength, as shown in figure 4.39 [368]. This demonstrates that at these low cryogenic temperatures, even small amounts of power transmitted into the system can have the ability to significantly heat the mass.

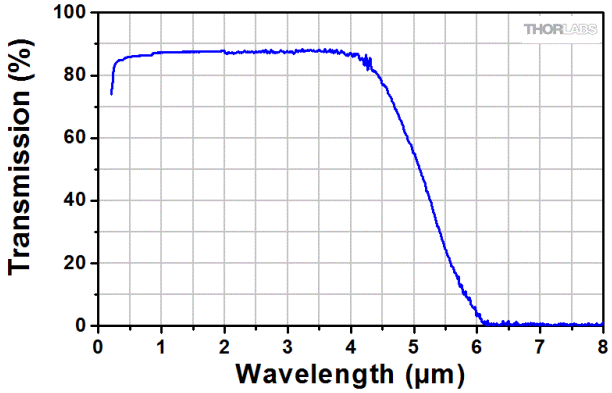

Figure 4.39: Transmission of 5 mm thick, uncoated sapphire windows, the same used in the Leiden cryostat ports. Image modified from [368].

Suspension 10 encountered the same cryogenic environment as suspension 9 since the port windows were also blanked. Suspension 10 equalised at 9.86 K, 0.6 K below suspension 9. The silicon ribbons for suspension 9 and 10 are similar in dimension and are cut from the same silicon wafer. It is therefore reasonable to assume the conductivity of the silicon ribbon is very similar. However, this 0.6 K difference at the bottom of the mass between suspensions likely arises from the fact the conductivity of the two different ribbons will not be exactly the same. This is investigated in the following section.

4.5.2 Cryogenic suspension ribbon conductivity

Each of the 3 suspensions has one variable change between them. Suspension 8 and 9 are the exact same suspension except suspension 8 had the sapphire port windows fitted to the cryostat, allowing an extra heat exchange mechanism in and out of the cryostat IVC. Suspension 9 and 10 are different suspensions with different connection techniques (macor-epoxy and HC bonding), masses and suspension support structures but experience the same cryogenic environment due to the port windows being blanked.

The thermal conductivity, k, of each suspension ribbon can be calculated using the steady state method, which is defined from appendix B as [369]:

$$P = \alpha T^3 A \frac{dT}{dL} \tag{4.1}$$

where P is the heat current in thermal equilibrium, A is the cross-sectional area of the ribbon and dT is the change in temperature measured over distance dL. Here, $k = \alpha T^3$, where T is

taken as the average temperature of the ribbon sensors and α is a scaling factor. This relation arises from the fact the heat steps applied generally keep the silicon ribbon within the specific heat dominated regime as shown in figure 2.14 of chapter 2. Any larger heat steps that push the silicon ribbon outside this T^3 relation are not used for conductivity analysis in this section.

At steady state equilibrium, all suspensions have some thermal gradient across them due to some ambient heat load, which is discussed further in section 4.5.4. As such, one must derive further equations to account for this, as shown in appendix B. Equation B.9 was used for determining α and hence the thermal conductivity, k, for each silicon ribbon suspension.

4.5.2.1 Heater power input and heat losses

It is important to determine the amount of power being put into the bottom of the suspension via the resistive heater mounted to the bottom of the mass. To do this, one refers to the heater circuit diagram in the inset of figure 4.21 in section 4.4.5.

Since the 1908P DMM constantly measured and recorded the heater circuit current and voltage, the DMM accuracy for DC current and DC voltage gives the errors required to determine the heater power error [370].

To determine the heater power, P_h, one can measure the current through the heater circuit, and use $P_h = I^2 R_h$ to determine the power dissipated within the resistor. Determining R_h requires care since one needs to know the value of R_h accurately.

The yellow lead wires running from the DMM to the cryostat each have a lead resistance of $0.05\,\Omega$, giving a total lead resistance of $0.1\,\Omega$. The thick red copper wires have a total resistance of $0.05\,\Omega$. The final stage of the heater wires were Lake Shore Quad-Twist WQT-36 phosphor bronze cryogenic sensor wires [335], these were measured to have a resistance of $4.15\,\Omega$ in total. Cryogenic wires are designed to have high electrical resistivity, and low thermal conductivity, to minimise heat flow into cryogenic systems, with the resistivity of the WQT-36 phosphor bronze wires being $4.02\,\Omega$m at 305 K, dropping to $3.34\,\Omega$m at 4 K [371]. Furthermore, including the wire thermally anchored to the poles, these wires were approximately 1 m in length, providing consistent resistance measurements with the data sheet.

The $12\,\Omega$ wire wound mount resistor on the bottom of the mass was a 12R5 WH20. The 5 % standard error on this heater was superseded when measured using the 1908P DMM, with the associated error stated in [370]. The heater resistance gave a value of $11.947\,\Omega$. Wire wound mount resistors are known to have stable resistivity across wide temperature ranges [372] [373] [374].

When the cryostat was closed, the resistance between the yellow lead wires was measured to be $16.4\,\Omega$, roughly consistent with the above measurements. When the cryostat was cold, this measurement dropped to $15.8\,\Omega$, roughly consistent with the $0.7\,\Omega$ drop in resistivity of the WQT-36 phosphor bronze wires. Care was taken to remove the measurement lead wire resistance for every measurement above.

One therefore assumes a fixed value for R_h of $11.947 \pm 0.014\,\Omega$, taking the resistive errors given in [370].

With the current measured for every temperature step, one can combine the errors of ΔI and ΔR_h to determine the error on the heater power δP_h. The average combined power error across all heat steps in all 3 suspensions was 0.2 %.

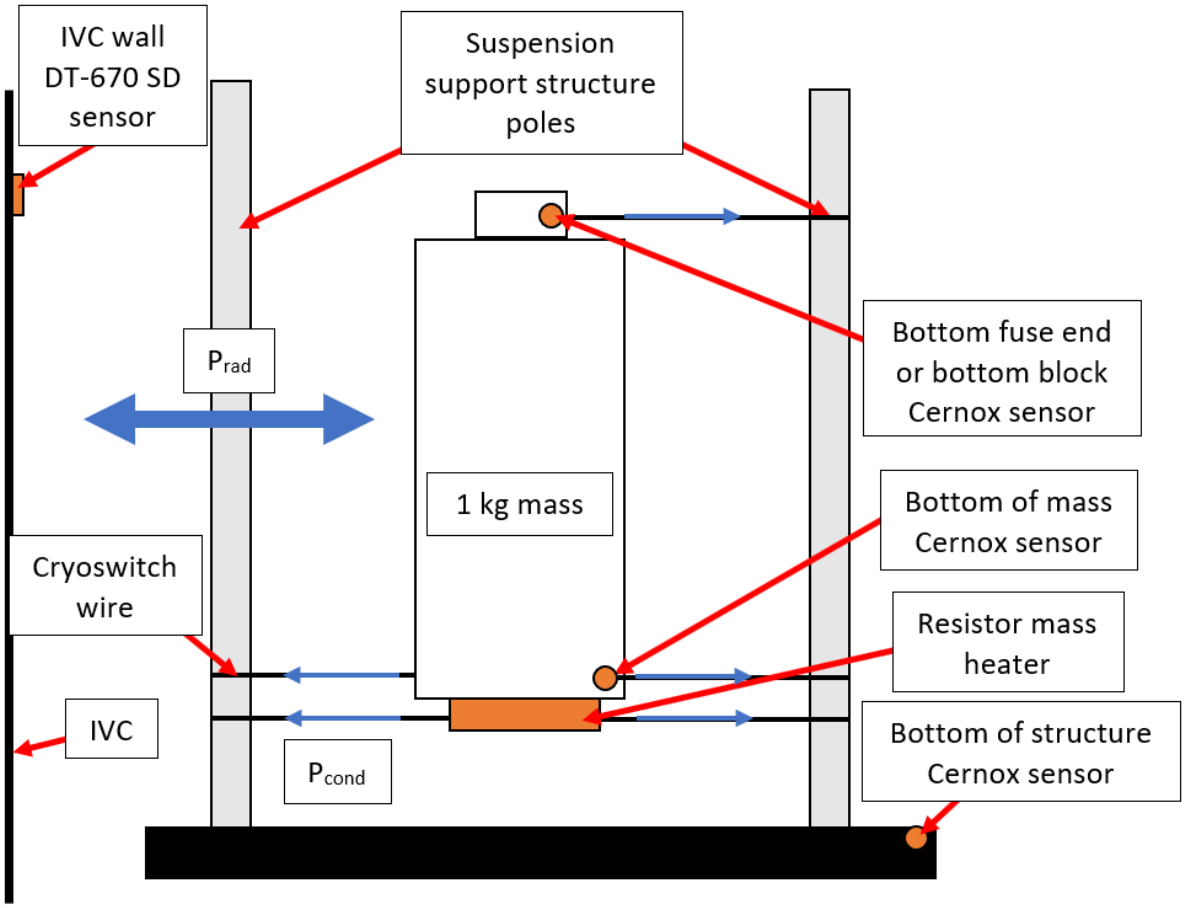

Figure 4.40: Schematic model for radiative and conductive losses.

Furthermore, it is also important to determine the heat losses from the suspension mass due to both conduction and radiation. For this analysis, only the mass losses are considered.

Conduction As previously stated, the suspension structure poles were always colder than the mass for every temperature step. Therefore, anything touching the mass, other than the silicon ribbon suspension itself, would help conduct heat away from the mass when heat was applied at the bottom. As shown in figure 4.40, there are 5 wires touching the mass at all times. Two wires are for the Cernox temperature sensors on the bottom fuse end (or block) and the bottom of the mass. Two wires are for the power supply to the heater on the bottom of the mass. Finally, in both the old and new cryoswitch designs, one wire is always attached to the mass, with the other wire attached to the copper contact plate.

Conductive losses are therefore estimated in the following way using equation B.2 in appendix B. One firstly assumes the temperature of the cold poles, where each of the wires are thermally anchored, are the same temperature as the bottom of the suspension structure since this is the closest sensor to the poles. It is likely the poles are slightly colder than the bottom of the suspension structure as logically they must sit somewhere between the coldplate temperature of 4 K and the temperature measured on the bottom of the suspension structure. The temperature of the mass is taken as the average temperature measured on the bottom fuse end (or block) sensor and bottom of the mass sensor. The ΔT value is the difference in temperature between the average temperature of the mass and the assumed temperature of the cold poles. This estimation provides a larger ΔT and hence is the conservative estimate on conduction losses since this assumes a higher conductive loss than what is expected for each heat step.

The estimated length of each conducting wire was 15 cm. Each wire consists of 4 wires of 36-AWG (127 µm wire diameter). Therefore, 5 wires, each consisting of 4 36-AWG wires, gives a total conduction area of 2.53×10^{-7} m^2.

Finally, the thermal conductivity of the wires must be estimated. Data is provided by Lake Shore for the thermal conductivity of the phosphor bronze wires at a variety of fixed temperatures [371]. This data was plotted between the 4–80 K temperature ranges and a 2nd order polynomial fit applied in order to interpolate the data for the thermal conductivity at the wire temperature for each heat step. The wire temperature was taken as the average temperature between the average mass temperature and the estimated pole temperature, taken as the bottom of the suspension structure temperature as before.

This provided an estimate of the conductive losses for each temperature step.

Radiation Similarly, radiative losses must be estimated for each temperature step. The radiative losses can be estimated using the following equation [369]:

$$P_{rad} = A\varepsilon\sigma(T_o^4 - T_s^4) \tag{4.2}$$

where P_{rad} is the radiative heat flow out of the mass, A is the emitting surface area of the mass, ε is the emissivity of the mass, σ is the Stefan-Boltzmann constant, T_o is the temperature of the emitting object and T_s is the temperature of the surroundings.

Due to the complexity of radiative loss estimations a number of assumptions were applied to the model shown in figure 4.40.

It was assumed that the entire surface area of the mass was simplistically emitting to the surroundings, noting that suspension 8 and 9 will have a different emitting area to suspension 10, due to the change from a denser stainless steel to an aluminium mass.

The emissivity, ε, was estimated by fitting 2nd order polynomials to data provided for stainless steel and aluminium in [375], and extrapolating or interpolating for the average mass temperatures respectively. This gave a range of emissivities for stainless steel of 0.009-0.032 and

for aluminium of 0.003-0.006.

It was also assumed the surroundings, in this case the IVC, were at the same temperature as measured by the IVC wall sensor. This is reasonable to assume since the actual line of sight of the radiating mass would be to the IVC wall at the same height as the IVC wall sensor. Similar to the conduction case, the mass was consistently hotter than the IVC wall across all temperature steps.

It should be noted here that, as discussed in section 4.4.10, the IVC wall sensor had detached from the wall. To mitigate for this, the temperature of the IVC wall sensor for all equivalent temperature steps applied for suspension 10, was applied to suspension 9. The IVC wall sensor is the T_s value. The T_o value is taken as the average mass temperature.

For every temperature step, radiative losses were compared to conduction losses. For suspension 8, where port windows were present, radiative losses were never more than 5 % of conductive losses for an average mass temperature <20 K. Comparatively, for suspensions 9 and 10 where no port windows were present, radiative losses were never more than 3 % of conductive losses for an average mass temperature <20 K.

With the above information, the combined loss estimations could be calculated for each step. Across all 3 suspensions, for the heat steps used to calculate thermal conductivity, the losses, herein defined as P_{losses}, were found to be between 21–238 µW.

With radiative and conductive losses now calculated, one can estimate the heat current through the silicon suspension by taking P_h and subtracting the combined estimated losses for each temperature step.

4.5.2.2 Dimensional data

Dimension	Suspension 8 Suspension 9	Error (%)	Suspension 10	Error (%)
Thickness (µm)	775	2.8	775	2.8
Width (mm)	3.51	3.7	3.50	3.7
Conduction area, A (m^2)	2.72×10^{-6}	4.7	2.71×10^{-6}	4.7
Length between sensors, L (mm)	168	0.6	89	1.1

Table 4.4: Dimensional data and corresponding errors for each silicon ribbon suspension.

The dimensions, A and L with associated errors, are recorded in table 4.4. The large ± 1 mm error arose due to the fact the DT-670 ribbon sensors were measured in place on the hanging suspension with a ruler. This error also takes into account the 3.175 mm long DT-670 SD package was wrapped in PTFE tape to secure it in place, making exact determination of the centre of the sensor difficult. The silicon ribbon cross-sectional dimensions were measured before assembly using the same RS PRO 150 mm digital callipers discussed in section 3.4.2.5 of chapter

3 [265]. The same error of $\pm 130\,\mu m$ when measuring the width of the sample and $\pm 22\,\mu m$ when measuring the thickness of the samples was applied.

4.5.2.3 Temperature data

The raw temperature data for each heating step for each suspension is shown in appendix C.

All suspension have two DT-670 SD sensors on the top and bottom of the ribbon. Suspension 8 has an uncalibrated DT-670 sensor on the top ribbon sensor. The bottom sensor is calibrated. For both suspensions 9 and 10, both sensors are calibrated DT-670 sensors.

Error analysis takes the sensor errors given for the two calibrated DT-670 SD sensors in tables 4.1 and 4.2. For the uncalibrated sensor, a similar set of tables are produced in appendix C.4 to provide the absolute temperature error. Errors are linearly interpolated between given temperature points in these tables.

These errors are then propagated through using the methodology shown in equation B.9 of appendix B.

4.5.3 Conductivity results

This section aims to calculate the thermal conductivity of the silicon ribbons from suspension 8, 9 and 10 from the temperature distribution data given the power, dimension and temperature information above.

It must be noted here that while there is a temperature gradient across all suspensions at cold thermal equilibrium, arising from an unknown ambient load, equation B.9 of appendix B only requires the *change* in power between ambient and any heat step in order to calculate α. While no power is input through the resistor, this ambient load creates losses, $P_{ambient\,losses}$, at the suspension's cold thermal equilibrium due to this temperature gradient.

Therefore, ΔP, in equation B.9 of appendix B can be expressed as:

$$\Delta P = P_{heat\,step} - P_{heat\,step\,losses} - P_{ambient\,losses}. \tag{4.3}$$

Therefore the ΔT terms in equation B.9 of appendix B are defined as $T_{heat\,step} - T_{ambient}$ for the top and bottom ribbon sensors.

Due to the underlying specific heat-dominated $k = \alpha T^3$ relationship, only heat steps where the average ribbon temperature was deemed to be within this regime (<25 K), were used to determine α, as shown in appendix C. For data where the average ribbon temperature was >25 K, and especially for extremely high heat steps, the estimates of α were found to be no longer valid, particularly due to the crude radiation estimates which become more significant and require greater accuracy as the temperature increases. This relationship also shows that α should be a constant value for each heat step. For each suspension, an α value was calculated for each heat step. If one then takes the mean of these individual α values for each suspension, one

can determine the thermal conductivity at any temperature within the specific heat-dominated T^3 regime.

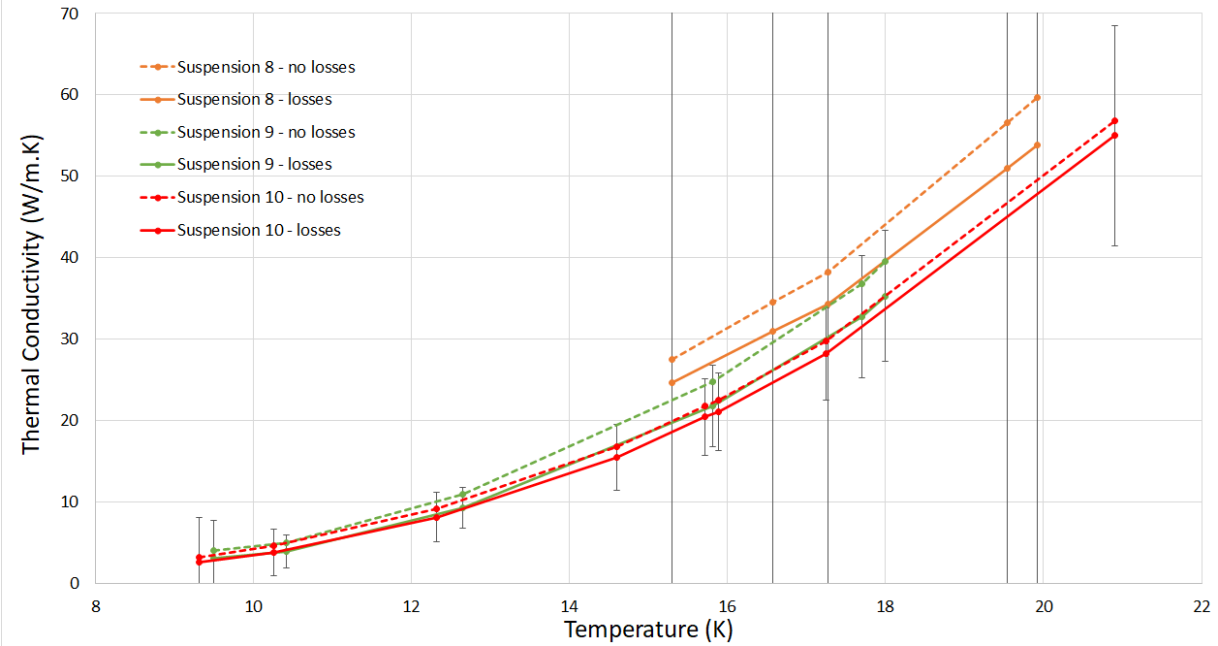

Figure 4.41: Thermal conductivity as a function of temperature, with and without estimated losses, for each heat step and each calculated α value and its corresponding individual maximum error.

Figure 4.41 shows the thermal conductivity, with and without estimated losses, for suspensions 8, 9 and 10 for their respective heat steps.

The error bars are applied to the conductivity data that includes estimated losses, and are the maximum percentage error obtained from equation B.20 from appendix B at each power step for each α value, applied to the corresponding calculated conductivity value k. Due to the uncalibrated sensor used for suspension 8, these errors ranged from 115–233 %, which are off the scale of figure 4.41.

Some correlation was found in the middle regions of data, for suspensions 9 and 10, between the standard deviation from the mean α value for each heat step and the corresponding calculated error in α, $\delta\alpha$, however this correlation broke down at the lowest and highest temperature steps. This is observed in the error bars in figure 4.41, where the percentage error from $\delta\alpha$ is directly applied to k. For the lowest power steps, this correlation was likely degraded due to the estimation of power losses having a more significant affect at these low powers. For the highest power steps this was likely due to a combination of radiative effects coming into play alongside the high temperatures starting to leave the specific heat-dominated T^3 regime.

One immediately notes each suspension has a higher conductivity if one was to remove the estimated losses. Indeed, if one overestimates ΔP, one ends up overestimating α, which increases the calculated thermal conductivity, k as can be seen in equation B.9 in appendix B.

This also shows it is easier to mistakenly measure a higher thermal conductivity on a sample rather than a lower one. This is likely one contributing factor as to why suspension 8 has a higher mean α value compared to suspensions 9 and 10 arising from the extra radiative power input through the port windows.

The fact that suspension 9 and 10, inclusive of losses, match relatively well is an indicator that these loss estimations are reasonable.

Suspension	α	Error max, $\delta\alpha_{max}$ (%)	Error min, $\delta\alpha_{min}$ (%)
8 - epoxy-macor, port windows	6.79×10^{-3}	961	234
9 - epoxy-macor, no port windows	4.84×10^{-3}	113	25
10 - HC bonded, no port windows	4.75×10^{-3}	143	22

Table 4.5: Calculated mean α values for each suspension with total combined maximum and minimum errors calculated from appendix B.

The mean α value constants for suspension 8, 9 and 10, including the maximum and minimum errors, as defined in appendix B, are shown in table 4.5.

To take the most conservative maximum error on the mean, $\delta\alpha_{max}$, requires finding the maximum absolute value given from equation B.20 in appendix B, which arises from the lowest power heat steps. This leads to the significant maximum errors shown in table 4.5. These maximum errors arise directly due to the $\delta\alpha_4$ term in equation B.18 of appendix B where the $\frac{1}{T_{term}^2}$ term dominates. For suspension 8, these errors are exacerbated further due to the δT_{term} being significantly larger since the top ribbon sensor was an uncalibrated DT-670 SD sensor, which produces an absolute temperature error $\approx$15 times that of a calibrated sensor at 10 K.

For comparative purposes, if one were to take the minimum absolute errors, which typically arise from the middle heat steps where the combined errors are smallest, these errors would reduce to the $\delta\alpha_{min}$ values recorded in table 4.5.

Suspension 8's higher mean α value is shown table 4.5. This likely arises due to the excess heat transfer through the port windows that were not fitted for suspension 9 and 10. The standard deviation of α values for the heat steps in suspension 8 was approximately one order of magnitude lower than suspensions 9 and 10. This higher ambient heat load is likely masking the true α variation measured across each heat step, particularly at lower power steps, artificially inflating the mean α value. Therefore, the reason for this tight standard deviation is likely a result of the excess heat transfer from the port windows. Indeed, one found that by reducing the ambient load loss estimations for suspension 8, these α values reduced closer to the α values of suspension 9 and 10. One is also reminded that suspension 8 and 9 are exactly the same suspension, except suspension 8 experiences a different cryogenic environment due to the port windows being present.

Fundamentally, the α values recorded in table 4.5, for suspension 9 and 10, are as predicted by the underlying physics, where two silicon ribbons, cut from the same wafer with similar

dimensions, should have a similar measured thermal conductivity and hence α constant when measured in the same cryogenic environment.

One can now apply these α values in order to calculate the thermal conductivity of the silicon suspension ribbon across the 4–25 K specific heat-dominated regime. Furthermore, one can then compare this data to other experimentally measured values of silicon.

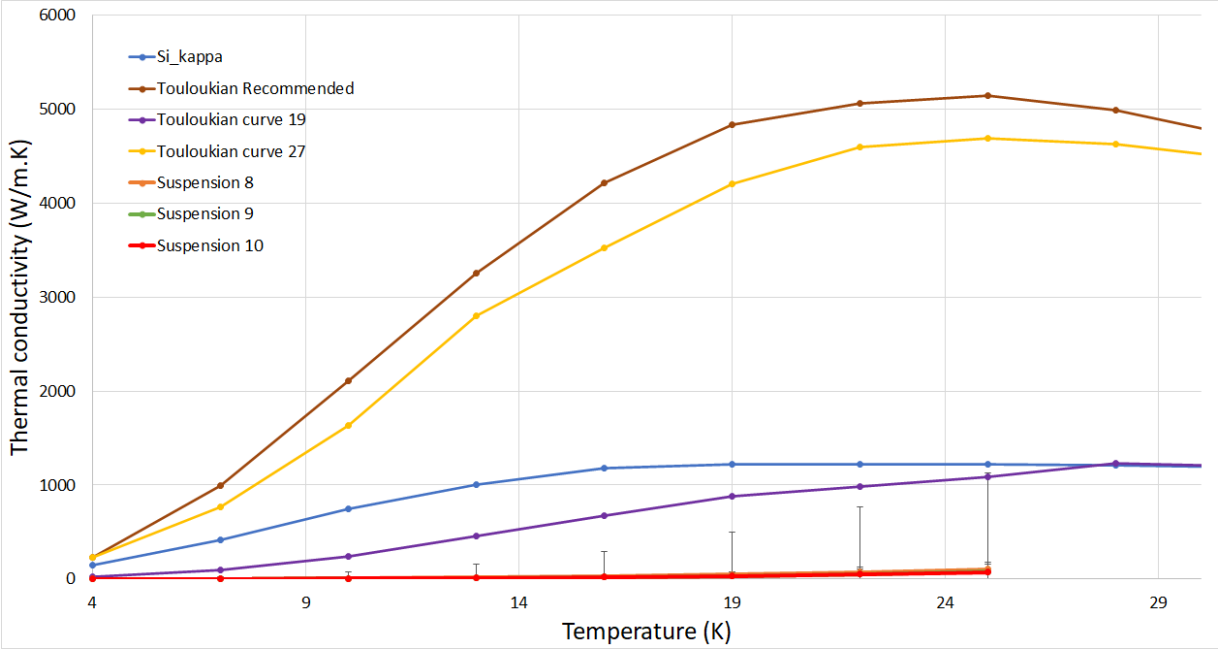

Figure 4.42: Thermal conductivity as a function of temperature of silicon suspension ribbon compared to experimentally measured thermal conductivity from sources [319] [376]. Thin black error bars shown correspond to $\delta\alpha_{max}$ error applied.

Figure 4.42 shows the thermal conductivity of the ribbon suspensions 8, 9 and 10 calculated using the mean α values in table 4.5. The data in figure 4.42 is plotted alongside experimentally measured thermal conductivity data of silicon samples obtained from Hull's Properties of Crystalline Silicon [319] denoted as "Si_kappa" and work conducted by Touloukian [376]. Curves 6, 19, 26, 27, 33-37 and 46 from Touloukian's work were found to be comparable boron doped (p-type) silicon samples within a similar resistivity to the silicon suspension ribbons. These curves are shown in appendix D. However, many curves were not provided over the temperature range of interest for comparison. As such, curves 19 and 27 that are representative of the high and low thermal conductivity cases were chosen. Touloukian's recommended curve for high purity silicon was also included [376]. The maximum percentage error from $\delta\alpha_{max}$ has been applied to the conductivity values, k, purely for completeness here. These wide error bars associated with suspension 8's large uncertainties, arising from the uncalibrated sensor, should not be mistaken for under-measuring the silicon thermal conductivity, since suspensions 9 and 10, with more accurate sensors and $\delta\alpha_{max}$ errors applied do not come close to matching the measured thermal conductivity data provided. Since the silicon ribbon suspension between 8 and 9

is the exact same, one should therefore only consider suspension 9 and 10 when comparing to the experimentally measured data.

One immediately notes that the silicon ribbon's conductivity in suspensions 8, 9 and 10, where 8 and 9 are the exact same ribbon, is significantly lower than even the least conductive experimentally measured sample within this temperature range.

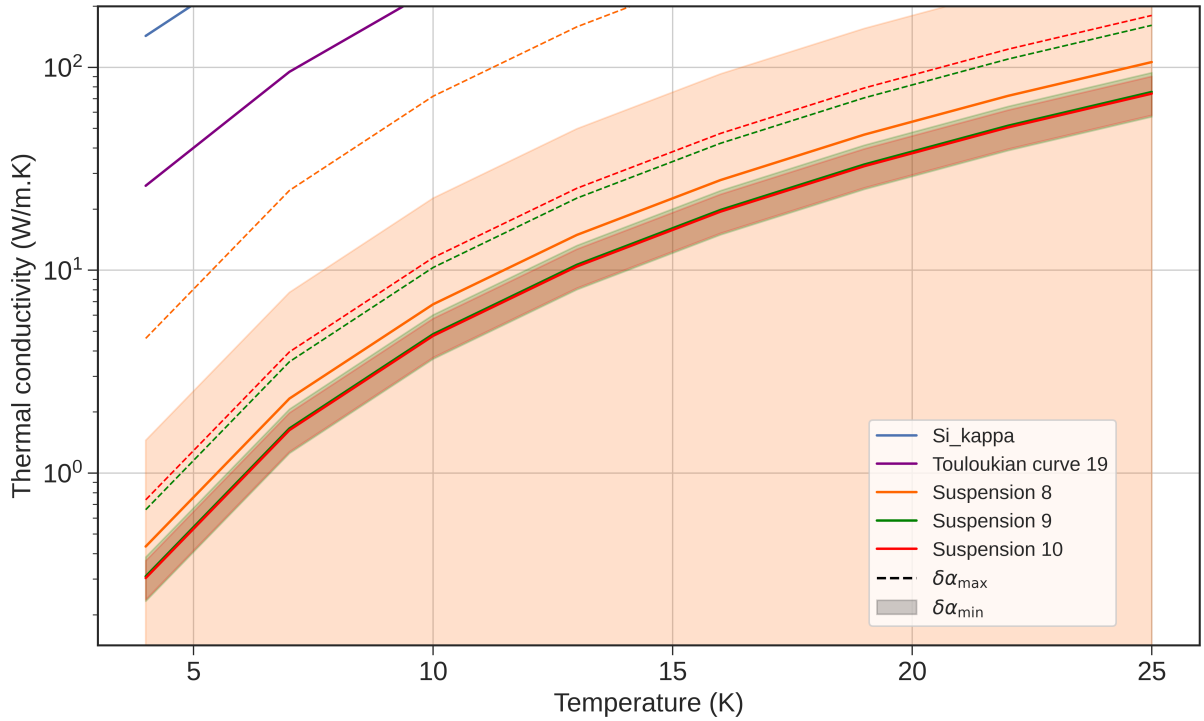

Figure 4.43: Thermal conductivity as a function of temperature for each suspension using each suspension's mean α value plotted alongside measured thermal conductivity of silicon [376].

Figure 4.43 zooms in on figure 4.42 to more clearly compare the ribbon conductivities compared to measured silicon conductivity data. Once again, suspension 8 errors, both $\delta\alpha_{min}$ and $\delta\alpha_{max}$, are off the scale of figure 4.43 due to the large errors associated with the uncalibrated sensor.

There are 3 primary considerations that mean the measured thermal conductivity of the silicon ribbon suspension is lower than expected when compared to other experimentally measured silicon.

Firstly, as discussed in section 4.4 and sections 3.4 and 3.8.2.2 of chapter 3, the silicon samples are cut using lasers which produce a thermal oxide layer near the cut edges. This thermal oxide layer indicates significant heating and melting of the silicon in order to pierce through the wafer to cut the sample. Melting of single crystalline silicon will produce polycrystalline or amorphous silicon which is known to have a significantly lower thermal conductivity [377] [378] [379]. While the depth of polycrystalline or amorphous silicon into the single crystalline sample is unknown, it is suspected this may contribute to the reduction in conduc-

tivity due to the reduction in conduction area of the more highly conductive single crystalline silicon. Future work, for example X-ray crystallography, should determine the depth and structure of the laser damaged silicon in order to both investigate damage depth for strength work conducted in chapter 3 and quantify the reduction in area experienced by the single crystalline silicon region.

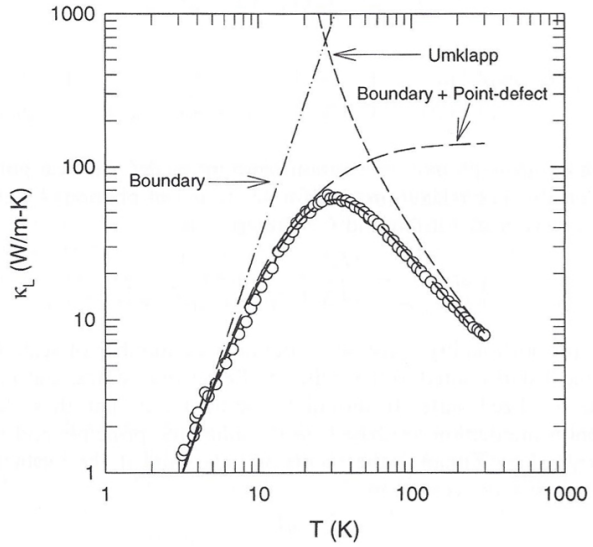

Figure 4.44: Example of the effects of boundary and point defect or impurity scattering on the phonon heat transport [380].

Secondly, boundary or geometrical scattering can not be ignored, particularly for suspension elements that are typically very long and thin. Boundary scattering can become a significant limiting factor of thermal conductivity at low temperatures, as shown in figure 4.44. One can determine the boundary scattering limit from the following equation [381]:

$$l_m = \frac{3k}{C_v v} \tag{4.4}$$

where l_m is the phonon mean free path, k is the temperature-dependent thermal conductivity, C_v is the volumetric specific heat capacity and v is the Casimir sound velocity in the material, which for silicon is $5660\,\mathrm{m\,s^{-1}}$ [382] [383].

If one takes the highest and lowest thermal conductivities measured across all 3 suspensions, the corresponding phonon mean free path, l_m, is 2.1–3.6 µm, from an average ribbon temperature of 9.3–19.9 K respectively, significantly shorter than the thinnest dimension of the silicon suspension ribbons at 775 µm.

To find a scenario where the phonon mean free path approaches dimensions on the scale of the silicon suspension ribbon would require one to use data from Touloukian's recommended curve. The phonon mean free path for an average ribbon temperature of 10–20 K ranges from 1024–190 µm respectively. This recommended curve is for high purity silicon, with typical

resistivity in the order of multiple 100s or 1000s of Ωcm, far greater than the stated resistivity of the suspension ribbons which are 1–100 Ωcm, boron doped (p-type) (100) silicon [376] [384].

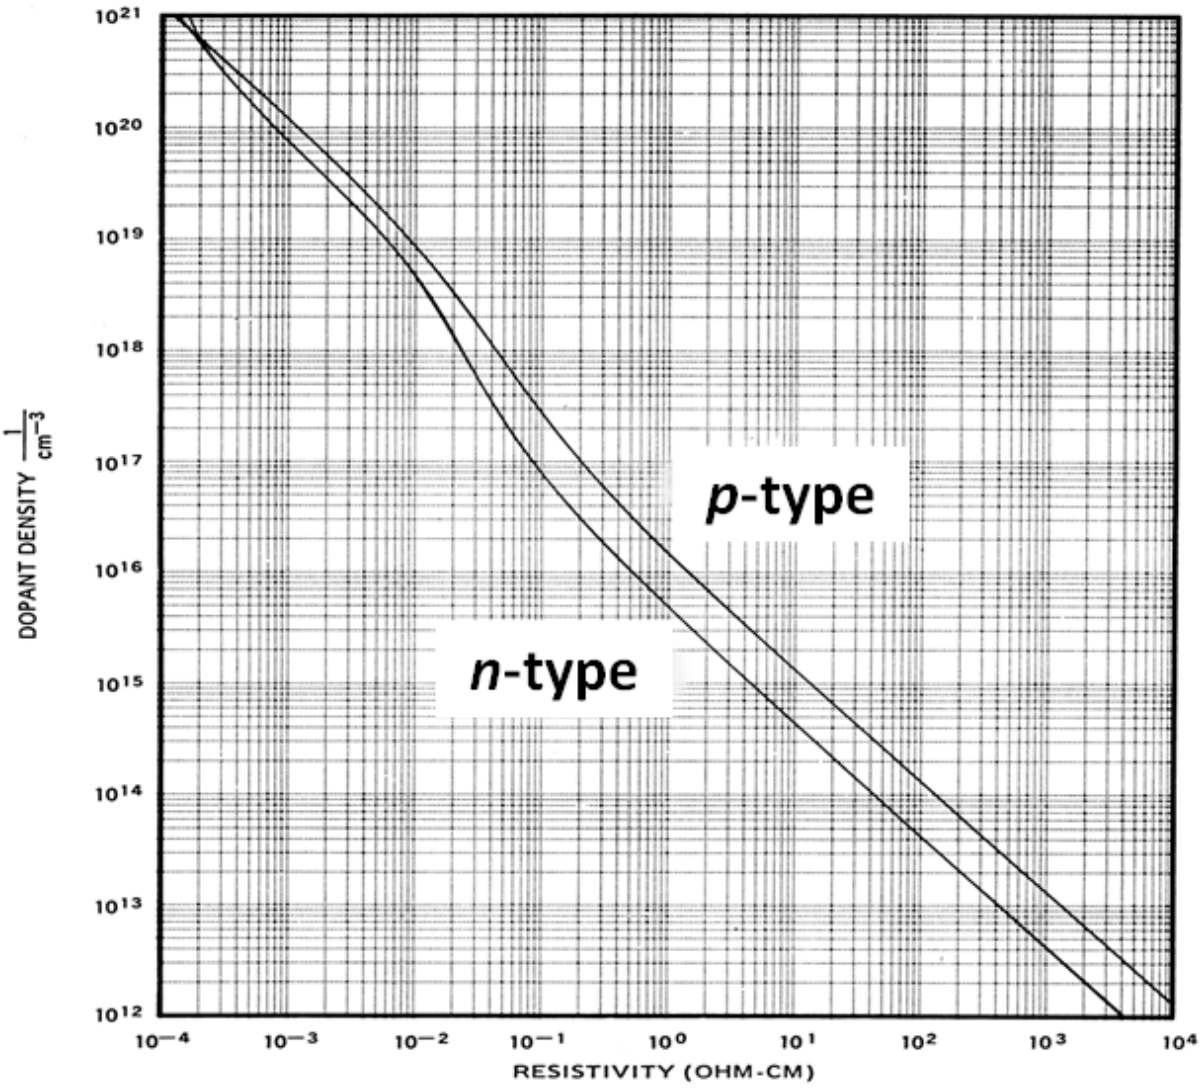

Figure 4.45: Resistivity of p-type and n-type silicon versus dopant density [385].

Thirdly, the dopant concentration and resistivity of the silicon ribbons will have a significant impact on the thermal conductivity. As stated in section 2.3.5 of chapter 2, heat transfer in crystalline silicon is dominated by phonon transport [201]. Phonon scattering from point defects or dopant impurities in a sample can inhibit thermal conductivity measurements, particularly at low temperatures where the Umklapp process is not a limiting factor, as shown in figure 4.44 [202] [380].

Sample	Dopant concentration (Boron atoms per cm^3)	Resistivity (Ωcm)	Sample cross sectional dimensions (mm)
Curve 19	1×10^{15}	15	3.00×3.00
Curve 27	4×10^{14}	35	0.63×0.64
Suspension 8/9	1.5×10^{14} – 2×10^{16}	1-100	3.51×0.775
Suspension 10	1.5×10^{14} – 2×10^{16}	1-100	3.50×0.775

Table 4.6: Dopant concentration, resistivity and cross sectional dimensions of silicon samples used to produce Touloukian's thermal conductivity curves in [376] and silicon ribbon's used in this thesis. The resistivity values used have been estimated from figure 4.45.

Table 4.6 shows the dopant concentration, resistivity and sample sizes of the silicon samples used to produce curves 19 and 27 in Touloukian's data [376] alongside the silicon ribbons used for suspensions 8, 9 and 10. The resistivity values for Touloukian's curves have been estimated from figure 4.45. With no independent measurement of the silicon ribbon suspension thermal conductivity, one can only infer the dopant density from the resistivity range as shown in figure 4.45. The inferred dopant density for the silicon ribbon suspensions are shown in table 4.6.

It is clear to see from figure 4.42, the difference between the thermal conductivity for curve 19 compared to curve 27. If one were to assume the silicon ribbons are at the lower end of the 1–100 Ωcm resistivity range, and hence higher dopant density of 2×10^{16} boron atoms per cm^3 this would likely significantly reduce the thermal conductivity of the sample due to the lack of sample purity.

It is believed these factors ultimately reduce the conducting cross-sectional area of the single crystalline silicon and inhibit phonon transport within the ribbon sample, reducing the thermal conductivities as shown in figure 4.42.

Further work is required to determine by how much each of these factor in and if these mechanisms alone can account for the discrepancy between the measured thermal conductivity of the silicon ribbon suspension compared to the thermal conductivity of other experimentally measured samples with similar material properties.

4.5.4 Ambient heat load

The fact a temperature gradient exists across suspension 8, 9 and 10 when at cold thermal equilibrium, implies there is an unknown ambient heat load present on the mass.

One can estimate the ambient heat load in each suspension by defining the following equation:

$$P_{ambient} - P_{ambient\,losses} + P_{sensor} = P_{heat\,current} \tag{4.5}$$

where $P_{ambient}$ is the unknown ambient heat load on the system, $P_{ambient\,losses}$ is the estimated heat loss at cold thermal equilibrium and $P_{heat\,current}$ is the heat current through the ribbon at

cold thermal equilibrium, calculated using equation B.5 in appendix B and using the α values listed in table 4.5 for each suspension. P_{sensor} must also be accounted for due to the sensor self heating of the bottom ribbon sensor which sits between the mass, where one assumes the majority of ambient heating applies and the rest of the ribbon above. This was taken as 16 µW quoted by Lake Shore at 4 K [343].

Suspension	Ambient heat load (µW)
8 - epoxy-macor, port windows	570.0
9 - epoxy-macor, no port windows	108.8
10 - HC bonded, no port windows	99.7

Table 4.7: Estimated ambient heat load on each suspension.

The estimated ambient heat loads for each suspension are recorded in table 4.7. The higher ambient load on suspension 8 is consistent with the fact this system had sapphire port windows in place during the experiment, allowing radiation to transmit through the system. Suspensions 9 and 10 should experience the same ambient load if experiencing the same cryogenic environment. The cause of this exact discrepancy in calculated thermal loads likely arises from the combination of estimated heat losses and slight variance in α values as shown in table 4.5.

To investigate the source ambient heat load further, one can only compare suspensions 9 and 10 since they shared the same cryogenic environment. Several conduction experiments were carried out on the different suspensions to investigate this ambient heat load. Conduction investigations were carried out in section 4.4.9, involving disconnecting the cryogenic switch to ensure no residual power was leaking into the system. This experiment was conducted only for suspension 8 and was never repeated for suspension 9 and 10. If this was the source of ambient heat load, the effect of disconnecting it may have been lost due to the higher ambient radiative window heat load in suspension 8. It is recommended that this experiment should be repeated in a cryogenic environment with the port windows blanked.

Heating of the thick copper heater wires was conducted on suspension 10 to investigate this potential conduction path introducing heat into the system with no observed heat leak into the system. This experiment should be conducted over a longer period of time in case there is a long thermal lag or the thick copper wires should be removed entirely.

In section 4.4.11.3, temperature sensors were individually disconnected and re-connected to investigate sensor self-heating inputting power into the system. Sensor self-heating was observed to be approximately 14 mK on a DT-670 SD sensor, which was found to be roughly consistent with the 16 µW of self-heating quoted by Lake Shore at 4 K [343]. With Cernox sensor self-heating measuring 10^{-7} W at 4 K their contribution was considered negligible [344]. The two DT-670 sensors alone, assuming maximum power dissipation, can not account for the total ambient load.

These mechanisms of conductive heat transfer, while unable to be ruled out entirely, are

unlikely to contribute significantly to the source of ambient heating of the mass, primarily due to the fact all wiring is thermally anchored to the suspension structure poles that are known to be colder than the 1 kg mass at all times. Due to conduction dominating at these lower temperatures by the fact that the estimated radiative losses from the mass as a percentage of conductive losses never exceeded 3 %, radiation was also not considered to be the primary driver of the ambient heat load for suspensions 9 and 10.

Two more sources were considered as a source for heating the mass at cold thermal equilibrium; cosmic ray flux and electromagnetic fields. For this work, the estimates were applied to the aluminium mass, used in suspension 10, due to the larger mass surface area and diameter, compared to the stainless steel mass, where these effects should be most pronounced. Heating from cosmic ray flux was estimated by using [369]:

$$dQ = mcdT \tag{4.6}$$

where dQ is the quantity of heat input from a cosmic particle, m is the mass, 1 kg, c is the specific heat of the mass and dT is the corresponding temperature rise in the mass from cosmic ray heating. To estimate dQ, one assumes a cosmic ray flux at sea level of $1.44 \times 10^{-6}\,\mathrm{m\,s^{-2}}$ [386], giving a particle strike rate of one particle per m^2 every 8 days. Scaling this to the top exposed half of the mass, this corresponds to a particle strike approximately once every 2 years. If one conservatively assumes high energy muons, of energy 100 GeV [387], or 1.6×10^{-8} J, and all the energy is deposited in the mass, this corresponds to a temperature rise, dT, of approximately 8 nK. Considering the strike rate and temperature rise, estimations of cosmic ray heating were considered far too small to be of any significance.

Magnetic field estimates, concerned with 50 Hz pickup from nearby lab equipment, estimated the magnetic flux and subsequent induced heating of the cryogenic mass. The electromagnetic pickup corresponds to a fluctuating field with a period of 20 ms. If one makes a simplistic assumption by treating the mass as a hollow ring, with a ring diameter equal to the mass diameter of 65 mm, one estimates a magnetic flux of 3.3×10^{-9} Wb. By dividing the magnetic flux by the period of fluctuation, one produces an induced voltage within the ring of 166 nV. Taking the resistivity of 6061-T6 aluminium as $4 \times 10^{-8}\,\Omega$m [388], one can scale this to the circumference of the ring, to give a resistance of $8.2 \times 10^{-9}\,\Omega$. This induces a current of 20.4 A. While this might appear excessively high, high inductive currents arising from magnetic fluctuations in rings are not uncommon [389]. The power induced into the mass is therefore 3.4 µW. These estimates show inductive heating to being on a similar order of magnitude as the ambient loads estimated in table 4.7.

While these estimations may go some way in accounting for the ambient heat loads they can not account for the total estimated ambient heat load that remains in the system. Further investigations are required to better characterise the cryogenic system and its environment in order to account for all unwanted sources of heat within the system.

4.6 Conclusion

The work in this chapter has ultimately demonstrated what is believed to be the world's first HC bonded silicon suspension. Furthermore, this work has demonstrated the successful hang of this HC bonded silicon suspension down to cryogenic temperatures and back up to room temperature without failure, surviving for over 8 months. The suspension remains intact at the time of writing, albeit locked up and supported ready for future work.

Iterations of each suspension have shown the importance of the coefficient of thermal expansion in jointing prototype silicon suspensions, particularly if epoxies or metals are used in any components of the suspension system.

Furthermore, cryogenics poses a significant challenge in the context of sensing, with unwanted and poorly characterised heat sources present throughout the system. Further work is required to better characterise the cryostat and improve any cryogenic wiring techniques for future iterations.

This work has also demonstrated the ability to extract thermal conductivity values from this cryogenic experiment. While the results are not in agreement with other experimentally measured thermal conductivities of silicon, a number of factors inhibiting phonon transport and minimising the cross-sectional conductive area are believed to be contributing to these low measured thermal conductivities. Future work should look to better characterise these silicon ribbons, particularly to understand the crystal structure across the full cross-section to provide insight on the surface damage, discussed in chapter 3, and in particular to provide a better understanding of thermal conductivity for long, thin suspension samples with a finite surface roughness.

This work paves the way for future iterations of silicon suspension prototypes that should look to increase the strength and minimise the roughness of the suspension elements, enhance robustness of the support structure and ultimately provide sufficient thermal extraction in order for the test mass to run at proposed cryogenic temperatures necessary for 3G detectors to achieve target sensitivity.

Chapter 5

FEA modelling of suspension thermal noise in 3rd generation gravitational wave detectors

5.1 Introduction

In order to understand noise within 3G detectors, one must determine losses and energies throughout the suspension elements. This is usually determined by mathematical modelling using FEA.

By modelling suspension elements, one can experiment virtually with a variety of parameters and investigate their impact on noise in gravitational wave detectors. Furthermore, by using work conducted in chapters 3 and 4, one can analyse the effects of real suspension stresses and designs respectively on the overall suspension thermal noise performance.

This work looks initially at tolerable stresses of a real silicon suspension followed by modelling of silicon's 3 crystallographic axis and subtleties of modelling short, thick ribbons compared with theory. Finally, directly modelled HC bonds between silicon suspension elements in ANSYS Mechanical APDL have been demonstrated for the first time for complete 4 fibre suspension geometries.

Previous FEA work on silica and some initial work on silicon has been conducted by Dr Alan Cumming [136] [144], Dr Rahul Kumar [390], Dr Karen Haughian [150] and Dr Liam Cunningham [364] [391]. Dr Liam Cunningham advised on the work involving HC bonds between silicon substrates.

5.2 FEA modelling in ANSYS Mechanical APDL

The general working principle of FEA is to take the geometry of interest, apply material properties, split the geometry into a series of finite elements, apply constraints such as loading or force

to the geometry and calculate the desired results, e.g. deformation, for each individual element, providing a final result for the whole geometry as a sum of its elements.

Hierarchically, as shown in figure 5.1, one can start by defining geometry keypoints, connecting the keypoints by lines and then connecting lines into areas or volumes. One can then define how to split the area or volume into finite element sizes, in a process known as "meshing". These elements are defined by a series of nodes that typically sit on the ends, mid-points or vertices of intersecting lines to define individual element geometries. This then defines the entire geometry as the sum of its elements. One can increase the accuracy of the model by increasing the number of nodes and elements however this is at the expense of finite computational power.

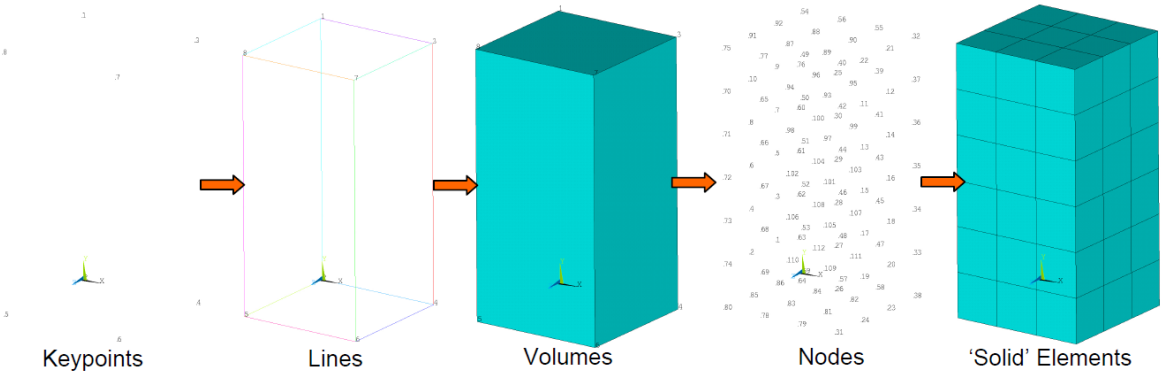

Figure 5.1: Hierarchical assembly of ANSYS model geometry and subsequent meshing to produce FEA model. Image obtained from [144].

Previous work on suspension fibres and ribbons notes the difficulty for FEA modelling long, thin geometries [144] [390]. This is typically due to the fact a high number of elements are required along the length of the suspension geometry for accurate modelling. This, in turn, can lead to scenarios where the model can not be solved due to PC memory issues or sheer length of computation time.

One can use beam elements for efficient fibre or ribbon analyses and solid elements to model the test mass, as has been used in previous work [144] [390]. For the inclusion of bonds to connect suspension elements together, beam elements have previously been jointed with shell elements to simulate the bond region, however it was discovered that shell elements are not suitable for modelling constrained surfaces. This is due to the inability of shell elements to model shear stresses [391]. Due to more powerful computing available, all FEA models produced in this thesis use solid elements, specifically SOLID185 and SOLID186 elements. SOLID185 elements provide 8 nodes per element, with one at each each corner vertex; where SOLID186 elements provide 20 nodes per element, with one at each corner vertex and additional nodes at each midpoint of the edge lines between vertices. Solid elements can model shear stresses and also give the ability to model material anisotropies for crystalline materials like silicon.

While ANSYS Mechanical APDL is a GUI based program, the majority of the work con-

ducted in this chapter was generated by using the ANSYS Parametric Design Language (APDL) [392]. This language is loosely based around C, C++ and Fortran programming languages. There are multiple advantages to using code for FEA, namely the ability to access legacy commands and elements, the speed and automation for building and solving models and the ability to easily change parameters within a model.

Thanks to progress in CPU power and memory size, as well as lower overall costs, FEA models generated today can compute, in reasonable time, models that historically were too intensive. Nonetheless, limits still remain and must always be taken into account when FEA modelling via meshing efficiencies and novel techniques, as will be used in this chapter.

5.3 Tolerable stress in the suspension

For experimental work carried out in chapter 4, it was necessary to determine the amount of stress the 1 kg suspension was able to take before breaking. By modelling the suspension geometry one can apply loads or deliberately misalign suspension elements and examine the stress throughout the system to help determine if the resultant real suspension geometry was viable.

It should be noted, for this section only, the fibre geometry used was 5 mm wide by 525 µm thick. The final, HC bonded, suspension geometry after procuring suspension length silicon ribbons from LML was 3.5 mm wide by 775 µm thick, as stated in section 4.4.11 of chapter 4. This change in dimension was due to LML's increased confidence in producing silicon ribbons that were thinner in width [257]. As such, all FEA work subsequent to this section uses these updated real ribbon geometries.

Furthermore, for this section only, any reference to "bonds", purely refers to two areas on each geometry that are fixed together in the model. No HC bonds are included in this section of work.

Modelling of the suspension was conducted with the following geometries, where the top and bottom blocks were attached to the ribbon along their full 10 mm height in the Y direction:

- Top block - $5 \times 10 \times 20$ mm
- Ribbon - 5×250 mm, 525 µm thick.
- Bottom block - $5 \times 10 \times 20$ mm
- 1 kg mass - 65 mm diameter, 120 mm length

The material properties were as follows:

- Top block, ribbon and bottom block - silicon - $E_{isotropic} = 157$ GPa, $\nu_{isotropic} = 0.27$, density $= 2328\,\mathrm{kg\,m^{-3}}$
- 1 kg mass - aluminium - $E = 68$ GPa, $\nu = 0.33$, density $= 2700\,\mathrm{kg\,m^{-3}}$

The assembled geometries are shown in figure 5.2(a). The model was then constrained on its top area by a fixed support and gravity was applied to the system, to simulate the effect of hanging the mass. The stress in the ribbon is 3.74 MPa, 6 % of the 60.3 MPa *control new* median strength found in chapter 3. With the cylindrical mass, it was discovered that the mass naturally hung off centre in the FEA simulation. While this was partly due to the fact the top and bottom blocks of the structure cause a mass offset, upon building up to the final model it was discovered that even a single ribbon, affixed centrally and directly onto the top of the cylindrical mass, still caused a deflection, as shown in figure 5.2(b). Due to the mathematical nature of meshing circles or cylindrical volumes in FEA, this produces a series of elements that are uneven in size around the centre of the circle or cylinder. These uneven element sizes lead to a small uneven mass distribution, in turn yielding a small asymmetry in mass through the longitudinal axis of the circle or cylinder, producing this deflection. Due to this, the mass geometry was modified to be equivalent to the 1 kg cylindrical mass, as shown in figure 5.2(c). Upon building, it was shown when removing the blocks and attaching the ribbon centrally to the mass that the suspension deflected centrally, as shown in figure 5.2(d).

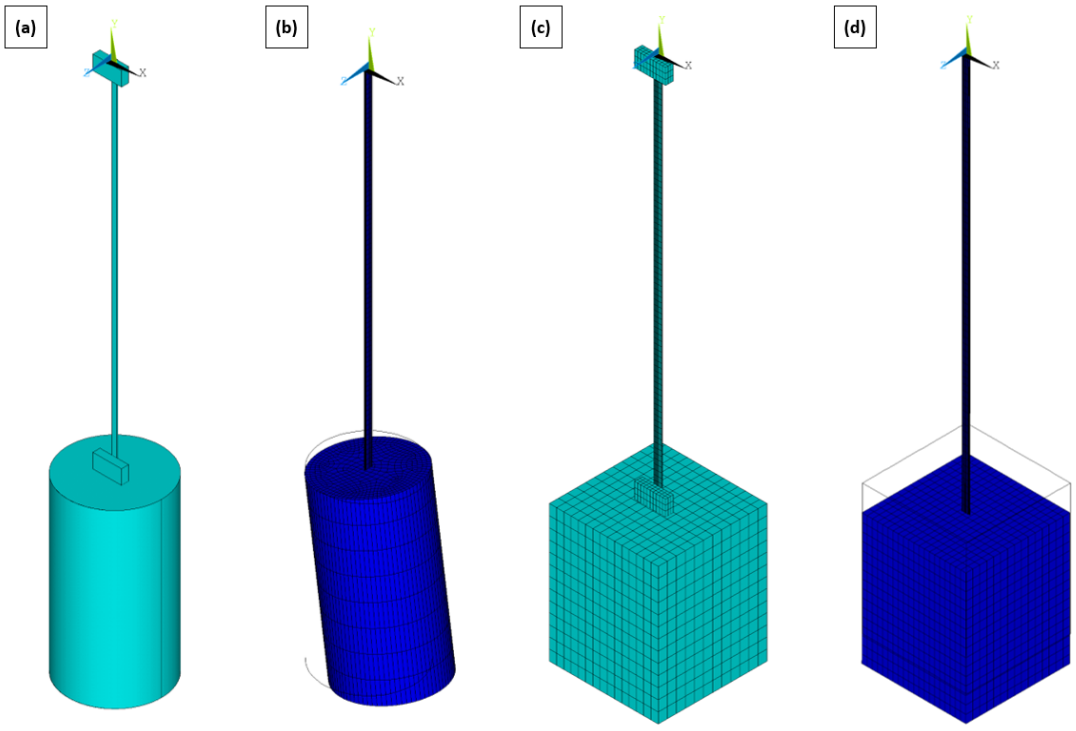

Figure 5.2: (a) Suspension volumetric geometry. (b) Symmetrical geometry with blocks removed and after loading, noting the deflection result. (c) Suspension model (a) with the circular mass replaced by an equivalent cubic mass, displaying solid elements. (d) Resultant straight deflection of cubic mass after loading, displaying solid elements. The offsets have been exaggerated for visual effect.

As is good practice with FEA modelling, one should always compare simple FEA models to theory where possible as a common-sense check. With a theoretical extension at the bottom of

the loaded ribbon of 7.19 µm, the model matched theory to within 0.03 %. The theoretical stress in the ribbon was 3.74 MPa, matching the model stress value to within 0.04 %.

Furthermore, one should also always check models for convergence. Convergence is discussed in more detail in section 5.5.3.1. All models in this thesis were checked for convergence.

5.3.1 Bottom block to mass bond tolerance

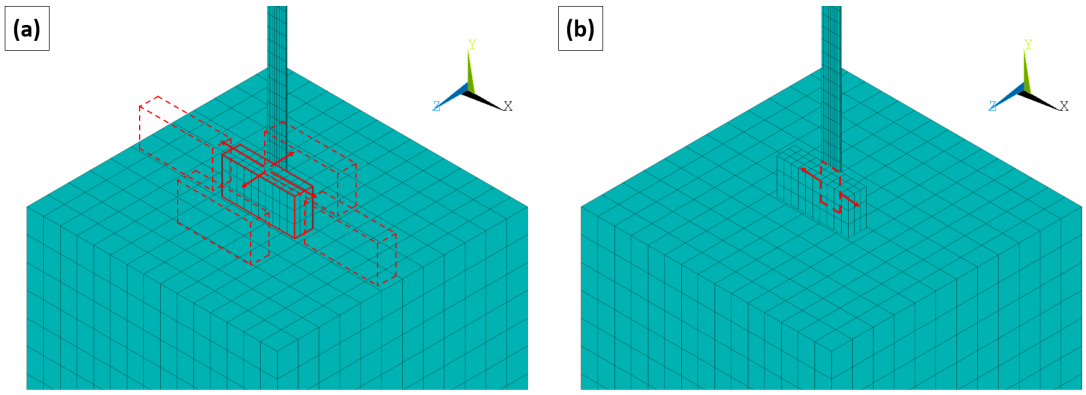

Figure 5.3: (a) Exaggerated illustration off-setting the bottom block to top of the mass bond due to misalignment in ±X and ±Z directions. (b) Illustration of off-setting the ribbon to bottom block bond due to misalignment in ±X direction.

With an ideal model built around the central axis of the ribbon including top and bottom blocks, it is possible to determine how initial misalignments due to machining tolerances or alignment errors, may affect stress within the model. Since the suspension ribbon is cut using the new LML laser technique, a conservative estimate of the maximum stress tolerated was taken from the weakest *control new* sample set, recorded in section 3.5.1 of chapter 3, at 33 MPa. With no misalignment built into the model, the suspension hangs centrally with a stress in the ribbon of 3.6 MPa, only 11 % of this conservative stress limit.

0.25 mm displacement direction	Stress (MPa)	Stress % of maximum
+Z	17.4	53 %
−Z	15.9	48 %
+X	5.9	18 %
−X	5.9	18 %
+Z +X	18.5	56 %
−Z +X	16.9	51 %
+Z −X	18.5	56 %
−Z −X	16.9	51 %

Table 5.1: Results of misalignment at bottom block to top of the mass bond.

The first investigation related to misalignment of the bottom block to top of the mass bond. This was simulated by deliberately off-setting the bottom block in the ±X and ±Z directions, as shown in figure 5.3(a). The ribbon and top block were kept centrally bonded to the bottom block at all times.

Conservative tolerancing estimates from the bonding jig, shown in section 4.4.11.1 of chapter 4, led to applying offsets of 0.25 mm in ±X and ±Z directions individually and together. The results of the corresponding misalignment are shown in table 5.1. The results are discussed in the following section.

5.3.2 Ribbon to bottom block bond tolerance

Displacement direction and amount	Stress (MPa)	Stress % of maximum
3 mm ±X	19.5	59 %
7.5 mm ±X	42.7	129 %

Table 5.2: Results of misalignment at the ribbon to bottom block bond.

The second investigation related to the ribbon and bottom block bond. This was offset in the ±X direction to represent a bonding misalignment at this interface, as shown in 5.3(b). Logically, the suspension should have lesser tolerance for misalignment in the ±X direction due to the fact the ribbon is thicker, and therefore stiffer, in this plane. As shown in table 5.2, excessive misalignment was applied, including an amount where the maximum stress was exceeded to understand how much offset was acceptable.

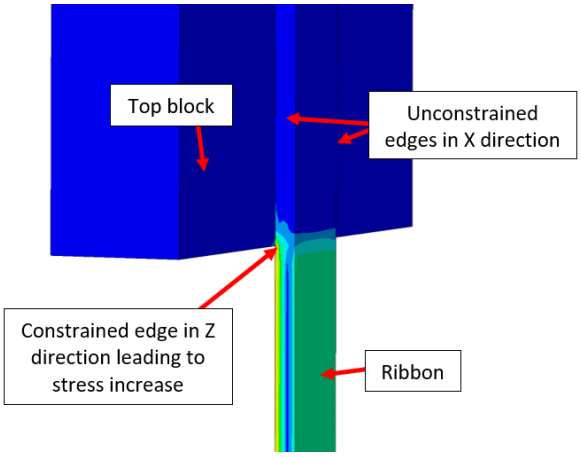

Figure 5.4: Stress contour plot focusing on top block to ribbon interface, highlighting areas of higher stress generated when the suspension is displaced in the ±Z direction.

The results from table 5.1 and table 5.2 appear to show these specific misalignments in the Z direction to be of greater importance than the X direction. This appears counter-intuitive

where one would expect the thinner, less stiff Z direction to be more compliant and the thicker X direction to be stiffer where small misalignments would significantly increase stress. The sensitivity in stress for misalignment in the Z direction arises from the fact the ribbon is attached one side to the top block, constraining the ribbon in this direction, as shown in figure 5.4. Stress is generated near the top of the ribbon interfacing with the bottom corner of the top block. This effect does not come into play when moving the block purely in the ±X direction as the ribbon is able to expand or contract in this direction freely.

5.3.3 Block positioning tolerance

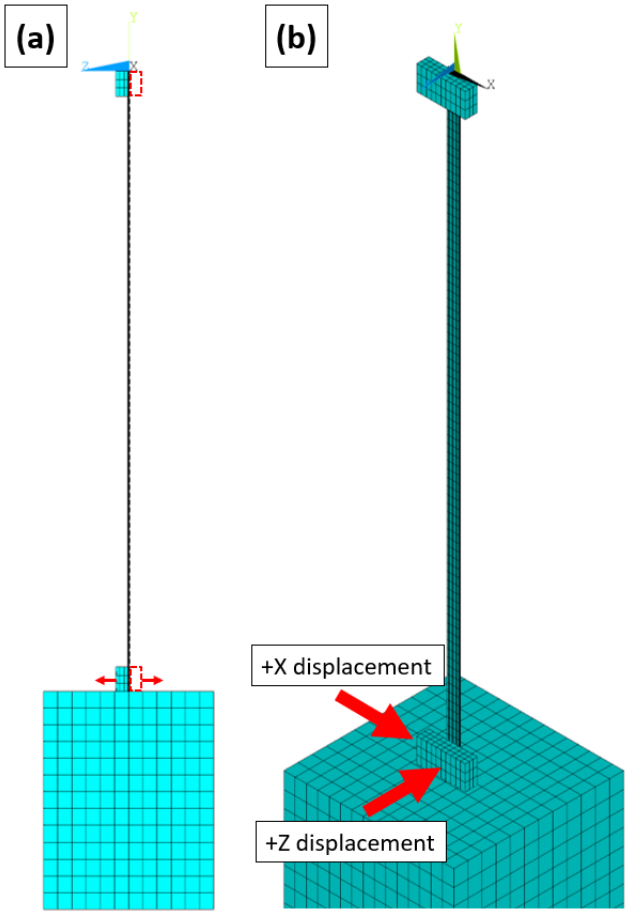

Figure 5.5: (a) Illustration of alternate block positions tried in combination with off-setting the bottom block to top of the mass bond in ±Z direction. (b) Illustration of displacing the suspension in a combination of ±X and ±Z directions.

A third investigation looked at changing the bonding location of the top and bottom blocks, as shown in figure 5.5(a). This was done by attaching the blocks to opposite sides of the ribbon and also incorporating a bonding misalignment in the ±Z direction of 0.25 mm. The results of the corresponding misalignment are shown in table 5.3.

Placing the blocks on opposite side increases the stress by 20 % compared to having blocks

on both sides. This is due to the asymmetry of mass forcing the mass to bend in opposite direction to where the bottom block is placed. This resultant displacement compresses the top of the ribbon further into the edge of the top block, exacerbating the effect shown in figure 5.4. Furthermore, this would lead to impracticalities for designing a bonding jig. Combining results from table 5.1 and table 5.3, one sees that misalignment of the bottom block on the top of the mass by up to 0.25 mm in the $\pm$Z direction has the greatest single impact, with stresses at around half the maximum tolerance.

Configuration and displacement direction	Stress (MPa)	Stress % of maximum
Blocks on same side	3.7	11 %
Blocks on opposing sides	4.5	14 %
Blocks on opposing sides +Z	16.5	50 %
Blocks on opposing sides −Z	16.7	51 %

Table 5.3: Results of blocks placed on opposing sides combined with and without a 0.25 mm offset in $\pm$Z direction.

5.3.4 Suspension displacement tolerance

Displacement and direction	Stress (MPa)	Stress % of maximum
1 mm +Z	5.9	18 %
3 mm +Z	10.0	30 %
6 mm +Z	16.2	47 %
3 mm +X	60.9	185 %
0.70 mm +X	17.3	52 %
0.74 mm +X	18.1	55 %
0.89 mm +X	20.9	83 %
2 mm +Z 0.70 mm +Z	18.9	57 %
2 mm +Z 0.70 mm −Z	18.9	57 %
2 mm −Z 0.70 mm +Z	21.0	64 %
2 mm +Z 0.70 mm +Z	21.0	64 %
2 mm +Z 0.74 mm +Z	19.5	59 %

Table 5.4: Results of successive displacement of centrally aligned, symmetric suspension.

The final investigation assumed a symmetrically bonded, centrally aligned suspension that is deliberately displaced at the bottom of the ribbon in the ±X and ±Z directions, as shown in figure 5.5(b). This represents the suspension swinging due to external excitation or accidental contact with the suspension while hanging. The results are shown in table 5.4 in order of a successive displacement process applying reasonably expected displacements.

The successive displacement process from table 5.4 provided values to help inform mass stop positions in both ±X and ±Z directions for the suspension support structure first discussed in section 4.3 of chapter 4. As expected, pushing the ribbon in the stiffer X direction leads to larger stresses compared to the Z direction. This tolerancing work informed the bonding jig design and assembly procedures discussed in section 4.4.11.1 of chapter 4.

The analysis in this section was not repeated for the real suspension ribbon geometry of 3.5 mm wide by 775 µm thick, which for reference, has a slightly reduced ribbon stress of 3.6 MPa. This was justified for the following reasons. Firstly, any bottom block to top of the mass misalignment, as shown in figure 5.3(a), would arise from the bottom block alignment plate discussed in section 4.4.11.2 of chapter 4. This was known to have a machining tolerance of ±100 µm, well below the displacement applied in table 5.1. The real ribbon is thinner in the stiff direction and so it is reasonable to assume the stress will be reduced for an equal amount of displacement. However, the real ribbon is thicker in the soft Z direction which will increase the stress in the ribbon for this particular misalignment and so a model was created with the real ribbon geometry to check how this affects the stress in the ±Z misalignment. This was found to increase the stress by only 0.4 MPa to 17.8 MPa, 54 % of the breaking stress limit. This small increase can be explained by the trade off between the stiffness increase from ribbon thickness against the relative displacement offset effect which is lessened for a thicker ribbon since the suspension is designed to align with the central axis of the ribbon.

The bond misalignment shown in figure 5.3(b) and table 5.2 was measured and modelled directly in FEA with real bond offsets, angular misalignments and the real suspension geometry in order to determine the maximum stress in the suspension. Each of the 4 suspensions produced in 4.4.11.1 of chapter 4 underwent this analysis in order to determine the lowest stress suspension with the best chance of survival when hung.

The blocks were never placed on opposing sides due to the analysis conducted from figure 5.5(a) and table 5.3 so this required no modelling. Furthermore, as was shown from the results above, the any ±Z misalignment will not significantly increase the stress in the real suspension geometry.

Finally, from figure 5.5(b) and table 5.4 results, one notes the excessive ±X displacement required to over-stress the suspension. With a thinner ribbon width, this allows for greater ±X displacement of the mass before breaking. The suspension is relatively insensitive in the ±Z direction when displacing the mass and can take substantial deflection before approaching even half the conservative breaking stress. For the real suspension geometry, as can be seen

throughout chapter 4 suspension builds, the mass stops were never >0.5 mm away from the mass when hanging, providing confidence that the real suspension geometry would never be over-stressed.

This body of work confirmed that machining tolerancing and alignment was not a concern for the proposed bonding jig design or bottom block alignment plate shown in section 4.4.11.1 of chapter 4.

5.4 Modelling of silicon's crystal axes

As discussed in section 2.3.2 of chapter 2, silicon is a crystal consisting of 3 primary axes. Each of these three axes have different material properties, such as Young's modulus and Poisson's ratio, unique to each axis. The aim of this section is to understand how to model these axes within the FEA environment correctly, as well as determining if this technique is necessary for accurate results, since Young's modulus and Poisson's ratio, in particular, will affect deflections and frequencies of the model. This will help to make models as efficient as possible by removing unnecessary extra computation where possible.

Work conducted by Hopcroft *et al.* neatly solves this problem, providing a useful guide on how to correctly model the axes of silicon and apply their properties to an FEA model via the use of a simple stiffness matrix [174].

The stiffness matrix encoding the Young's modulus and Poisson's ratio for a (100) silicon wafer is given as [174]:

$$c_{(100)} = \begin{bmatrix} 194.5 & 35.7 & 64.1 & 0 & 0 & 0 \\ 35.7 & 194.5 & 64.1 & 0 & 0 & 0 \\ 64.1 & 64.1 & 165.7 & 0 & 0 & 0 \\ 0 & 0 & 0 & 79.6 & 0 & 0 \\ 0 & 0 & 0 & 0 & 79.6 & 0 \\ 0 & 0 & 0 & 0 & 0 & 50.9 \end{bmatrix} \times 10^9\,\text{Pa} \tag{5.1}$$

As noted by Middlemiss *et al.*, ANSYS differs in its standard definition of the stiffness matrix [393]. This can be remedied by switching matrix element 44 with 66 in equation 5.1 [393]. One can then enter this into the material properties section of ANSYS.

This matrix yields the following values for each of the three axes material properties [174]:

$$\begin{gathered} E_x = E_y = 169\,\text{GPa},\ \ E_z = 130\,\text{GPa} \\ \nu_{yz} = 0.36,\ \ \nu_{zx} = 0.28,\ \ \nu_{xy} = 0.064 \end{gathered} \tag{5.2}$$

It is important to note, this matrix is only valid for a (100) wafer aligned within the FEA model along the corresponding axes of X, Y and Z. If the FEA model is to be built for any other (100) silicon rotation or wafer orientation, the matrix must also be rotated. As shown in section

3.4.1 of chapter 3, the silicon ribbons are cut along the standard <110> axes, corresponding to the X and Y axes in the diagram and will be built in the same axes for FEA work in this section. As such, no rotation or modification of the material properties matrix is required.

5.4.1 Modelling an unloaded ribbon

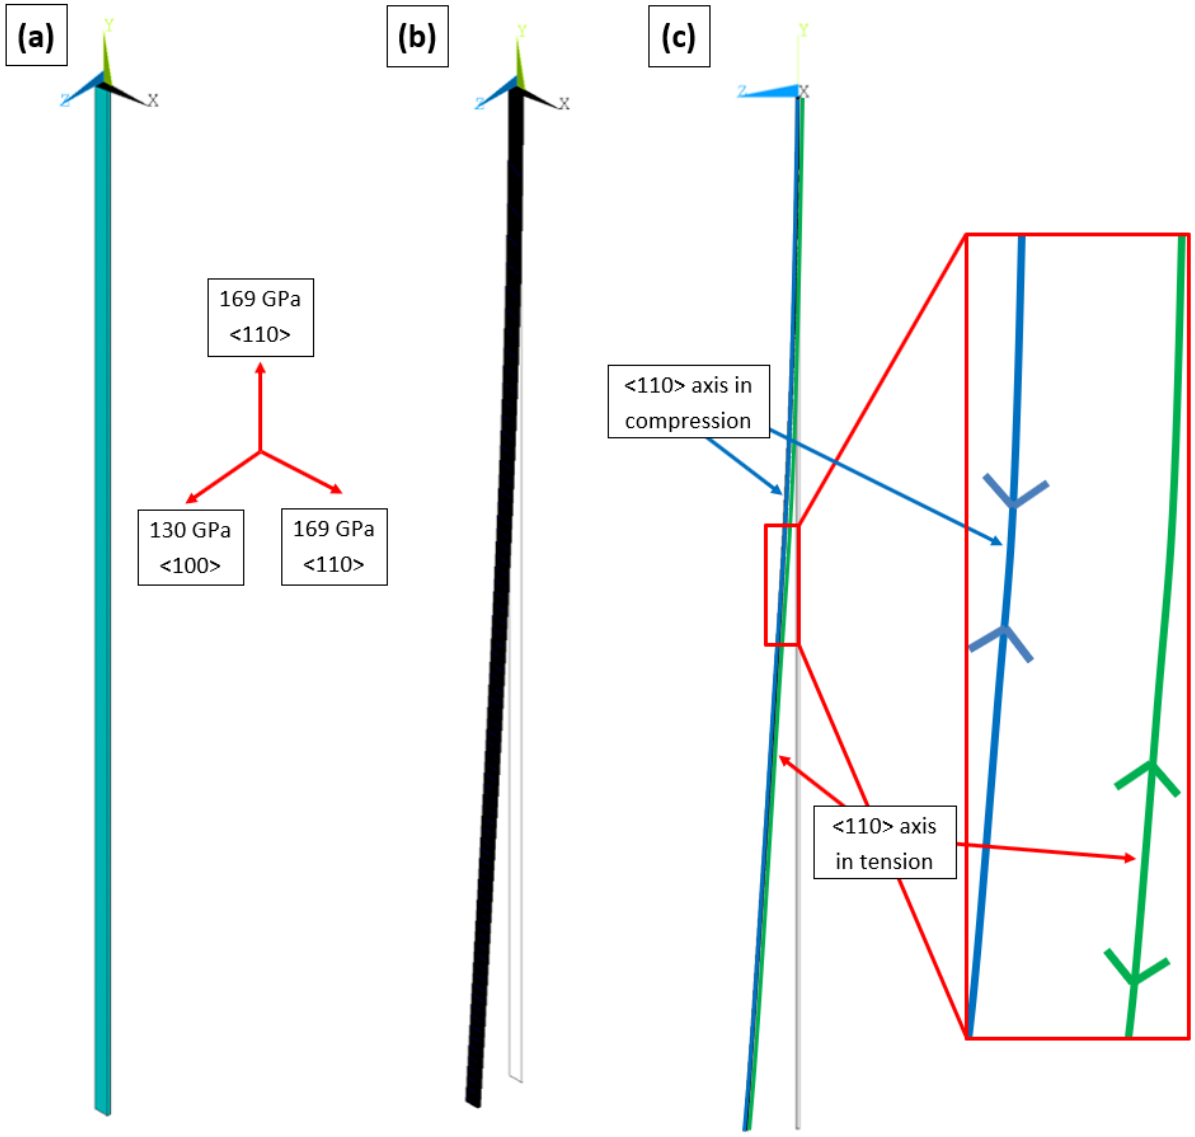

Figure 5.6: (a) Silicon ribbon geometry, crystal axes and corresponding Young's moduli used for the isotropic and anisotropic investigation. (b) First mode shape for both isotropic and anisotropic models. (c) Crystal axes involved in the bending of the first mode shape.

With anisotropic properties applied to the model, the first step was to model a simple silicon ribbon and compare resultant model frequencies to theory. Figure 5.6(a) shows the simple FEA geometry of the silicon ribbon. As discussed at the end of section 5.3, the real suspension ribbon geometry used was 3.5 mm wide by 775 µm thick and all subsequent analysis in this chapter will

use these geometric values.

In order to provide a baseline check, this was also done for an isotropic silicon ribbon. The isotropic Young's modulus and Poisson ratio was chosen as the root-mean-square (rms) value. From equation 5.2, these values are $E_{rms} = 157\,\text{GPa}$ and $\nu_{rms} = 0.27$.

The frequency of the n^{th} bending mode of a simple beam or cantilever of thickness, t, length, L, Young's modulus, E and density, ρ, is given by [394]:

$$\omega_n = (k_n L)^2 \frac{t}{2\sqrt{3}L^2} \left(\frac{E}{\rho}\right)^{\frac{1}{2}} \tag{5.3}$$

where $k_n L = 1.875$ $(n = 1)$, 4.694 $(n = 2)$, 7.853 $(n = 3)$, 10.996 $(n = 4)$ and 14.137 $(n = 5)$. For $n > 5$, this can be approximated as $k_n L = (2n-1)\frac{\pi}{2}$ [395]. Figure 5.6(b) shows the first mode shape of the silicon ribbon geometry.

Mode	Isotropic model frequency (Hz)	Theoretical frequency (Hz)	% difference
Young's modulus (GPa)		157	
1	16.46	16.45	-0.09 %
2	74.30	74.28	-0.02 %
3	103.15	103.09	-0.06 %
4	288.80	288.52	-0.10 %
5	465.18	465.55	0.08 %
6	565.90	565.69	-0.04 %
7	935.37	935.03	-0.04 %
8	1300.60	1303.01	0.19 %

Table 5.5: Isotropic silicon model mode frequencies compared to theory.

Mode	Anisotropic model frequency (Hz)	Theoretical frequency (Hz)	% difference	Theoretical frequency (Hz)	% difference
Young's modulus (GPa)	From matrix equation 5.1	130		169	
1	17.07	14.97	-12.33 %	17.06	0.04 %
2	77.07	77.07	-0.01 %	77.07	0.01 %
3	106.96	93.80	-12.30 %	106.95	0.01 %
4	299.48	262.54	-12.33 %	299.35	0.04 %
5	482.46	483.01	0.11 %	483.01	-0.11 %
6	586.81	514.76	-12.28 %	586.91	-0.02 %
7	969.93	850.84	-12.28 %	970.10	-0.02 %
8	1348.50	1351.89	0.25 %	1351.89	-0.25 %

Table 5.6: Anisotropic silicon model mode frequencies compared to theory.

The first 8 modes were analysed and the resultant frequencies for the isotropic ribbon were compared to theory as shown in table 5.5 and were shown to have excellent matching.

Similarly for the anisotropic ribbon, the first 8 modes were analysed and the resultant frequencies were compared to theory as shown in table 5.6.

Since bending of mode 1 took place in the Z direction, as shown in figure 5.6(b), one might naively assume a Young's modulus of 130 GPa for the theoretical frequency. As can be seen in table 5.6, this is an incorrect assumption. Bending involving deflection in the Z direction also took place in modes 3, 4, 6 and 7. Instead, a Young's modulus of 169 GPa should be used since bending in this mode, took place in the form of tension and compression along the <110> crystal axes, as shown in figure 5.6(c).

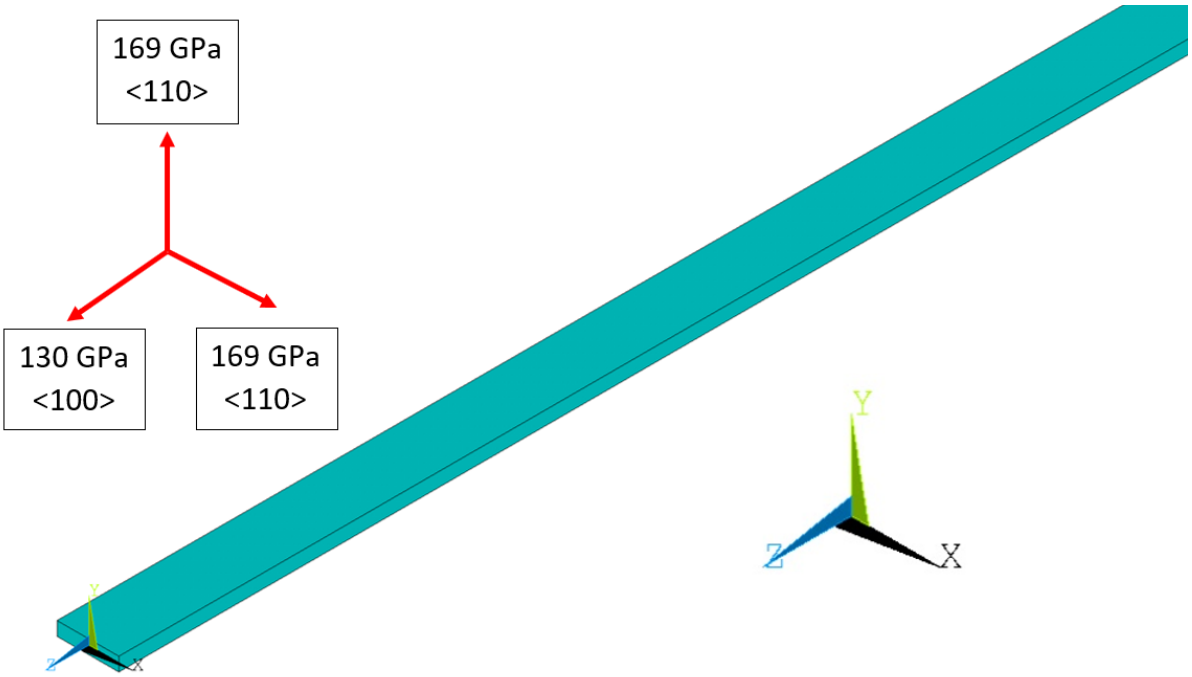

Figure 5.7: Modified ribbon geometry with same stiffness matrix providing a Young's modulues of 130 GPa along the ribbon's length.

To confirm the results, another model was generated with the same anisotropic matrix, however the ribbon geometry was orientated so the 250 mm length of the ribbon ran along the Z axis, as shown in figure 5.7. This meant the faces along the length of the ribbon should parallel to the <100> axis. These should be in tension and compression for the first mode and so the Young's modulus used to calculate the corresponding frequencies should be 130 GPa. This was found to be correct and the model matched theory, with the largest discrepancy at 0.47 % across all modes.

5.4.2 Modelling a loaded ribbon and determining bending points

The next step is to apply tension to the silicon ribbon. This can be done in ANSYS by modelling a MASS21 mathematical point mass element, which has no physical volume or geometry within the model. This 1 kg mass was attached onto the bottom of the ribbon, central to the cross-section.

Equation 5.3 is not valid for a ribbon under tension. Instead, one can use the equations 2.9, 2.10 and 2.11 for determining the frequency as detailed in section 2.4.1 of chapter 2.

Any pendulum material with a finite stiffness will have some effective length, shorter than the material length of the pendulum, as shown in [144] [396]. Since fibres and ribbons have a finite stiffness they do not bend at the attachment point. The "bending point" of a fibre or ribbon therefore needs to be defined, and this is commonly taken as intersection point between the vertical and the linear extrapolation back from the lower section of the fibre, where no bending occurs (and therefore the fibre is straight in this region). A schematic of this is shown in figure 5.8. This is the same technique used in section 4.1 of [144].

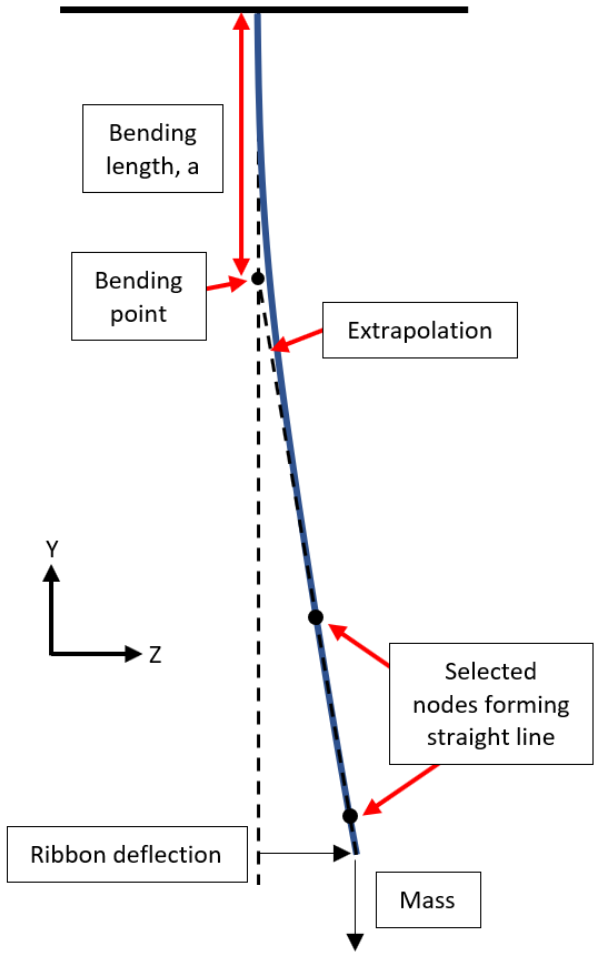

Figure 5.8: Example of bending length extrapolation from a deflected model in ANSYS.

One can determine the bending point in ANSYS using the methodology shown in section

4.1.3 of [144]. This technique displaces a pendulum under tension by a fixed amount in the X or Z direction. By then extracting nodal positions along the length of the displaced ribbon, one can extrapolate the bottom straight line portion back to the vertical to calculate the bending point, as shown in figure 5.8. One can determine this analytically using the following equation [396]:

$$a = \sqrt{\frac{EI_{(x/z)}}{T}} \tag{5.4}$$

where a is defined as the bending point, which when subtracted from the material pendulum length, gives the effective length of the pendulum which is used for pendulum frequency calculation. One notes the effective length of a pendulum is determined only by the cross-sectional area, tension and Young's modulus. From equation 5.4, one can see that a ribbon with a non-square cross-section will have different bending points, depending on the bending direction, meaning care is needed to extract parameters like mode frequencies correctly.

Parameter	Anisotropic		Isotropic	
Pendulum length, L (m)	0.25	1.25	0.25	1.25
Young's modulus (GPa)	169		157	
Pendulum X thicker direction - mode 2				
Bending point, a (m)	0.218		0.211	
Effective Length, L_{eff} (m)	0.032	1.032	0.039	1.039
L/a	1.145	5.723	1.188	5.938
Theoretical frequency (Hz)	7.890	0.540	6.312	0.536
FEA frequency (Hz)	1.860	0.491	1.816	0.489
% difference	-76 %	-9 %	-71 %	-9 %
Pendulum Z thinner direction - mode 1				
Bending point, a (m)	0.048		0.047	
Effective Length, L_{eff} (m)	0.202	1.202	0.203	1.203
L/a	5.169	25.847	5.363	26.816
Theoretical frequency (Hz)	1.236	0.464	1.225	0.463
FEA frequency (Hz)	1.110	0.455	1.106	0.455
% difference	-10 %	-2 %	-10 %	-2 %

Table 5.7: Loaded anisotropic and isotropic silicon ribbons, both 0.25 m and 1.25 m pendulum lengths, comparing theory to ANSYS results for the frequency of the first and second modes. Red texts highlights significant mismatch between theory and modelling.

Analysis was conducted on 4 models to investigate effects on anisotropy and pendulum length. Modal analysis was conducted on the first 25 modal frequencies. The first mode involved bending in the thinner Z direction, the second mode in the thicker X direction. These 2 modal frequency results are recorded in table 5.7.

As noted by Cagnoli *et al*., equations 2.9 and 5.4, are valid assuming $\frac{L}{a} \gg 1$ [396]. As can be seen in table 5.7, for the stiffer X direction $\frac{L}{a} \not\gg 1$ and so the theory does not match the ANSYS results, however in the Z direction this theory is valid and yields an acceptable error between theory and FEA. The reason for this discrepancy is the theory negates higher order terms for simplicity. These higher order terms become significant for thicker bending directions with longer bending lengths. ANSYS analysis includes non-linear effects through stiffness modelling, as such one must include higher order terms in the theory for the purposes of comparison. Alternatively, one can make the suspension longer so that $\frac{L}{a} \gg 1$ for both X and Z directions. One can see in table 5.7 that both anisotropic and isotropic ribbons are in agreement with this simplified theory when the ribbon is extended to 1.25 m long, allowing $\frac{L}{a}$ to always fulfil the $\frac{L}{a} \gg 1$ criteria needed for the approximation to be valid.

5.4.3 Comparison of anisotropic to isotropic models

Furthermore, table 5.7 shows that as the length of the ribbon to bending thickness scales, the effect of anisotropy is significantly reduced, most evident in the Z direction results. Even for short, 0.25 m, suspensions the difference in frequency between anisotropic and isotropic models was minimal when the $\frac{L}{a} \gg 1$ condition is met.

Finally, one notes that anisotropic and isotropic models produce similar levels of accuracy when compared to theory. While the isotropic models match the respective theory slightly better than anisotropic models in the thicker X direction, this is likely due to over-simplification of the theory. Both ANSYS models will take into account higher order terms and so care should be taken when drawing conclusions that isotropic is "more accurate" as one approaches the limit of the simplified theory.

To verify using isotropic material properties, the next 7 higher mode frequencies, to include the violin and bounce modes, as discussed in sections 2.4.2 and 2.4.3 of chapter 2 respectively, were compared in table 5.8.

As can be seen in table 5.8, anisotropic models do not deviate significantly compared to isotropic models at these higher order modes of interest. This is fundamentally due to the fact that pendulum, bounce and violin mode deformations only involve bending in the <110> axes as discussed in section 5.4.1. For efficiency purposes, it was deemed future work would utilise the simpler isotropic properties of silicon. This also avoids issues with model alignment in the XYZ co-ordinate system.

Mode	Anisotropic	Isotropic	% difference
Pendulum length, L (m)	0.25	0.25	
Violin 1 thinner Z direction	112.8	111.0	1.6 %
Bounce	206.1	199.3	3.3 %
Violin 2 thinner Z direction	292.9	285.8	2.4 %
Violin 1 thicker X direction	348.3	336.6	3.4 %
Violin 3 thinner Z direction	561.5	545.3	2.9 %
Violin 4 thinner Z direction	923.3	894.5	3.1 %
Violin 2 thicker X direction	1104.8	1066.4	3.5 %
Pendulum length, L (m)	1.25	1.25	
Violin 1 thinner Z direction	16.6	16.5	0.2 %
Violin 1 thicker X direction	21.6	21.3	1.4 %
Violin 2 thinner Z direction	33.8	33.7	0.3 %
Violin 3 thinner Z direction	52.5	52.2	0.6 %
Violin 2 thicker X direction	54.8	53.5	2.3 %
Violin 4 thinner Z direction	73.1	72.5	0.8 %
Bounce	95.5	92.1	3.6 %

Table 5.8: Anisotropic and isotropic modal frequency results for modes 3 to 9 for both 0.25 m and 1.25 m pendulum lengths.

5.5 Modelling hydroxy-catalysis bonds

The work in this section is interested in developing direct modelling techniques for HC bonds between silicon suspension elements. With a successful suspension prototype designed and built in chapter 4, one can model the suspension thermal noise. As discussed in chapter 2, thermal noise from HC bonds may be a limiting factor for 3G gravitational wave detectors. Therefore, one can model the suspension thermal noise from the prototype suspension and scale this to an ET-like detector. It is important to be able to directly model these bonds in order to understand their contribution to the overall thermal noise of a silicon suspension. ANSYS has the ability to determine the strain energy in a volumetric region of a model. By dividing this by the total strain energy of the whole geometry, one can determine the ratio of energy in the bond to the whole suspension, in turn calculating the loss contribution from the HC bond and the resultant bond noise. The work in this section creates a new, superior model for HC bonds with the most accurate estimation of the HC bond contribution to suspension thermal noise to date.

5.5.1 Indirect modelling of HC bonds

Previous work involves an indirect modelling technique of HC bonds developed by Elliffe [152] [397]. Elliffe identified issues with FEA modelling of extremely thin elements between larger substrates. Since HC bonds are on the order of 10s to 100s of nm thick, and due to the relative weakness of bond strength compared to the bulk material, bond areas are typically on the order of cm^2 [150] [152]. To directly model these bonds would require the generation of extremely thin elements with excessively high aspect ratios.

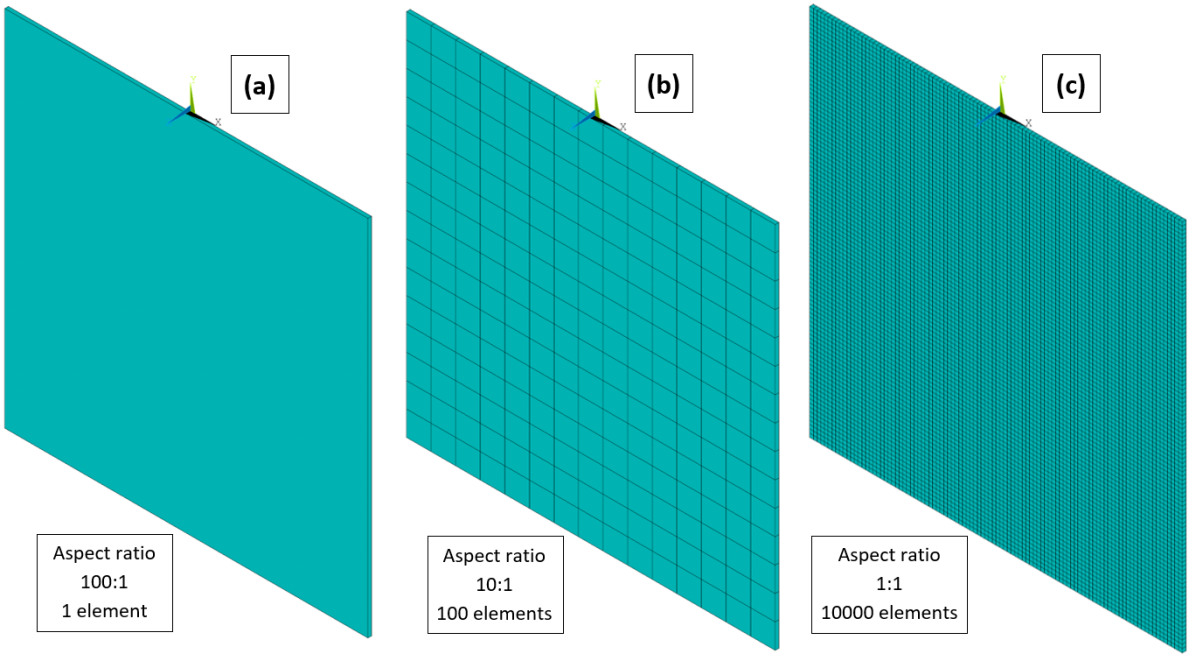

Figure 5.9: Illustration of the problem meshing thin geometries in FEA.

To take a less extreme example for visual clarity, figure 5.9(a) shows a geometry 1 unit by 1 unit square, with a thickness 100 times less in the X direction. The aspect ratio is 100:1 if one meshes this geometry as 1 large element. To improve this, one can mesh the Y and Z directions by splitting the geometry into 10 $\times$ 10 elements, as shown in figure 5.9(b), keeping 1 element thick in the X direction, reducing the aspect ratio to 10:1 with 100 elements. If one wanted to reduce this further, to a ratio of 1:1, this would increase the number of elements to 10000, as shown in figure 5.9(c). One can see how the number of elements grows as the square when trying to improve the aspect ratio. Considering HC bonds are on the order of nm, for a 1 cm $\times$ 1 cm bond area, figure 5.9(a) would have an aspect ratio of 10^6:1. To reach an acceptable aspect ratio, such as in figure 5.9(b), would require 10^{12} elements.

Extremely high aspect ratios can lead to errors within the FEA computation, unphysical results or even a failure of the model to solve. FEA software like ANSYS typically limits this aspect ratio to 20:1 for solid elements [392] [398]. Aspect ratios above this can be modelled but come with an associated warning to check the accuracy of the results.

Instead, Elliffe developed a technique whereby one models a greater bond thickness, on the order of 10s or 100s of μm, at a variety of thicknesses. One can then plot the bond energy ratio to bond thickness and extrapolate this data down to the nm scale. This relies on extrapolating orders of magnitude outside the modelled range of data. It also relies on the assumption that the extrapolation is linear throughout the extrapolation range. Furthermore, this method was only applied to exceptionally simple geometries and was never applied to complex suspension geometries. The work conducted in this section shows that bond deformation does not conform to these simplistic assumptions and develops a more accurate method for cases where extrapolation breaks down.

5.5.2 Previous direct modelling of HC bonds

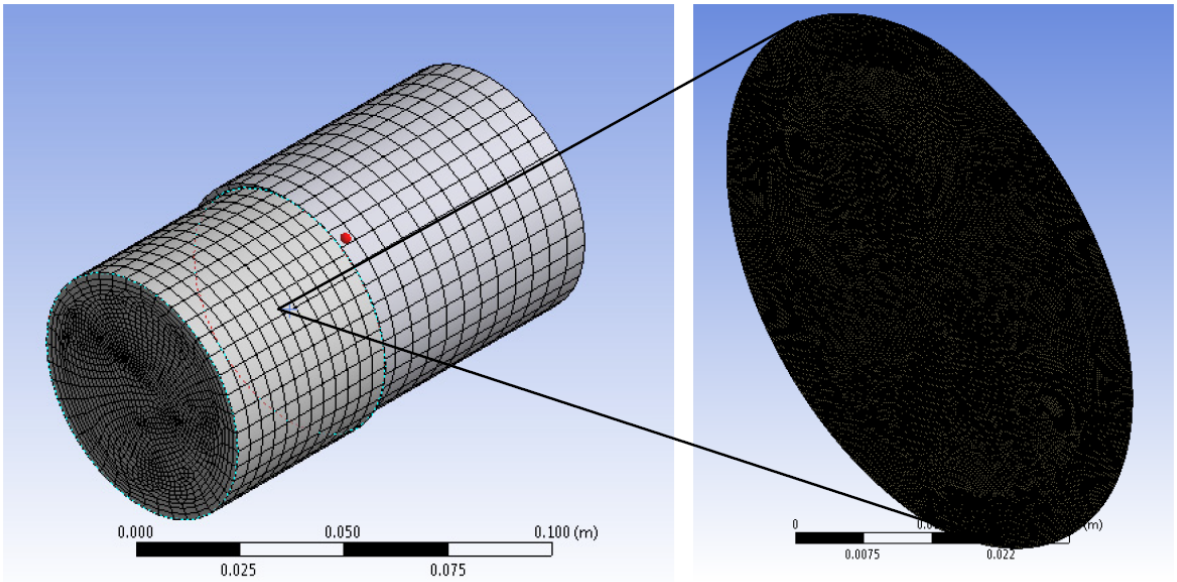

Figure 5.10: ANSYS Workbench model of 50 mm and 70 mm silicon cylinders connected by a HC bond. Image obtained from [150].

Previous work has developed direct modelling of HC bonds between silica suspension elements [364] [391]. While this work was in agreement with the extrapolation technique, this was still only for simplistic geometries, so one must still develop a technique to directly model bonds in real, complex suspension geometries, where stresses and deformations are more complex. This also allows for the investigation of how real bond geometries deform under different modal perturbations. With the advent of more powerful computing, direct modelling for more complex geometries is possible, allowing a greater array of results, and without the need for multiple extrapolation models and their inherent assumptions.

Work by Haughian demonstrated direct modelling of order 100 nm thick bonds between two large silicon cylinders in ANSYS Workbench [150]. Haughian first modelled a solid cylinder,

120 mm in length. A second model was then created, split into 50 mm and 70 mm cylinders connected directly together. These cylinders were slightly offset to reflect the same offset in the real experimental sample. A third and final model, included a bond between these two cylinders, as shown in figure 5.10. This allowed for comparison between each model to investigate the effects of the bonding algorithm as well as the introduction of the bond itself on the modal frequencies of the system. It should be noted however, this work still relied on extrapolation for the thinnest bond geometries [150]. This work was recreated in ANSYS Mechanical APDL to verify the results are the same as the direct modelling techniques used subsequently. The work replicating Haughian's models matched all modal frequencies and energy ratio values to within 0.2 %.

During the course of this verification, it was noted that the same APDL code run on different specific computers produced slightly different results. For the same circular geometry and mesh, on the same version of ANSYS, on the same operating system, another computer might produce up to 18 % more elements. This anomaly was traced back to CPU architecture. It appears that somewhere between Intel's 3rd generation (circa 2012) and 6th generation (circa 2015) processors the CPU architecture changed, in turn, changing the way the meshing algorithm selects the optimal mesh for a given set of meshing controls [399] [400]. Ultimately, the results produced from these different meshes only varied at the 5th or 6th decimal place, but this is worth noting when running APDL code on different machines for comparative purposes.

5.5.3 Modelling HC bonds between silicon substrates

The intention of the work in this section is to directly model HC bonds between silicon suspension elements and compare to the results of using Elliffe's extrapolation technique, applied to more complex suspension geometries, in ANSYS.

The same single ribbon 1 kg suspension geometry from the HC bonded suspension in chapter 4 was used. The model bond regions were also modified to reflect the real suspension, where a 0.5 mm gap existed at the top of the top block and bottom of the bottom block, as discussed in section 4.4.11.1 of chapter 4. This created two bond areas at the top and bottom of the ribbon that were 9.5 mm high by 3.5 mm wide.

As discussed in section 4.4.11.1 of chapter 4, silicon can not be HC bonded without a sufficient oxide layer [190]. The initial work in this section explores three bonding scenarios; no oxidation, oxidation of the silicon ear only and oxidation of both the silicon ear and silicon ribbon. This produces HC bond thicknesses of 61 nm, 211 nm and 361 nm respectively.

The HC bonds have unique material properties which must be incorporated into the model. These are E_{bond} = 18.5 GPa, ν_{bond} = 0.17, density = 2202 kg m^{-3} as defined by [401].

5.5.3.1 Modelling approach

As discussed in section 5.5, one typically has two options when modelling HC bonds in FEA: using elements with aspect ratios greater than the ANSYS preferred limit of 20:1 or using an extremely high number of elements. This work used a compromise, deploying both options, using the convergence curves to investigate the optimal combination. This can be probed by solving a model with this single element and recording the parameter of interest, in this case the bond energy ratio. One can then increase the number of elements representing the bond area and record the resultant bond energy ratio. A classic convergence curve should follow a dog leg shape, at the point where the model results settle to values that are almost constant, independent of mesh choice and therefore element count. This allows both for the determination of a valid result, as well as selecting the most efficient number of elements near the apex of the curve, as shown in figure 5.11. By selecting the lowest element density in the region where the results are constant, one selects the most efficient, accurate model.

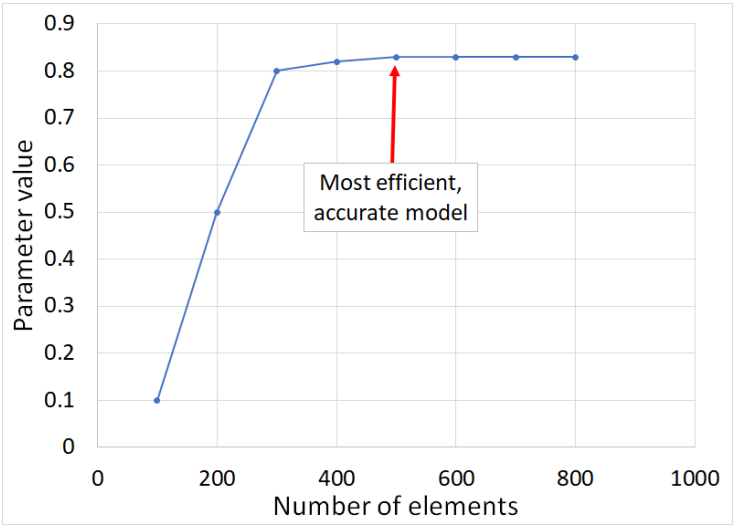

Figure 5.11: Example of a convergence curve used to illustrate a converged model and the most efficient model to select in the region where results are constant.

5.5.3.2 Modelling the first HC bond

The first step was to create a silicon suspension with the HC bond only at the top in order to find the element density where models converged for each of the parameters of interest, independently. The parameters selected were the modal frequency, extension and bond energy ratio. For efficiency, only the first 6 modes of the model would be analysed.

When attempting to solve for modal analysis, it was noted the first models tested did not converge for either 61 nm, 211 nm or 361 nm bond thicknesses. While frequency and extension appeared to come close to converging, the bond energy ratio appeared to keep rising. An example of the convergence curve for a 211 nm bond's energy ratio for mode 1 is shown in figure 5.12(b).

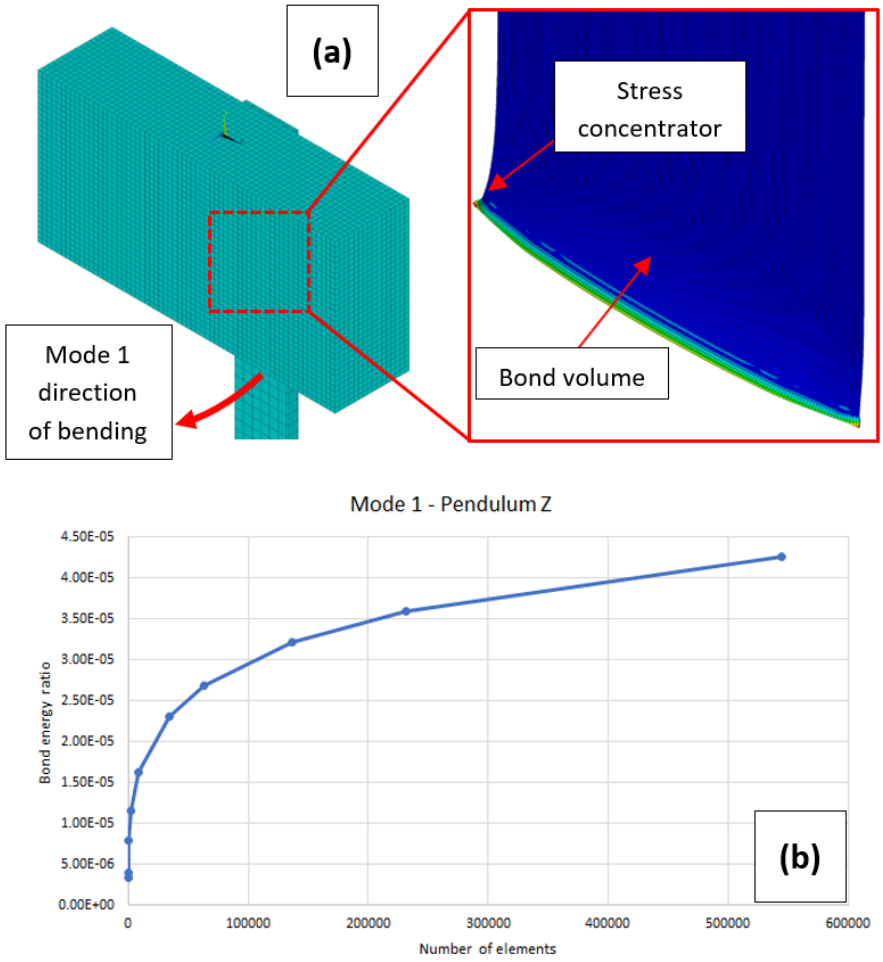

Figure 5.12: (a) Stress concentrator corner effect in a 211 nm thick bond geometry caused when the bond is squeezed between the ribbon and the top silicon block for mode 1. (b) Resultant non-convergence curve for the bond energy ratio for mode 1.

The issue was traced to the bonding algorithm used to connect elements together. Work conducted by Haughian previously used an augmented Lagrange algorithm [150]. This contacting method works via a series of iterations that monitor the interaction between the solid body surface elements. Each iteration changes pressure and stress parameters so that the final separation between these surface elements is smaller than some set tolerance. Heavily distorted or deformed meshes can lead to the requirement of additional iterations [402]. While ANSYS recommends this method to avoid over-constraining, it was found unsuitable for the specific modelling required, as penetration between the corner of the HC bond and the top silicon block produced erroneous stress concentrator artefacts as shown in figure 5.12(a). This bond geometry in the inset of 5.12(a) arises from the softer bond being squeezed between the ribbon and the top silicon block during mode 1, which is the pendulum Z direction. The stress in this model would increase with the number of elements due to this corner effect, known as a stress singularity, and so the model will never converge [403].

Instead the pure Lagrange multiplier contact method was used to attach geometries together.

This method ensures zero penetration between contacting geometries but comes at increased computational cost [402].

It should also be noted, for all models that included bonds, the adjacent ribbon and blocks are more densely meshed compared to models in the previous sections. This is to minimise errors that arise from contacting densely meshed to sparsely meshed geometries together, for example, contacting multiple small elements from the bond to one relatively large element from the ribbon [402].

5.5.3.3 Bond convergence in the XY plane

For reference, the corresponding modes and mode shapes are shown in figures 5.13(a)-(f). For the 3 bond thicknesses, the Lagrange contact-based models quickly converged in frequency in all 6 modes, a selection of which is shown in figures 5.14(a)-(c). The convergence of the model extension is shown in figures 5.14(d)-(f). Figures 5.14(g)-(i) show the corresponding convergence curves the bond energy ratio for each mode shown in figures 5.14(a)-(c).

Generally, the most computationally efficient model was found to have an element area size value, herein referred to as the "asize", of 15.625 µm, equating to 136192 elements in the bond. The value is generated by using the "ESIZE" ANSYS command, applied to an area, which sets the default element edge length based off the user-defined asize value. ANSYS will try to make a square element in the XY direction defined by the asize dimension, although if the area dimension is not divisible by this defined size the command will be overridden to adjust the elements to fit. Since the 9.5 mm by 3.5 mm was not exactly divisible by the 15.625 µm asize value in both directions, the resultant bond element dimensions were 15.625 µm in the Y dimension and 17.157 µm in the X direction. Taking the larger of the two dimensions, for a 61 nm, 211 nm and 361 nm thick bond, this corresponds to an aspect ratio of 281:1, 81:1 and 48:1 respectively. This naturally generated aspect ratio warnings in the ANSYS software since the ratio was greater than 20:1. However, comparing the convergence of the extension and frequency in models with bonds to models with no bonds, verified these models globally.

Furthermore, we see clear convergence of the HC bonds from figure 5.14. Comparing the most efficient 15.625 µm asize model to the highest element model, with an asize value of 7.8125 µm, the largest difference in parameter results across all 6 modes, all bond thicknesses and across frequency, extension and bond energy ratio was 1.38 %. This corresponded to the bond energy ratio error for a 61 nm thick bond for modes 1, 4 and 6. All of these modes involve pendulum motion in the Z direction, implying small errors may still arise from corner effect along the bond-top ribbon block interface as shown in figure 5.12(a). The next largest difference in parameter results, at 0.44 %, was the bond energy ratio error for a 211 nm thick bond for modes 1 and 4, providing further evidence of this edge effect, which is reduced as the bonds are made thicker, relative to the block and ribbon.

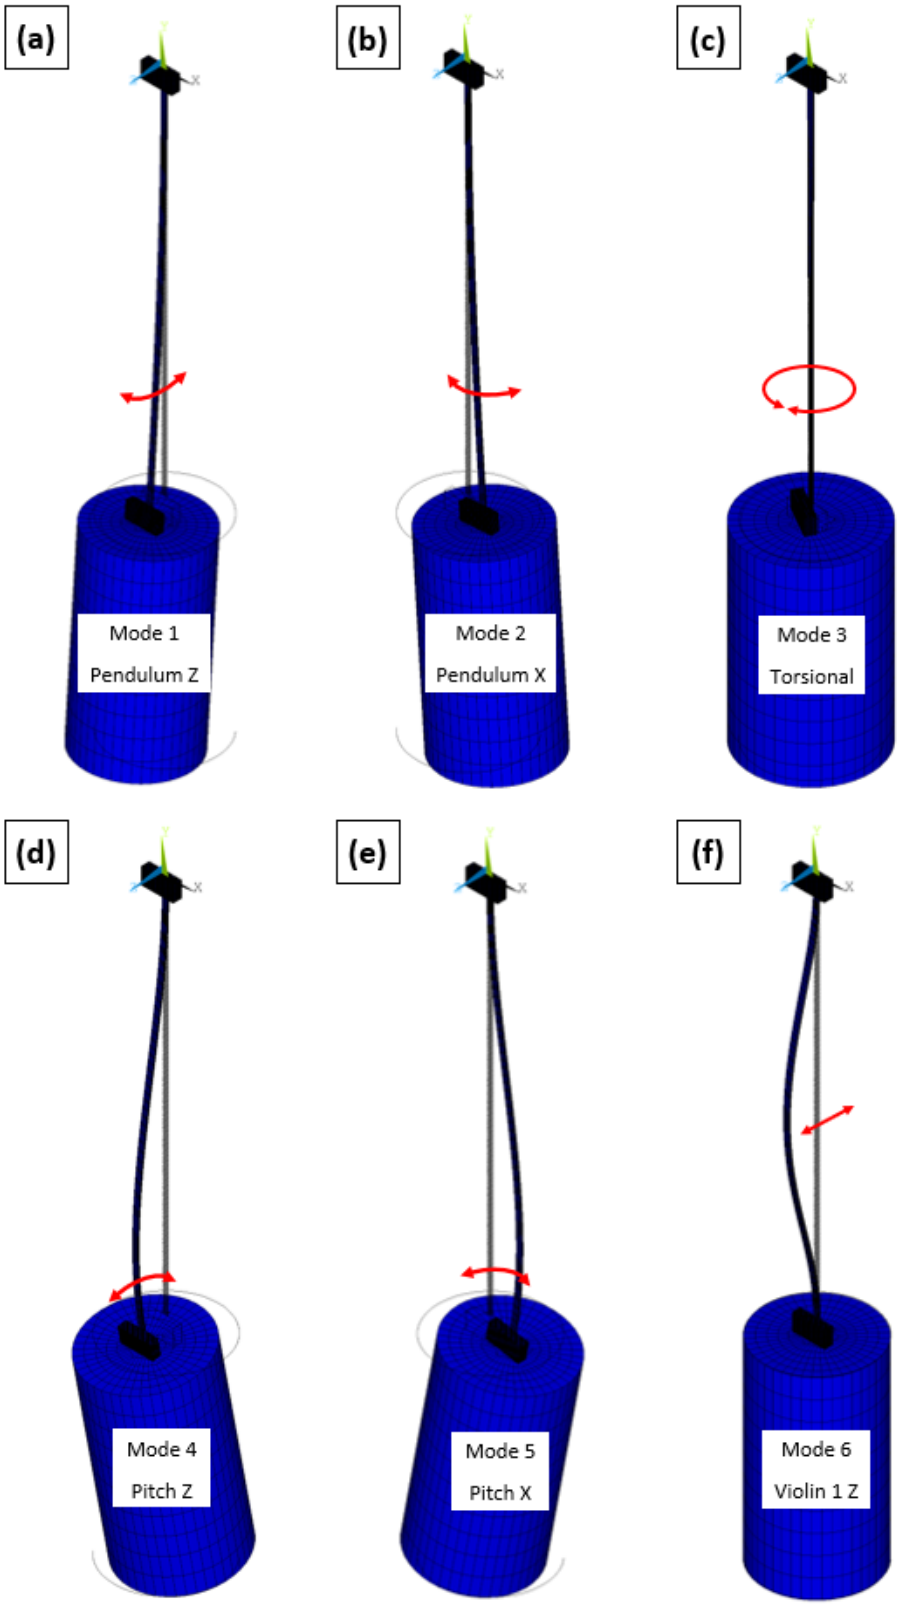

Figure 5.13: Illustration of the first 6 mode shapes of the suspension assembly valid for all bond thicknesess.

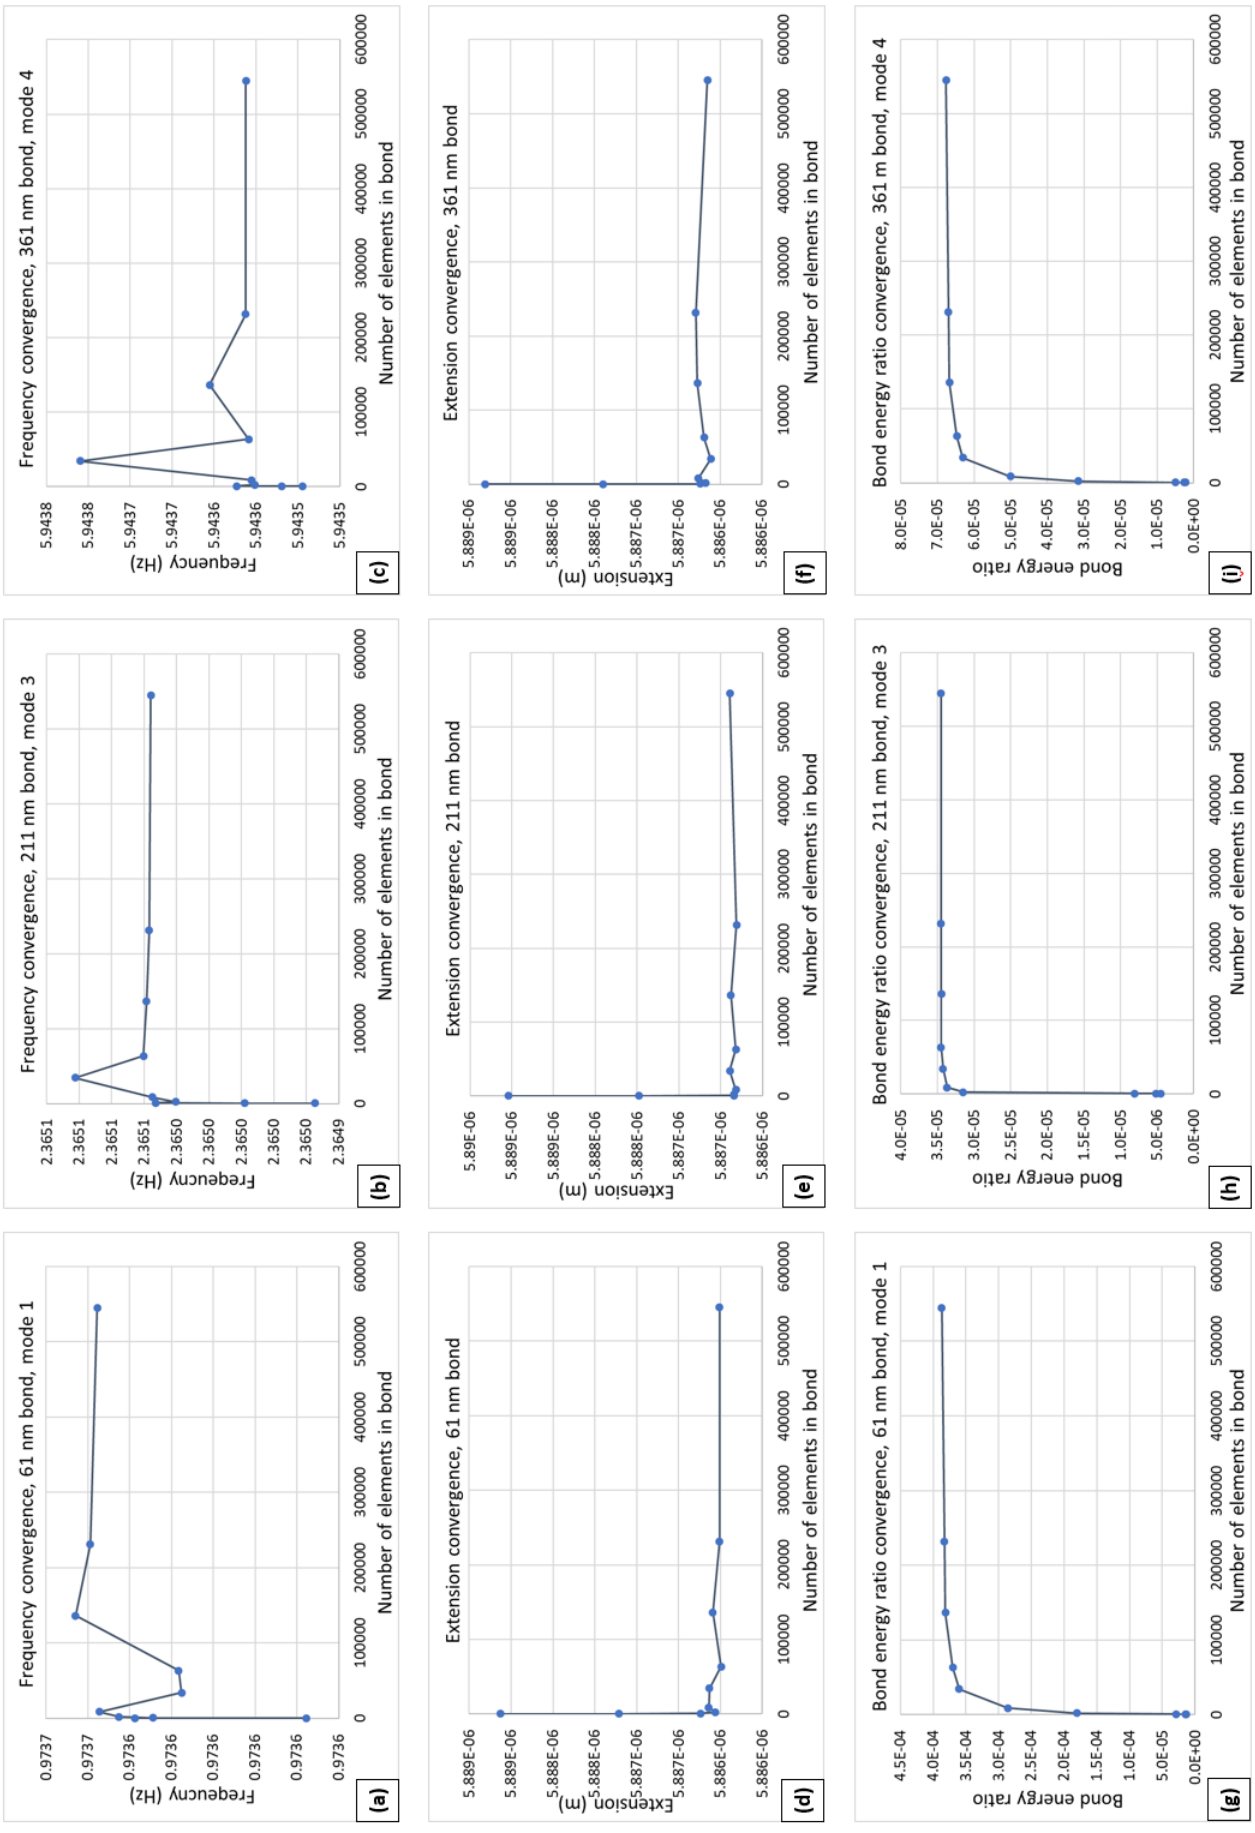

Figure 5.14: Convergence graphs from 61 nm, 211 nm and 361 nm bond thicknesses representing frequency, extension and bond energy for modes 1, 3 and 4 respectively.

5.5.3.4 Bond convergence in the Z plane

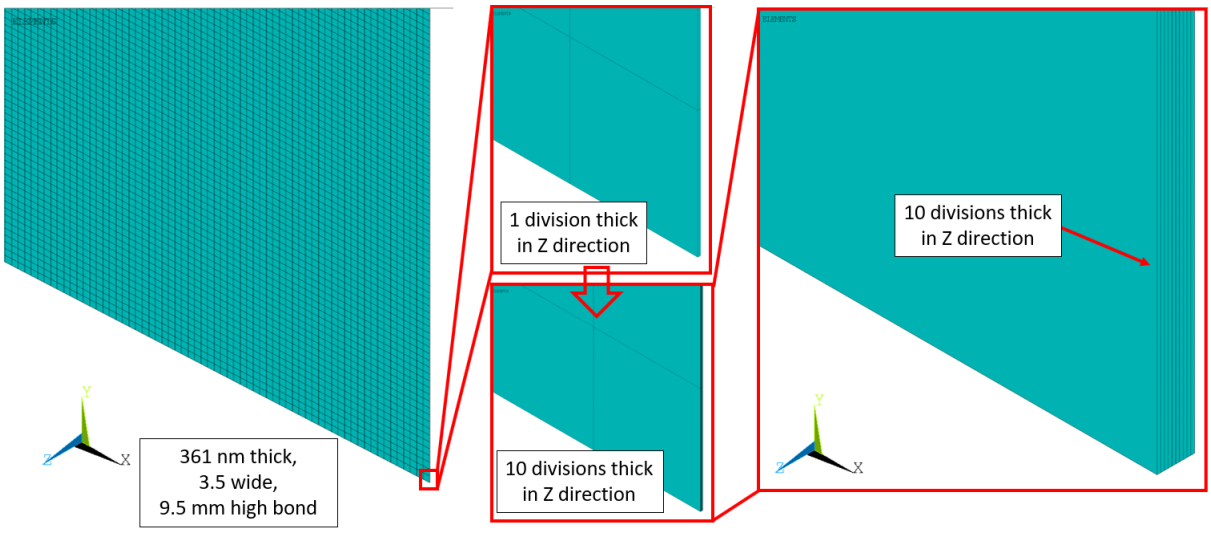

Figure 5.15: Example of bond convergence check by increasing the number of elements from 1 to 10 through the Z direction, the thickness of the bond.

In order to complete verification of the models it is also important to investigate convergence in the Z direction, through the thickness of the bond, as shown in figure 5.15. The most efficient models from the bond energy ratio convergence plots were re-run again, splitting the Z direction 10 elements deep, leading to a 10× increase in the number of elements. No change was observed in the resultant values and so it was determined that having 1 element representing the bond thickness was sufficient.

5.5.3.5 Improving model efficiency

For greater model efficiency, all model elements were replaced with 8 node SOLID185 elements instead of 20 node SOLID186 elements. Since the model solves accurately in a reasonable time, this was not applied to HC bonds due to verification of 1 element thickness being sufficient. The difference in the bond energy ratio between the two models was 0.2 %.

By changing the HC bond elements from SOLID186 to SOLID185, one removes the mid-nodes in the Z direction. Due to the large number of elements in the bond, efficiency increases are most likely to be found here. During the development of the single ribbon bonded models, SOLID185 elements were used in the HC bonds and verified as equivalent to SOLID186 for these bond geometries. This was particularly useful for extremely large bond areas with a large number of elements (as will be discussed later for model 3). However, for the 4 ribbon suspension work later on, SOLID186 elements were used throughout the model for ease of building and due to reasonably short solve times.

5.5.3.6 Bond deformation

Since bonds of this thickness or geometry for silicon suspensions have never been directly modelled before, the bond deformation for 6 modes was checked in order to better understand the dynamics.

The resultant bond deformation for each mode from the 361 nm bond are shown in figures 5.16(a)-(f). These bond geometries results correspond to each mode shape shown in figures 5.13(a)-(f). These deformations are exactly the same for the two thinner bonds.

One can see bond deformation is non-trivial, for example the pendulum X mode does not deform in a trapezoid manner as one might simplistically expect. The ability to observe deformation dynamics of real HC bond thicknesses is one of the many advantages gained from being able to directly model bonds.

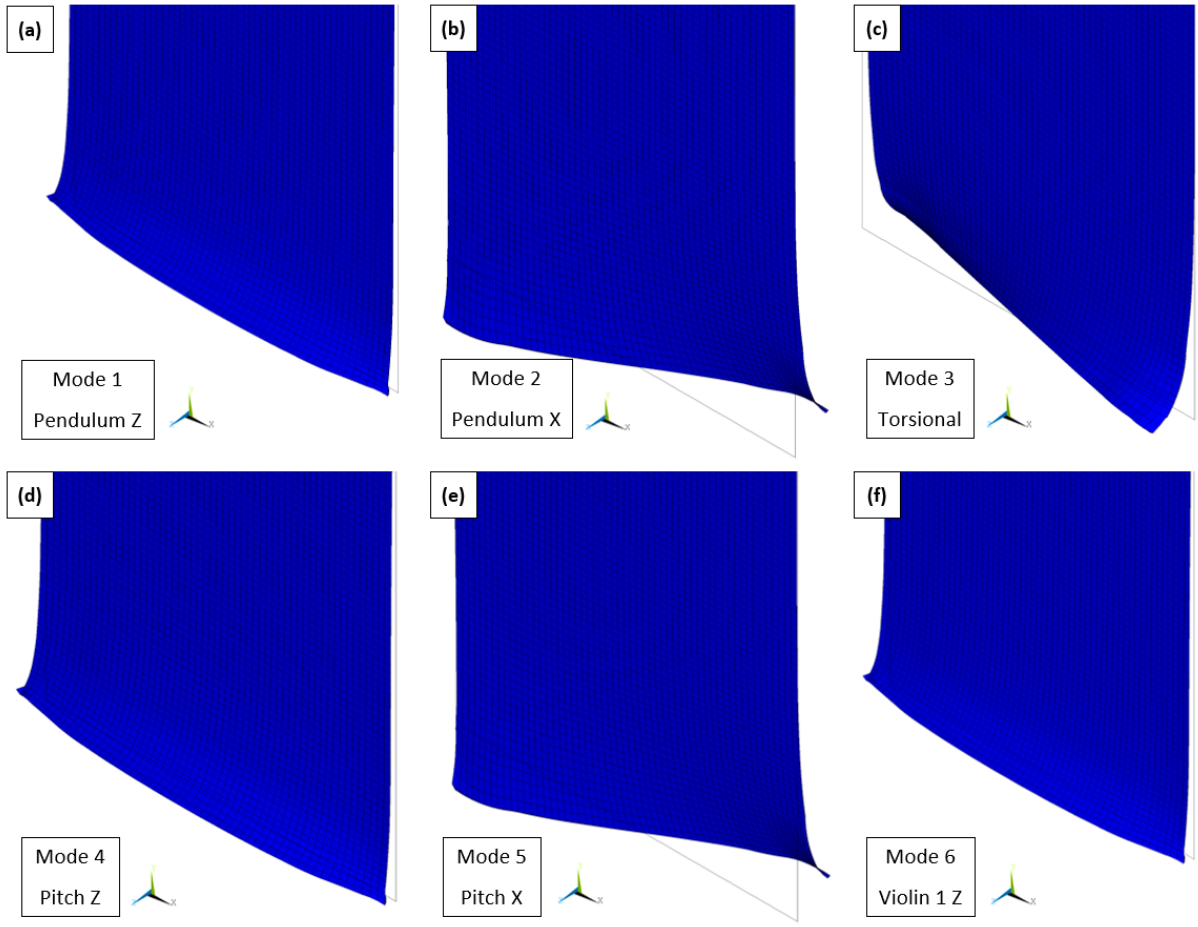

Figure 5.16: Illustration of the first 6 mode shapes of the top 361 nm bond which correspond to figures 5.13(a)-(f). These images are focused towards the bottom area of the top bond where the geometry deforms most.

5.5.3.7 Comparison of direct modelling to extrapolation

Applying Elliffe's technique to extrapolate back to nm scales requires modelling of ideally 3 or more thicker bonds, typically on the order of μm. Bond thicknesses of 1 μm, 10 μm and 100 μm were modelled. Only 3 convergence models were run for each thickness, selecting the largest, smallest and middle asize values to confirm the suspected dog leg shaped convergence. This corresponded to bond element aspect ratios of 12:1, 1:1 and 9:1 respectively. The latter aspect ratio increases again since the bond gets thicker in the Z direction whereas the XY asize value is fixed.

It was found, the middle asize value, set at 23 μm was sufficiently accurate, with the largest error, when compared to the highest element model, arising from the 1 μm bond thickness, specifically the mode 2 bond energy ratio. This is consistent with the fact this bond thickness has the largest aspect ratio. Furthermore, the smallest errors were encountered in the 10 μm model due to the better aspect ratio. Modelling was also conducted to confirm one element for bond thickness in the Z direction was sufficient.

These models typically converged faster due to lesser elements in the bond, taking between 90 to 130 minutes to solve. By comparison, the time it took to directly model and solve the first 6 modes of the thinnest, most efficient 61 nm bond model was 146 minutes. The bond contained 136192 elements. This took approximately the same length of time, ± 10 %, for each of the thicker extrapolation bonds. Even the most efficient extrapolation model, requiring a minimum of 2 models to produce a linear extrapolation with poorer accuracy, took longer than running one directly modelled bond.

One can take the energy ratio in each bond and plot them against element thickness for each of the 6 modes. This can then be extrapolated down to a 61 nm bond, for example, and compared to the directly modelled bond. The results of extrapolated values compared to directly modelled values for modes 1 and 2, as representative for primary X and Z modal directions, are shown in figure 5.17.

One notes more energy appears to be in the bond in figure 5.17(a), the pendulum Z mode, than in figure 5.17(b), the pendulum X mode, most clearly seen when looking at the 100 μm thick bond. This appears consistent with the fact the pendulum X mode is much stiffer in this bending direction and so more strain energy will be present in the bond.

One noted that previous extrapolations assumed a linear relationship [152] [397]. Both the linear and power fits shown in figures 5.17(a) and (b) do not provide a satisfactory fit across both modes. Comparing the linearly extrapolated data to the direct model value, mode 1 linear fit is 61 % higher than the directly modelled bond, with mode 2 47 % higher. Instead, a power law was fitted to investigate if this was more representative of the relationship between bond thickness and bond energy ratio. Comparing the power law extrapolated data to the direct model value, mode 1 linear fit is 54 % lower than the directly modelled bond, with mode 2 92 % lower.

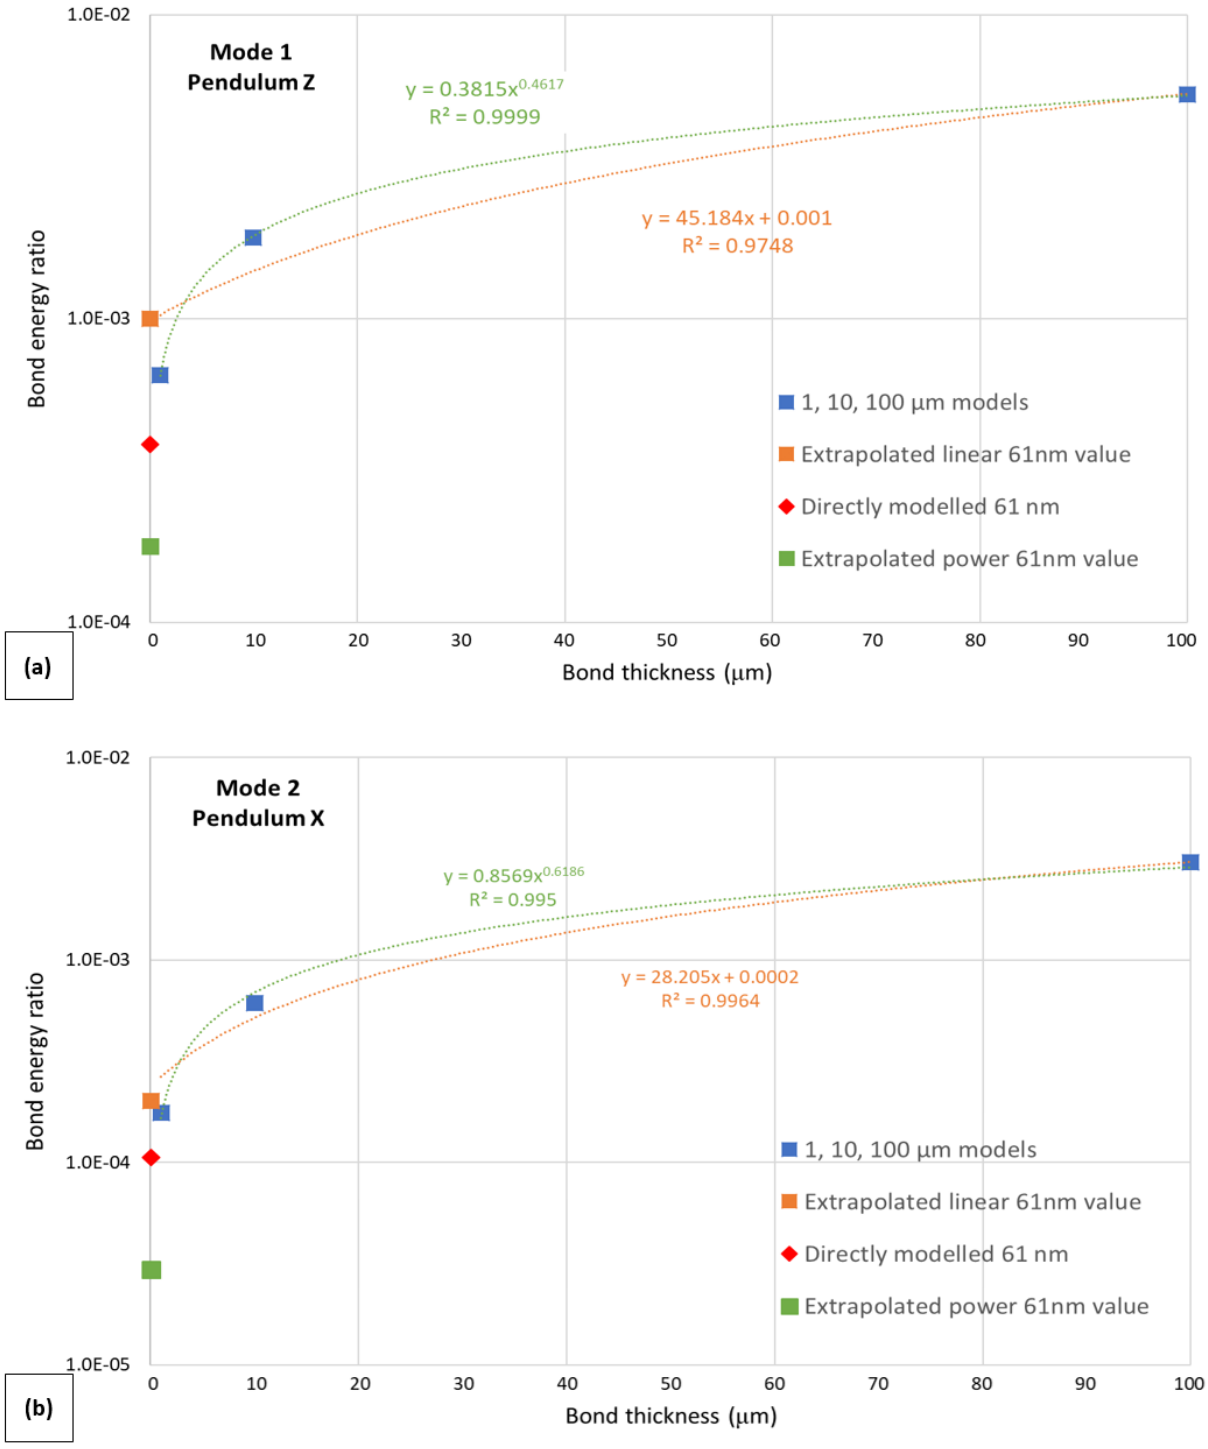

Figure 5.17: Linearly and power law extrapolated values for a 61 nm bond compared to directly modelled values for (a) mode 1 and (b) mode 2.

Looking at both graphs in figure 5.17, one notes the benefits of direct modelling over extrapolation, particularly as both modes do not fit the assumed linear relationship. While the power law may better fit the data, confidence in this fit is still low as the physics of bond energy ratios as a function of volume is not well understood at these extremely thin bond thicknesses. Furthermore, the overall computation time for direct modelling, plus bond deformation data shows

this method to be more efficient and accurate.

As shown in figure 5.18(a), it was noted when comparing the directly modelled 61 nm bond to 211 nm and 361 nm bonds that the bond energy ratio for mode 1 (as well as modes 4 and 6) were *higher* than the 211 nm, appearing to contradict the assumption that thinner bonds should have less energy. This is not the case in figure 5.18(b), where the convergence curves support this assumption.

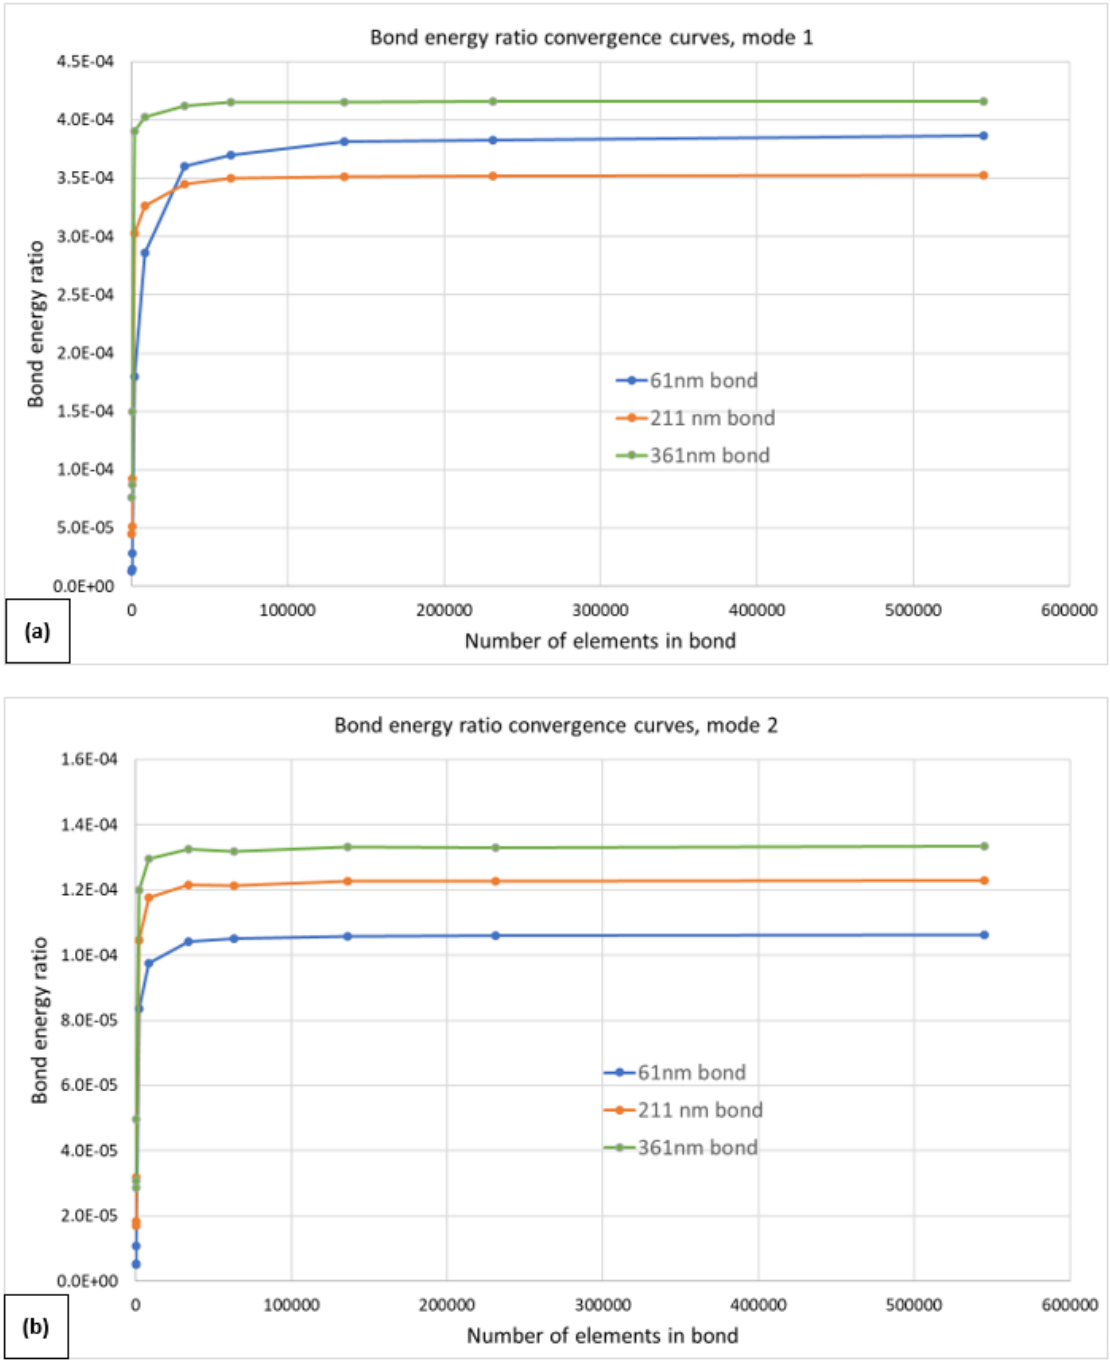

Figure 5.18: Convergence curves for 61 nm, 211 nm and 361 nm bonds for (a) mode 1 and (b) mode 2.

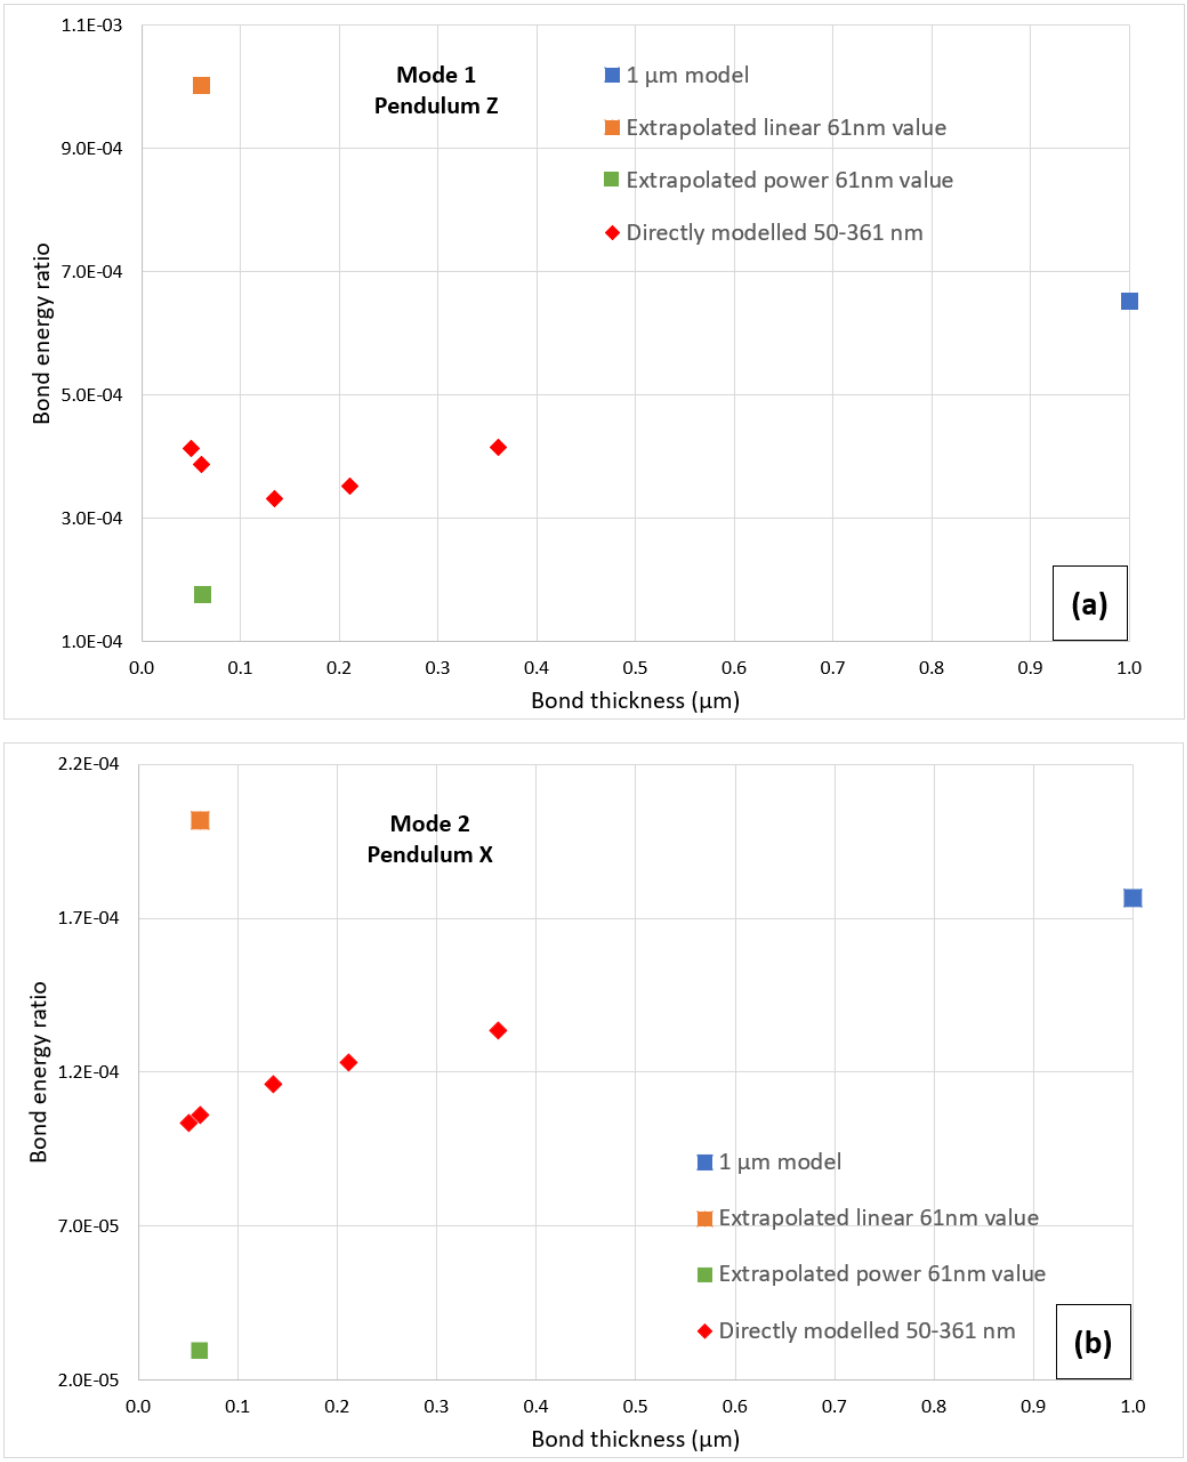

Figure 5.19: Inclusion of 50 nm, 135 nm, 211 nm and 361 nm bonds into figure 5.17, focused on area of interest, for (a) mode 1 and (b) mode 2.

To try and probe this relationship, the 211 nm and 361 nm bonds were included. Two more bond thickness models, 50 nm and 135 nm, were created to better understand how the energy changes around the thinnest of bonds. As shown in figures 5.19(a) and 5.19(b), the 50 nm bond conveniently matches the thinnest bonds created in [191] and the 135 nm bond thickness was

selected as an average thickness between the 61 nm and 211 nm bonds.

One can see the effect of rising bond energy in figure 5.19(a), around the minima 135 nm bond. Conversely, mode 2 in the stiff X direction, shows a proportional relationship between bond energy and bond thickness, as shown in figure 5.19(b).

While it is possible that this seemingly anomalous result is at or near the limit of the ANSYS modelling, the fact it applies to the three Z directional modes, and not X directional modes, as shown in figure 5.18(a), implies this is possibly the resurgence of the corner effect issue witnessed in figure 5.12.

Fundamentally, this anomaly is of no concern, since consistently strong HC bonds between silicon require oxidation of at least one of the silicon components, yielding a minimum bond thickness of >200 nm [401]. While a bond thickness of 61 nm between two unoxidised pieces of silicon is possible, preliminary results show hydrogen bubble formation for this thickness of bond can cause poor bonds with high mechanical loss [404]. As such, it is believed this modelling technique is valid for current HC bond thicknesses likely to be used in 3G detectors.

The above demonstrates the first direct FEA modelling of real thickness HC bonds within a suspension geometry.

5.6 Thermal noise contribution from HC bonds

Using this new modelling technique, the contribution of HC bonds compared to other thermal noise sources was estimated for proposed 3G gravitational wave detector suspensions.

The 361 nm bond thickness was used for this section of work. This assumed a 150 nm thermal oxide thickness on the silicon ears and suspension ribbon as discussed in section 5.5.3.

The following subsections take the suspension as built in chapter 4, and iterate each model, one variable at at time, towards an ET-like suspension.

5.6.1 Model 1 - Double bond model

Since the suspension was bonded to an ear-like structure at the top and bottom, the next step was to create the double bond model, as shown in figure 5.20(a). This model can be quickly verified by checking the top and bottom bond energy ratios for particular modal deformations. For torsional (mode 3) or violin (mode 6) mode, one noted the top and bottom energy ratios match to within 0.3 %, as expected. Likewise, for pendulum modes 1 and 2, one would expect most of the energy to be in the top bond rather than the bottom, as there is little bending in the bottom bond for a single ribbon pendulum. Modes 1 and 2 top bond energies are approximately 2 orders of magnitude greater than bottom bond energies.

This model served as the base model onto which a series of single variable iterations can be applied to move from a 1 kg low-stress suspension to a geometry representative of an ET-scale suspension.

5.6.2 Model 2 - Scaled mass model

The next step was to increase the mass on the bottom of the suspension. For comparative purposes, the mass geometry was kept the same, so a pseudo-density of $125\,570\,\mathrm{kg\,m^{-3}}$ was applied to the mass volume to yield a mass of 50 kg, representative of a four ribbon suspension supporting a 200 kg test mass. This suspension geometry is shown in figure 5.20(b). This changed the resultant mode shapes and so the modal expansion was increased from 6 to 8 modes in order to capture all previous modes for comparison. The resultant mode shapes, in ascending order were; torsional, pendulum Z, pendulum X, pitch Z, pitch X, bounce, violin Z and violin X. The corresponding bond energy ratios in both the top and bottom bonds increased with the increase in mass as expected. This increased the suspension stress to 181 MPa, which with a safety factor of 3 applied [256], would require silicon suspension stresses of 543 MPa which, as shown in chapter 3, is currently unrealistic.

5.6.3 Model 3 - Scaled stress model

The next step was to scale model 2 to have the same stress in the ribbon as model 1, 3.6 MPa as practically demonstrated in chapter 4, while maintaining the ribbon aspect ratio. This lead to a substantially thicker ribbon, with dimensions 23.9 mm wide and 5.3 mm thick. As such, the bond area also increased from 9.5 mm$\times$3.5 mm to 9.5 mm$\times$23.9 mm. This also required the width of the 20 mm blocks to be extended to 25 mm to accommodate this new ribbon and bond geometry. Due to this increase in bond area, the model size increased substantially, taking around 5 days to build and solve. One notes, the model solve time does not scale linearly with the number of elements. This suspension geometry is shown in figure 5.20(c).

5.6.4 Model 4 - Scaled working stress model

The next step was to increase the stress to a realistic working stress of 40 MPa. Referring to section 3.9 of chapter 3, this stress value with a realistic safety factor of 3, provides an optimistic attainable stress sitting between the median and 1st quartile values for CMP treatment #1, CMP treatment #2, argon ion etch treatment #2 sample sets and well below the ITM KOH sample set. Furthermore, this value matches the suspension strength limited case in [136], inclusive of a safety factor of 3 [256]. Maintaining the same ribbon aspect ratio leads to a new ribbon width of 7.44 mm and thickness of 1.65 mm. The silicon blocks were returned back to the standard 20 mm width geometry. This gives a new bond area of 9.50 mm$\times$7.44 mm. This suspension geometry is shown in figure 5.20(d).

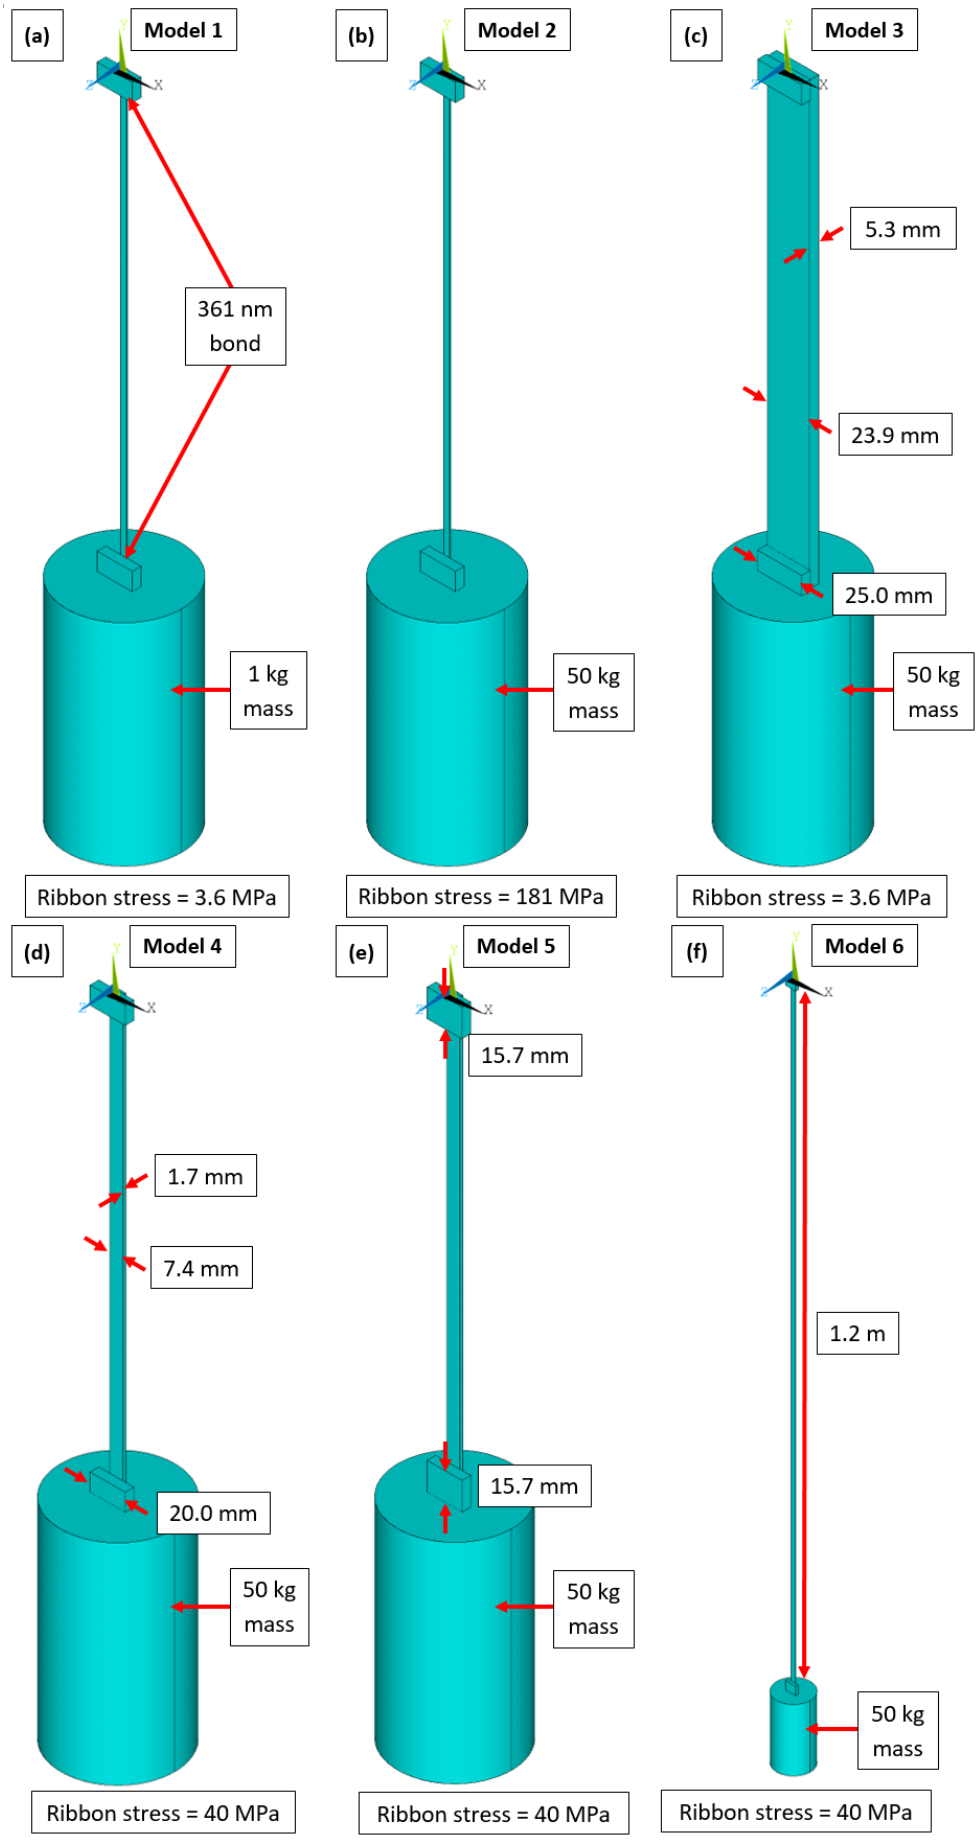

Figure 5.20: Illustration of suspension iteration steps to produce models 1, 2, 3, 4 5 and 6. Note dimensional values have been rounded to 1 decimal place.

5.6.5 Model 5 - Working bond stress model

The next step was to investigate the bond stress resulting from the geometry in figure 5.20(d). The resultant bond stress for this geometry is 6.9 MPa. The typical observed breaking stress for HC bonds between silicon substrates is 30 ± 17 MPa [186]. Taking the conservative lower bound value of 13 MPa, and applying a safety factor of 3, would require a maximum bond stress of 4.3 MPa. The bond stress in model 4 exceeds this, so the bond area was increased in model 5 to conform to the bond stress limit. One must therefore increase the bond area to satisfy a bond stress with a safety factor of 3. In order to keep the same ribbon aspect ratio and stress, the width of the ribbon must remain fixed leaving the one dimension able to increase the bond area being the height of the bond in the vertical Y direction. The model shown in figure 5.20(e), extends the length of the block to 15.71 mm, increasing the bond area to 15.21 mm$\times$7.44 mm, maintaining the 0.5 mm unbonded area at the top area of the top silicon block, giving the necessary a bond stress of 4.3 MPa.

5.6.6 Model 6 - Scaled length model

The final step was increasing the suspension length. The longer the suspension length the greater the dissipation dilution provided by the pendulum geometry, as shown in equation 2.23 in chapter 2. ET suspension lengths of between 1.5 m [127] and 2 m [126] have been proposed. For this reason, double the aLIGO suspension length was chosen, 1.2 m [256]. This shorter suspension length provides a conservative estimate of the dilution. The length of the suspension should not impact on the bond energy ratio for this final model. This suspension geometry is shown in figure 5.20(f).

5.6.7 Comparison of mechanical loss contributions in a single ribbon suspension

One can visualise the evolution of the bond loss for the soft pendulum Z and stiff pendulum X modes in each model and compare it to the next dominant mechanical loss factors from the suspension ribbon, namely surface, bulk and thermoelastic losses. The final ribbon dimension was used for calculating the corresponding loss values.

The mechanical loss comparison values are therefore:

- $\phi_{surface} = 1.03 \times 10^{-9}$, taking $h\phi_s$ as 5×10^{-13} from [153] and calculating using equation 2.8 from chapter 2
- $\phi_{bulk} = 1.09 \times 10^{-9}$, at 20 K [154]
- $\phi_{thermoelastic} = 2.28 \times 10^{-14}$, using equation 2.5 from chapter 2 with the following variables:

- $f = 10\,\mathrm{Hz}$
- $T = 20\,\mathrm{K}$
- $\rho = 2328\,\mathrm{kg\,m^{-3}}$
- $C = 3.41\,\mathrm{J\,kg\,K^{-1}}$
- $k = 1225\,\mathrm{W\,m^{-1}\,K^{-1}}$
- $E = 157\,\mathrm{GPa}$
- $\frac{dE}{dT} = -1.1 \times 10^{6}\,\mathrm{Pa\,K^{-1}}$
- $\alpha = -2.5 \times 10^{-9}\,\mathrm{K^{-1}}$

The value used for the material loss of the HC bond is $\phi_{bond} = 1.18 \times 10^{-3}$, at 20 K, taken from [110] [111]. The bond loss contribution for each mode is calculated from the respective bond energy ratios as follows:

$$\phi_{bond\,contribution} = \left(\frac{E_{top\,bond}}{E_{model}} + \frac{E_{bottom\,bond}}{E_{model}} \right) \times \phi_{bond}. \tag{5.5}$$

A plot comparing mode 1, the soft pendulum Z direction, and mode 2, the stiff pendulum X direction, for each model against the next dominant mechanical loss factors is shown in figure 5.21 and figure 5.22 respectively.

Comparing the overall values for figure 5.21 to figure 5.22, one notes the energies are consistently lower in figure 5.22. This is to be expected due to the aspect ratio of the ribbon, since this pendulum X direction is through the thickest ribbon dimension. This therefore pushes the bending point of the ribbon further away from the top bond, taking energy away from the top bond region. This has little effect on the bottom bond region since little bending is experienced at the bottom bond.

5.6.7.1 Model 1 results

Figure 5.21 and figure 5.22 immediately shows that the 1 kg HC bonded suspension, built in chapter 4, is completely dominated by bond loss. The next dominant loss mechanism, bulk loss, is two orders of magnitude below the bond loss. The bulk loss value used here comes from a commonly used, experimentally measured, data set taken from one sample over the full cryogenic temperature range in [154]. However, the surface loss is calculated using equation 2.8 in chapter 2, where the $h\phi_s$ value of 5×10^{-13} is determined experimentally from multiple silicon cantilevers with a variety of surface finishes [136] [153]. In reality, the surface loss is likely higher than the bulk loss for any one sample, as shown in figure 2.7 of chapter 2 [153]. The values used here for surface and bulk loss approach the best values for known silicon. Future analysis is required to ensure surface treatments that might be necessary, as studied in chapter 3, do not detrimentally affect the surface loss of the suspension ribbons or fibres. This is beyond the scope of this thesis but must be considered in future work.

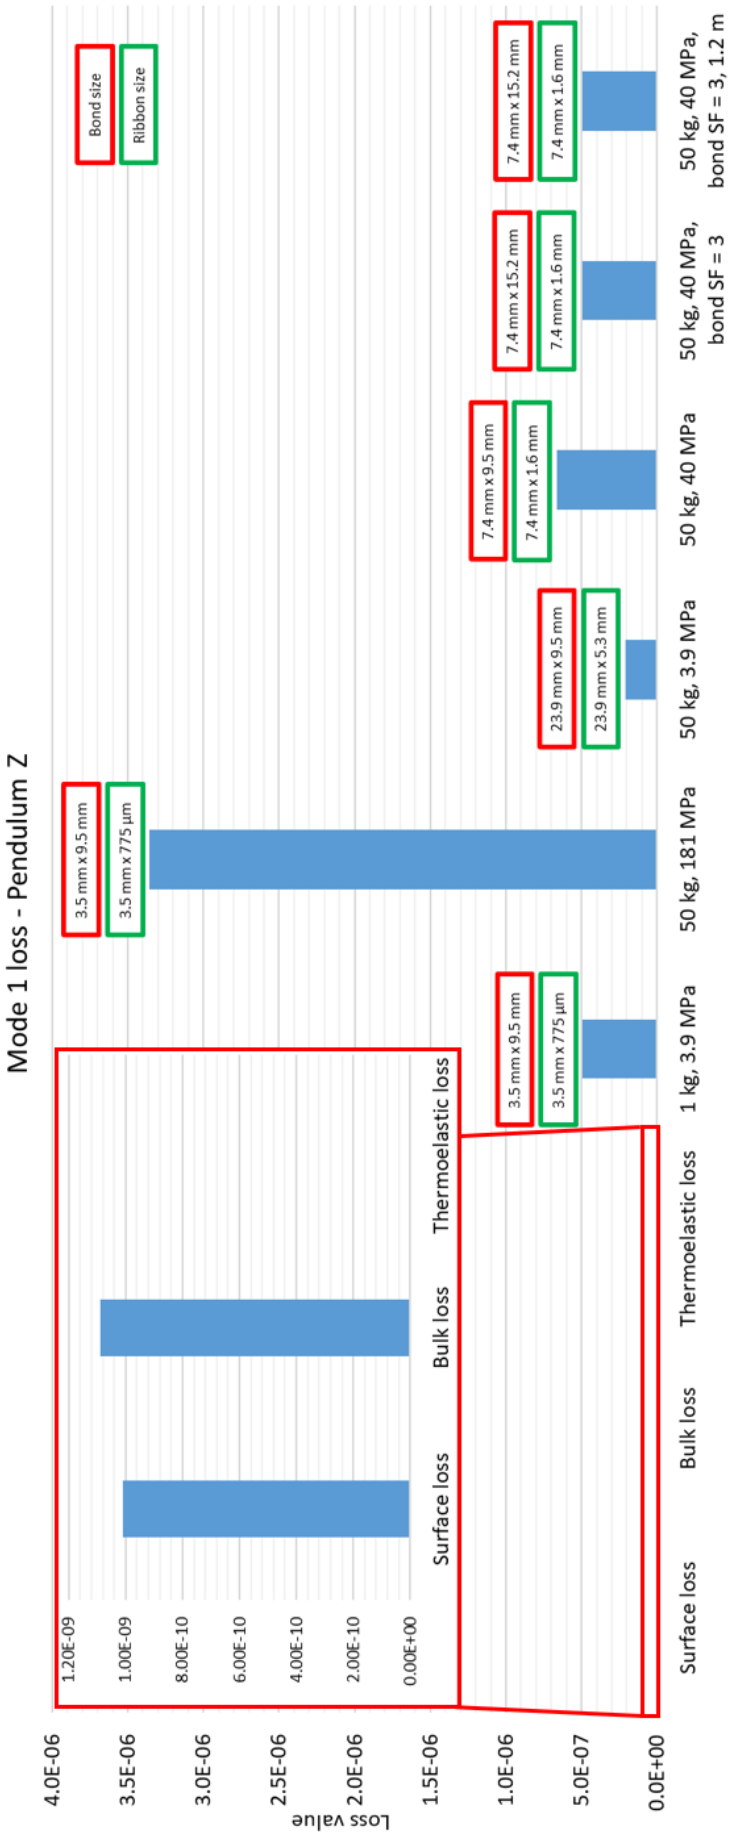

Figure 5.21: Total bond loss energies in mode 1, pendulum Z, for model iterations 1 to 6. Inset chart shows comparative surface, bulk and thermoelastic lost on a more visible scale.

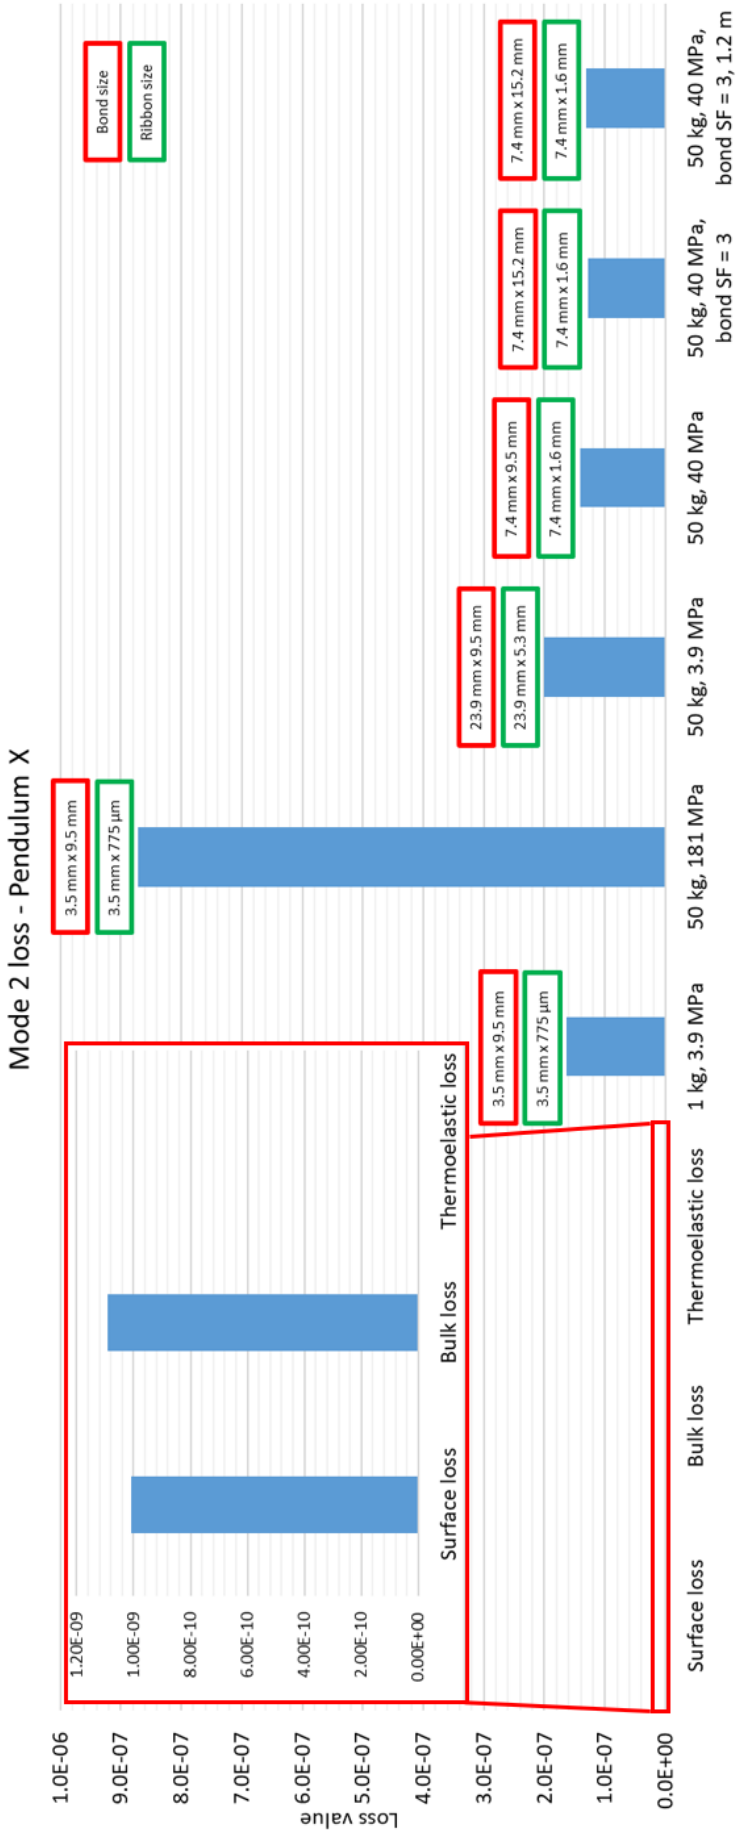

Figure 5.22: Total bond loss energies in mode 2, pendulum X, for model iterations 1 to 6. Inset chart shows comparative surface, bulk and thermoelastic lost on a more visible scale.

5.6.7.2 Model 2 results

For the model 2, changing the mass from 1 kg to 50 kg increases the energy in the bond region for both modes. This is due to the bending point moving closer to the top bond region due to the increase tension in the ribbon, which drives more energy into the lossy top bond region. The energy in the top bond was still two orders of magnitude greater than the bottom bond since there is comparatively little bending at the bottom of the single ribbon pendulum.

5.6.7.3 Model 3 results

Most notable was the energy distribution in the third model iteration with the excessively thick ribbon used to bring the ribbon stress back to the original 3.9 MPa. For the soft pendulum Z mode, the thickness of the ribbon has increased the second area moment of inertia, I_z, substantially by 4 orders of magnitude which, from equation 5.4, increases the bending point of the ribbon to a theoretical value of 1.4 m below the top of the ribbon. As previously discussed, this simplistic theory does not take into account higher-order effects which become important for conditions where $\frac{L}{a} \not\gg 1$ such is the case for this model. This result is important to note, as it offers one unexpected potential way of of reducing the energy within the lossy HC bond region.

5.6.7.4 Model 4 results

For model 4, one sees an increase in the pendulum Z direction bond loss due to the bending point of the suspension creeping back up towards the bond area, driving more energy into this region. Conversely, we see a small decrease in the pendulum X direction.

5.6.7.5 Model 5 and 6 results

There are little differences in the bond loss energies between models 4, 5 and 6. There is a small decrease from model 4 to 5, due to the extended height of the top and bottom block that stiffens the ribbon slightly, reducing energy in the bond region. Increasing the length of the pendulum should not affect the energy in the bond as evidenced by comparing models 5 and 6 since the bending point calculated from equation 5.4 is independent of length, L.

5.6.7.6 Dissipation dilution for models 1-6

Increasing the ribbon thickness has two primary advantages, helping to reduce stress within the suspension as well as providing a larger cross-sectional area for thermal extraction of the test mass. However, the disadvantages are substantial. Firstly, this will raise the baseline noise from the pendulum modes further into detection bands, as can be seen in the frequency increase, shown in table 5.9, of 0.72 Hz for pendulum Z, and 5.27 Hz for pendulum X, from model 1 to model 3. Furthermore, the dilution is reduced significantly as shown in table 5.9.

As discussed in section 4.2 of [144], one can calculate the dilution factor, D, from energy values, E, extracted and summed from ANSYS elements [109]:

$$D = \frac{E_{stored\,in\,gravitational\,field}}{E_{stored\,in\,wire}} = \frac{E_{model\,kinetic\,energy}}{E_{model\,strain\,energy}} \tag{5.6}$$

These values are also recorded in table 5.9 for the corresponding mode. It should be noted one can not compare the individual relative magnitudes of energies from model to model.

Mode	Frequency	Mode shape	Strain energy	Kinetic energy	Dilution factor
Model 1 - 1 kg, 3.9 MPa					
1	0.97	Pendulum Z	1.58	18.67	11.8
2	1.43	Pendulum X	21.59	40.11	1.9
Model 2 - 50 kg, 181 MPa					
2	0.91	Pendulum Z	0.18	16.31	91.5
3	0.96	Pendulum X	0.90	18.35	20.3
Model 3 - 50 kg, 3.9 MPa					
1	1.69	Pendulum Z	46.88	56.39	1.2
2	6.70	Pendulum X	868.21	886.74	1.0
Model 4 - 50 kg, 40 MPa					
1	0.95	Pendulum Z	0.94	17.66	18.9
2	1.15	Pendulum X	9.60	26.18	2.7
Model 5 - 50 kg, 40 MPa, bond SF = 3					
1	0.96	Pendulum Z	0.97	18.02	18.5
2	1.23	Pendulum X	10.04	29.76	3.0
Model 6 - 50 kg, 40 MPa, bond SF = 3, 1.2 m long					
1	0.45	Pendulum Z	0.05	4.03	81.3
2	0.46	Pendulum X	0.27	4.22	15.8

Table 5.9: Frequency, energy and dilution results of pendulum Z and pendulum X modes for each model iteration.

The frequencies in table 5.9 are consistent where the thicker, stiffer X pendulum direction consistently produces a higher frequency compared to the thinner Z dimension. One notes as the pendulum length is increased substantially, the frequencies drift toward one another. Looking at equation 2.10 in chapter 2, this is due to the pendulum effective length dominating over the difference in second area moment of inertia between X and Z directions.

One can see the dilution increases with increasing mass when comparing models 1 and 2. This is a direct result of the increase in mass, increasing the relative energy stored as gravitational potential energy when the pendulum reaches its maximum displacement.

Gain in bond loss for model 3 is 16$\times$ greater for the pendulum Z mode and 4$\times$ greater for the pendulum X mode compared to model 2. However, a similar comparison for the dilution shows

model 3 to be 76× lower in the pendulum Z mode and 20× lower in the pendulum X mode. Looking at equation 2.35 in [144], shows that dilution is inversely proportional to the square root of the second area moment of inertia. As this increases with larger ribbon dimensions, the dilution is reduced. Fundamentally, any benefit in bond loss from increasing ribbon dimensions, and thus stiffness, is immediately lost in the suspension dilution.

Dilution recovers by thinning the ribbon in models 4 and 5, with a slight decrease in dilution between these models for the pendulum Z mode arising from the extension of the top and bottom block heights. Increasing this bond area while keeping the ribbon the same length will ultimately have an effect of stiffening the ribbon slightly as more of its length is attached to the thicker, stiffer blocks.

Finally, model 6 sees a significant increase in dilution. This is directly related to the length of the pendulum as dilution scales linearly with pendulum length as shown in equation 2.23 in chapter 2. The average increase in dilution for both modes between models 5 and 6 is a factor of 4.8. If one considers the pendulum length has increased from 0.25 m, to 1.2 m, this is also consistent in being an increase of a factor 4.8.

5.7 Developing a 4 ribbon suspension

With a 3G detector-scale, double bonded, single ribbon suspension successfully modelled with energy distribution and losses understood, one can implement this into a 4 ribbon suspension to simulate the real detector system and model its suspension thermal noise. One then compares the suspension thermal noise of this detector to the proposed ET design thermal noise requirements.

The work in this section looks *only* at the suspension thermal noise arising from the silicon ribbons and not from the test mass, penultimate test mass or any proposed system from which the final-stage suspension is attached to or hung, for example a marionette with steel wires.

5.7.1 Building and verifying the 4 ribbon ET-like model

It was decided to copy the aLIGO mass design and scale dimensions appropriately, including the flats on the side of the mass. The aLIGO test masses are 200 mm deep, 340 mm in diameter with 95 mm flats centrally located along the full depth of the mass [405] [406].

A script was written in ANSYS APDL that automatically scaled and built a silicon mass from aLIGO's silica test mass dimension, with the density of silicon listed in section 5.3. Iterative selection of the diameter of the 200 kg silicon mass yielded a diameter of 572 mm, a depth of 336 mm with a flat height of 160 mm. The penultimate mass was created to be the same dimension and density as the test mass mirror.

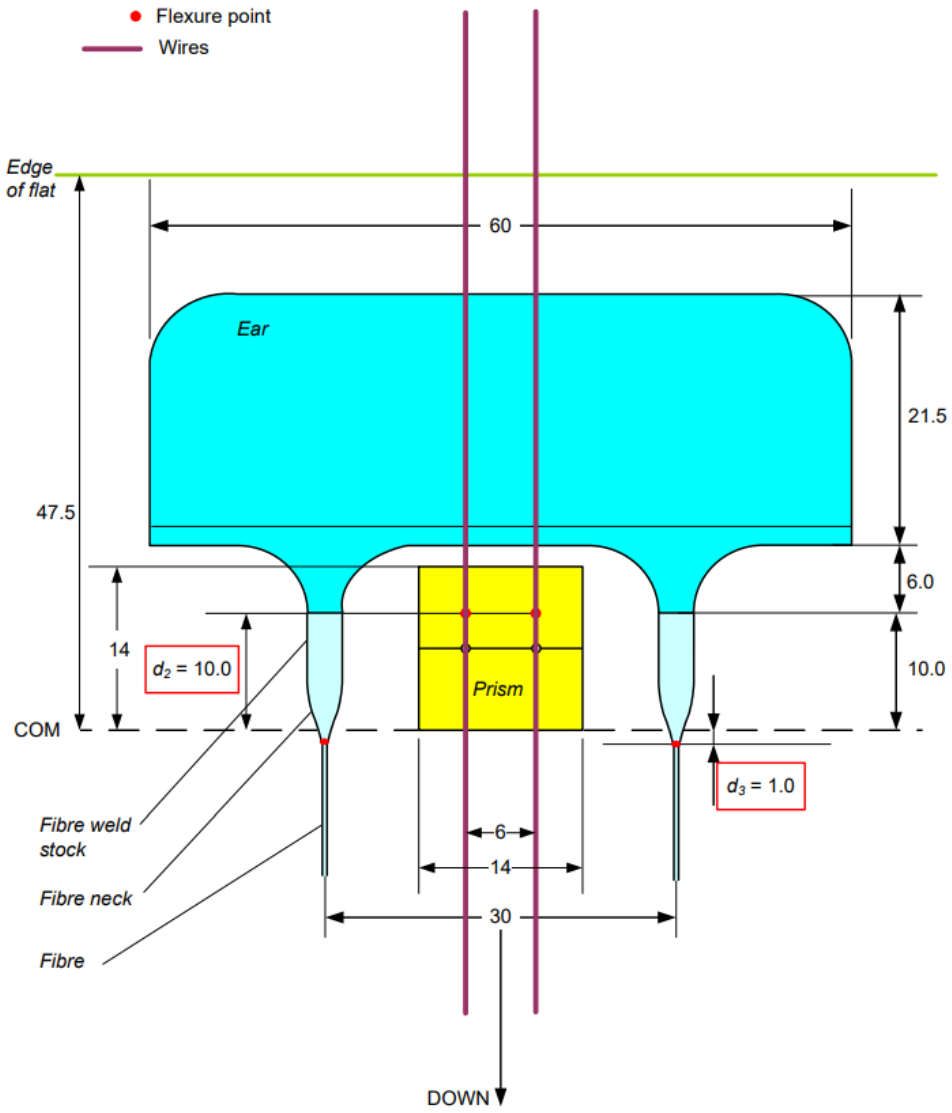

Figure 5.23: Drawing of aLIGO of d-values in relation to penultimate mass COM, silica fibres and HC bonded ear. The term flexure point, is interchangeable with bending point. The purple wires referenced in this diagram relate to the steel wires which hang the penultimate mass from the higher metal stages of the suspension system. Image obtained from [407].

The next step was to determine the bending point of the 1.2 m ribbon suspension in the primary interferometer direction, i.e. the thin ribbon, soft Z direction. This is important in order to scale aLIGOs "d-values", parameters that define the bending point of suspension fibres in relation to the centre of mass (COM) of penultimate and test mass [407]. This determines the pitch dynamics of the suspension [408].

A diagram of aLIGOs d-values for the penultimate mass is provided in figure 5.23. The d2 value is defined as the distance between the weld point and the COM of the mass. The d3 value is defined as the distance from the COM of the mass to the bending point of the suspensions fibre. One must adapt and define these values for the silicon suspension. As such, the d2 value is defined as the distance between the edge of the exposed ribbon (i.e. the bottom of the silicon

block ear) and the COM of the mass. The d3 value remains defined as the distance from the COM of the mass to the bending point of the suspensions fibre. The d2 and d3 values for the penultimate mass are mirrored for the test mass [408].

For the 1.2 m silicon ribbon suspension, determining the Z direction bending point of the ribbon requires careful calculation. aLIGO assumes infinite stiffness at the bottom of the ear on the penultimate test mass [407]. For this reason, a replica model of the 1.2 m single ribbon, shown in figure 5.20(f), was created without the top silicon block, which will now be referred to as the top ear for consistency, leaving only the bottom ear, ribbon and mass. HC bonds were not included in these bend point models for computational efficiency and the negligible impact they would have on this analysis due to their infinitesimal size. This model was displaced by 1 % of the suspension length, 12 mm in the Z direction while the top area of the ribbon was fixed. By then extracting nodal positions along the length of the displaced ribbon, one can extrapolate the bottom straight line portion back to the vertical to calculate the bending point, as described previously in section 5.4.2. This the same technique used in section 4.1.3 of [144].

This model produces a bending length, a of 30.60 mm. Comparison to theory, using equation 5.4 (since $a \ll L$ and hence $\frac{L}{a} \gg 1$), gives a is 29.86 mm, matching to within 2 %.

To account for the finite stiffness of the silicon ears, a second model was created. This model included only the top ear, the ribbon and the mass with the same technique shown in figure 5.8. This produced a bending length of 43.05 mm. However, since 15.21 mm of this model is constrained by the relatively stiff top ear, the bending length becomes 27.84 mm. By subtracting the difference between the two models, one can take into account the finite stiffness of the top ear to find the effective bending length of the ribbon-ear assembly. Going one step further than aLIGOs bend calculations, one can account for the ear flex by subtracting this difference from aLIGOs d2 value of 10 mm to give a d3 value of 7.24 mm. This 2.76 mm increase away from the COM is entirely consistent with introducing a finite stiffness and hence flex in the top ear. Taking the ribbon bending length of 30.60 mm this yields a d2 value of 23.36 mm from the COM of the penultimate mass to the bottom of the top ear.

Finally, the separation of the silicon ribbons were designed in line with figure 5.23 but were scaled to the new silicon mass depth of 336 mm which scales the centre-to-centre ribbon distance from aLIGOs 30 mm to 50.4 mm.

A drawing of the final ear positions and bending point parameters in relation to the COM of the penultimate mass is provided in figure 5.24.

Bonds were not introduced into this initial model in order to verify the simpler model against theory first. All 4 ribbon suspension models are constrained on the 4 top corners of the primary flats on the penultimate mass for static and modal solutions, consistent with [144]. The loaded extension of the model matched to within 0.2 % of theory.

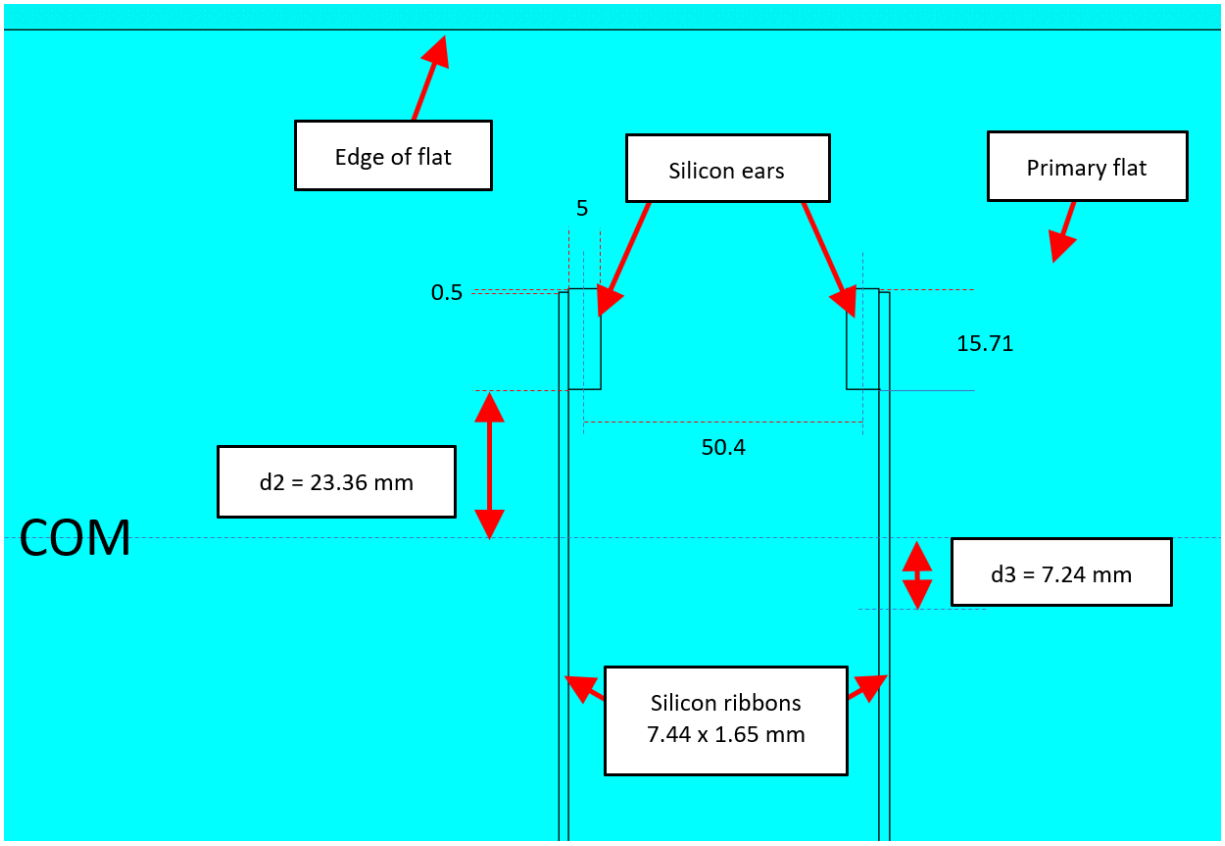

Figure 5.24: Side-on view of penultimate test mass ear positions and bend points of 4 ribbon silicon suspension scaled to match aLIGO. All dimensions in mm.

To compare model resonant frequencies to theory for the pendulum directions, it is important to note the effect of the bending points in the effective length of the pendulum. The silicon ribbon lengths are 1.2 m, with the d2 offset, this meant the distance between the penultimate and test mass COMs becomes 1.12 m. By taking into account the bending point below, for the penultimate mass, or above, for test mass, defined by the d3 value, this provides an effective length in the soft, pendulum Z direction of 1.11 m.

The stiff, pendulum X direction has a different bending point and hence different d2 and d3 values. As the primary mode of interest and largest contribution to suspension thermal noise, the d2 and d3 values are defined by the pendulum Z direction. The pendulum X values of d2 and d3 are therefore only a result of the chosen pendulum Z direction.

The same bend point extrapolation analysis was conducted using the same techniques shown in figure 5.8 for the stiff, pendulum X direction. Due to the 20 mm thickness, and hence relative stiffness, in the X direction of the ears, the flex in the ear was not taken into account for this analysis, unlike the Z direction. The displaced models yielded a bending length, a, of 142.9 mm in the stiff, pendulum X direction. This is more difficult to compare to theory using equation 5.4 (since $a \ll L$ and hence $\frac{L}{a} \gg 1$), and second order effects have increasing significance. However, using equation 5.4, only yields a 6 % difference between the model and theory. Now taking the assumption that the ear is infinitely stiff, then the effective bending length from the bottom of

the COM is the d2 value, defined by pendulum Z, minus the 142.9 mm bending length. By using the distance between the penultimate and test mass COMs, this provides an effective length in the stiff, pendulum X direction of 0.88 m.

Using these effective bending lengths in equation 2.9, one calculates a pendulum Z frequency of 0.48 Hz and pendulum X frequency of 0.57 Hz. The 4 ribbon model with no bonds produces modal frequencies of 0.47 Hz for the pendulum Z frequency and 0.52 Hz for the pendulum X frequency, matching theory to within 1.5 % and 8.4 % respectively. The larger error in the stiff pendulum X direction is primarily accounted for by the increasing significance in second order terms but there will also be a contribution from the assumption of infinitely stiff ears.

Furthermore, using equation 2.9, taking the stiffness in the Y direction, and equation 2.12, one can calculate the suspension bounce and violin mode frequencies. The model's bounce frequency was 28.5 Hz, matching to within 0.2 % of theory. The model's first violin mode in the Z direction was 59.0 Hz, matching to within 0.1 % of theory.

One should note, when calculating suspension thermal noise, for actual 4 ribbon detector suspensions, the bounce mode must be multiplied by a factor of $\sqrt{2}$, in order to account for the resonant vertical motion of the penultimate mass. This motion is not seen within the ANSYS model due to the constraints placed on the penultimate mass. All suspension thermal noise plots generated in this thesis apply this correction to the bounce mode.

Finally, it is important to compare the theoretical dilution to that of the model. Using equations 5.6, the extracted model kinetic and strain energies, provide a dilution in the pendulum Z direction of 37.1 and in the pendulum X direction of 6.6. Comparing this to theory using equation 2.23 and defining L as the effective lengths calculated above for each direction, the model dilution matches theory to within 0.3 % and 2.9 % for the pendulum Z and pendulum X directions respectively. The bounce mode has no dilution due to no gravitational restoring force, and hence no energy being stored in the lossless gravitational field. The model closely matched this prediction, producing a dilution of 0.99. The violin Z dilution was found to be 16.6, approximately half the dilution of the pendulum Z dilution. This is consistent with equation 2.14.

This series of checks, validates the 4 ribbon model suspension against theory.

5.7.2 Introducing HC bonds into the 4 ribbon model

One can introduce the top and bottom bonds into the system. For computational efficiency, it was decided that top and bottom 361 nm bonds between the top and bottom ears and ribbons would only be applied to one ribbon. This keeps the model size manageable allowing for simple scaling where one can multiple the top and bottom bond energy ratios by 4. It was assumed this was unlikely to significantly affect the dynamics of the suspension system due to the infinitesimal size of the bond compared to the suspension geometry. One can refer back to figure 5.12, for an idea of how the 361 nm bonds sits in between the ribbons and ears.

The final, ET-like, HC bonded silicon suspension is shown in figure 5.25(a). An ANSYS

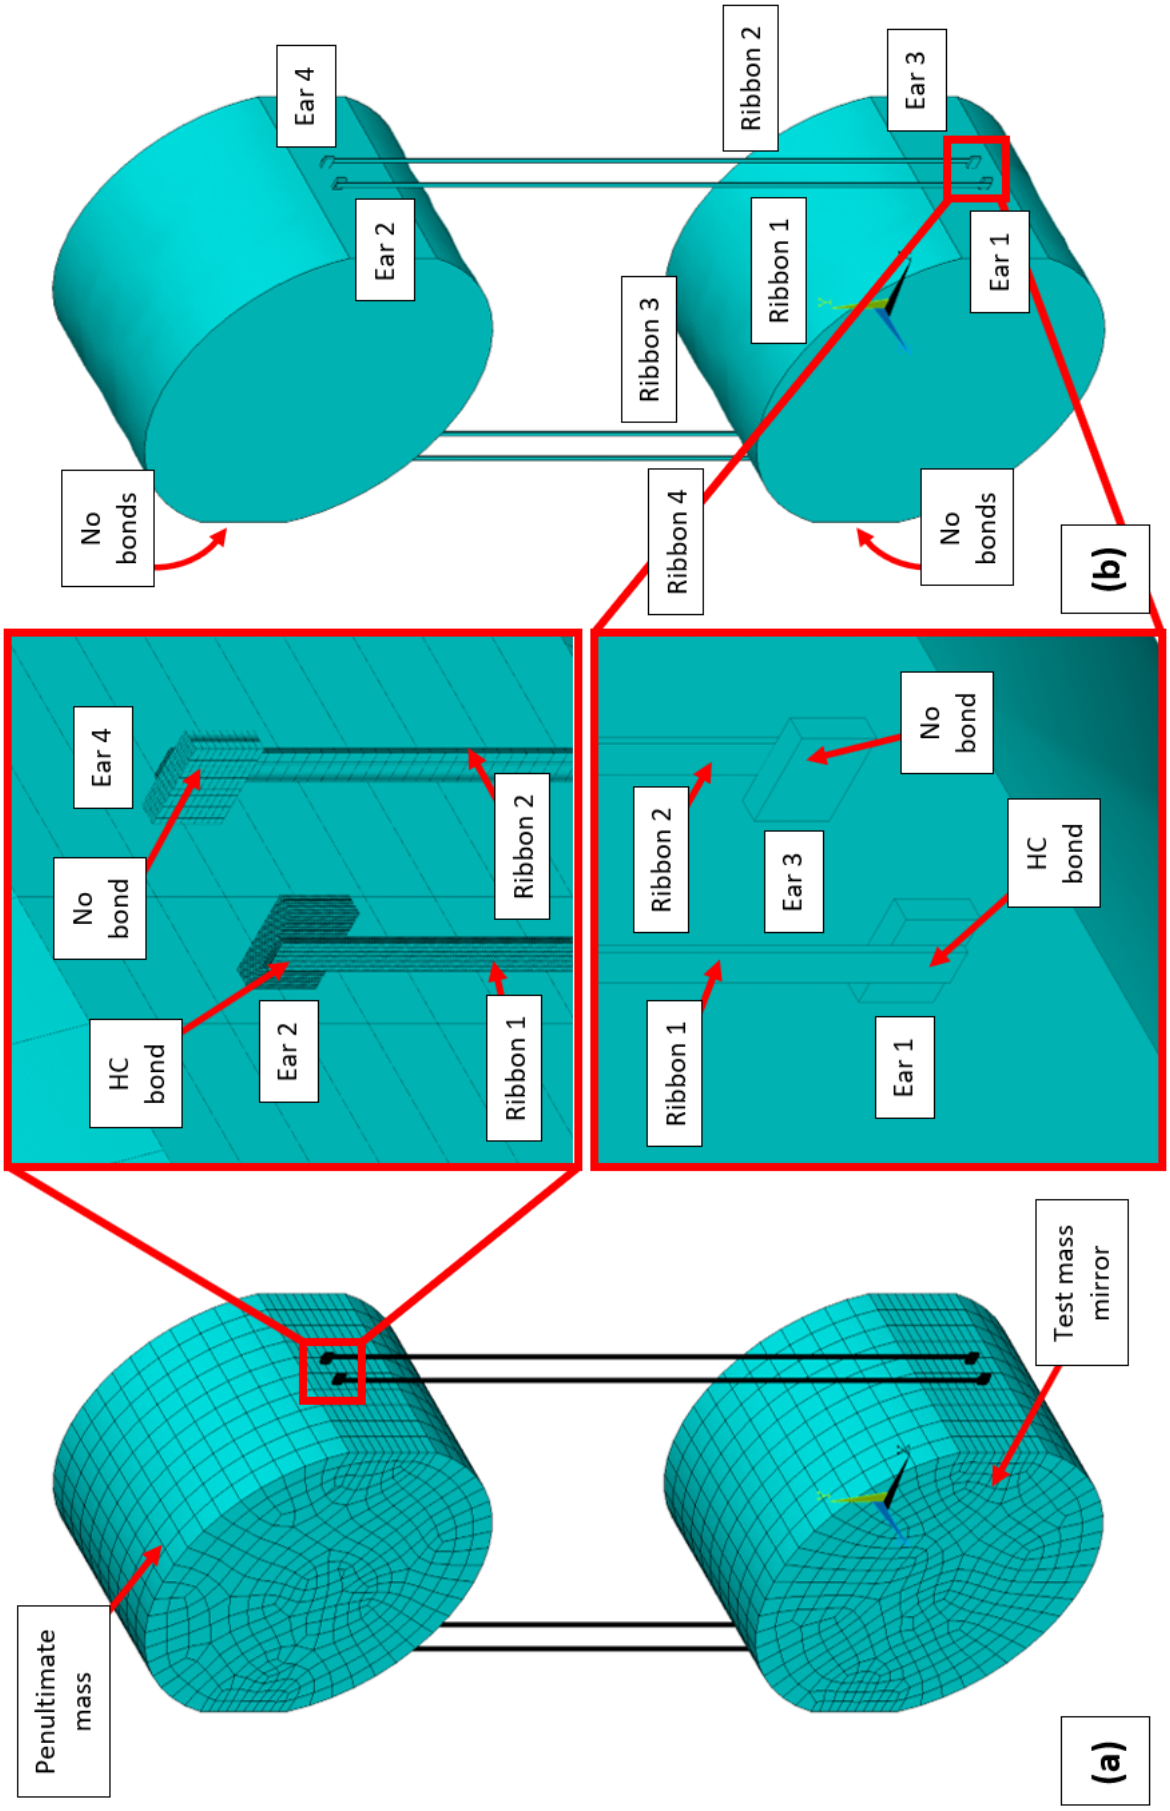

Figure 5.25: (a) 4 ribbon, ET-like, final stage silicon suspension scaled from aLIGO parameters showing model elements. Inset highlights mesh density differences between bonded and unbonded ribbons and ears. (b) Naming convention and HC bond locations in 4 ribbon model showing model volumes. Insets highlights bonds are present only on ribbon 1.

APDL script was written, as shown in appendix E, that built the whole suspension geometry in line with the above requirements, including bonds at the top and bottom between ribbon 1 and ears 1 and 2. The inset of figure 5.25(a) highlights the denser meshing required on ribbon and ear geometries where bonds are present as explained in section 5.5.3.2.

Due to the way the model was built, everything was referenced from the centre of the test mass with the Z direction parallel to the interferometer beam, consistent with all ANSYS work in this thesis. As such, a naming convention was defined in an anticlockwise fashion, as shown in figure 5.25(b), where ribbon 1 is in the +X +Z direction. Ribbon 1 has bonds between the ribbon and top and bottom ears. The remaining ribbons were affixed using an ANSYS multipoint constraint (MPC) bonding algorithm directly onto the ears.

Firstly, it is necessary to compare this 4 ribbon, HC bonded model to theory, where possible. The loaded extension of the 4 ribbon, HC bonded model matches theory to within 0.6 %. This slight increase is accounted for by the softer bond material deforming under shear load, compared to the silicon-silicon ribbon to ear interface in the no bond model.

Furthermore, one can compare this 4 ribbon, HC bonded model to the previously verified 4 ribbon model with no bonds. One expects the frequencies and dilutions of the 4 ribbon model with no bonds and 4 ribbon, HC bonded model to match. The pendulum Z, pendulum X, bounce and violin Z modal frequencies between the no bond and HC bonded models matched to within 0.1 %. Similarly, the dilutions for each modal frequency between the no bond and HC bonded models matched to within 0.3 %. This validates the assumptions that the infinitesimally small bonds inserted onto one ribbon of the HC bonded model do not significantly affect the suspension dynamics.

In order to ensure the HC bonds are modelled correctly, the HC bonded 4 ribbon suspension was compared to model 6 of the single ribbon bonded models in section 5.6.6. Once again, care must be taken to ensure valid comparisons are drawn and subtleties in suspension dynamics are not overlooked. It is necessary to compare bond energy ratios in each mode, which should match between the 4 ribbon, HC bonded model and the single ribbon, HC bonded suspension.

In order to provide a valid comparison for the pendulum modes, the total strain energy of 200 kg 4 ribbon system must be scaled to match the 50 kg mass of the single ribbon suspension. Due to symmetrical mass distribution, this is then divided by 4.

Furthermore, one has to consider the subtlety in the change in bending, where β is 1 for the single ribbon suspension but 2 for the 4 ribbon suspension, since bending took place only at the top for the single ribbon suspension. The bottom bonds can not be compared with single ribbon models, since there is minimal energy in the bottom bond for the single ribbon suspension as no significant bending occurs in this region. However, since bending occurs at the top *and* bottom for the 4 ribbon suspension, the energies in both the top and bottom bonds should be very similar. As such, only the top bonds are compared between the single and 4 ribbon suspensions. The bottom bond energy ratio is *then* compared to the top bond energy ratio for the ribbon suspension

to check validity. For the bounce mode, since no bending occurs, one only divides the total strain energy by a factor of 4 to account for mass scaling.

Pendulum Z	4 ribbon	4 ribbon (scaled)	Single ribbon	4 ribbon energy ratio	Single ribbon energy ratio
Top bond strain energy	6.23×10^{-6}		2.08×10^{-5}	4.20×10^{-4}	4.21×10^{-4}
Bottom bond strain energy	6.23×10^{-6}		4.27×10^{-9}	4.20×10^{-4}	N/A
Total strain energy	1.19×10^{-1}	1.48×10^{-2}	4.95×10^{-2}	% difference (top bond only)	0.3
Pendulum X					
Top bond strain energy	1.02×10^{-5}		2.64×10^{-5}	1.02×10^{-4}	1.00×10^{-4}
Bottom bond strain energy	1.02×10^{-5}		3.29×10^{-6}	1.02×10^{-4}	N/A
Total strain energy	8.01×10^{-1}	1.00×10^{-1}	2.63×10^{-1}	% difference (top bond only)	1.3
Bounce					
Top bond strain energy	1.80×10^{-2}		8.40×10^{-2}	4.49×10^{-6}	5.20×10^{-6}
Bottom bond strain energy	1.80×10^{-2}		8.73×10^{-2}	4.49×10^{-6}	N/A
Total strain energy	1.61×10^{4}	4.02×10^{3}	1.62×10^{4}	% difference (top bond only)	14.6

Table 5.10: Comparison of pendulum and bounce mode bond energy ratios between 4 ribbon and single ribbon HC bonded suspensions. The % difference shown for each mode is the comparison between the top bond energy ratios only.

Table 5.10 shows that the 4 ribbon HC bonded suspension matches the HC bonded, single ribbon suspension in the pendulum modes to within 1.3 %. One notes a larger error of 14.6 % on the bounce mode comparison. This likely arises from the fact there is an off-axis force present on the ear since it is affixed to the side of the masses. This applies a moment on the ear which is not present for the axially loaded in the bounce mode for the single ribbon suspension.

Furthermore, the top and bottom ribbon bond strain energies and ratios match well across all modes shown, indicating equal bending energy near the top and bottom of the suspension as predicted.

5.7.3 Determining violin mode energy distributions

In order to calculate the thermal noise for the violin mode, one must be able to extract the energy distribution in a single fibre or ribbon, in this case ribbon 1 attached with HC bonds. This

is due to the fact a symmetrically built 4 ribbon suspension will produce 4 overlapping violin modes. The energy in the suspension system therefore ends up evenly distributed across multiple fibres rather than distinct single fibres. In reality, this is not the case due to real asymmetry that will exist in even near-identical ribbons or fibres. This has been previously observed in GEO600 analysis [409]. However, it is possible to overcome this issue by artificially stiffening or softening the ribbons by a small amount. One can then check the strain energy contained within the ribbon volume in order to correlate the output violin Z mode from ANSYS to each ribbon [409].

Before applying this to the 4 ribbon, HC bonded suspension, a 4 ribbon suspension with no bonds was first generated. It was decided that ribbon 1 would remain at the standard isotropic silicon Young's modulus of 157 GPa, where ribbons 2, 3 and 4, would be artificially stiffened by 1 % with a Young's modulus of 158.6 GPa applied.

ANSYS output mode	Frequency (Hz)	Strain energy in ribbon				Total ribbon strain energy
		Ribbon 1	Ribbon 2	Ribbon 3	Ribbon 4	
Mode 7	58.97	4008.5	2.0	31.7	0.3	4042.4
Mode 8	58.99	31.1	0.9	4032.8	3.0	4067.7
Mode 9	59.12	0.3	1887.1	0.5	2187.0	4074.8
Mode 10	59.12	2.4	2184.9	3.0	1884.6	4074.9

Table 5.11: Strain energy of bonded suspension with artificially stiffened ribbons 2, 3 and 4, for the violin Z modes.

Table 5.11 shows the resultant strain energy distribution in each mode of each ribbon for this model. This method's validity can be checked by comparing the total ribbon strain energy summed from each mode. One expects mode 7 to contain most of the energy in ribbon 1 due to the remaining ribbons have a higher stiffness and hence higher frequency of violin mode.

Since ANSYS outputs modes in terms of ascending modal frequency order, with ribbon 1 being the "softest" by comparison, one would expect this to be the lowest frequency which is indeed the case. Due to ribbon 1 being softer, the loading of the suspension puts extra energy into the diagonally opposite ribbon 3. This leads to an uneven energy distribution throughout the violin modes, most clearly unbalanced between ribbons 1 and 3 as seen in table 5.11. Ribbons 2 and 4 both have the same Young's modulus and so the energy distribution is more evenly split on these diagonally opposite ribbons due to the more even loading between them. One notes that all the violin Z modal frequencies are closely matched with the 4 ribbon, HC bonded violin Z frequency of 59.04 Hz.

Furthermore, one can compare the total ribbon strain energy summed from each mode for this model against the total strain energy summed for the 4 ribbon, HC bonded suspension with symmetric Young's moduli. The averaged total strain energy across the 4 violin Z modes for the 4 ribbon, HC bonded suspension is 4023.8, comparable to the energies in table 5.11. It is

worth noting here that this averaged value is the total strain energy of the *entire* 4 ribbon, HC bonded suspension and not just the ribbons. This was purely selected for speed and convenience. However, logically the majority of the strain energy *must* be located within the ribbons due to their geometrical deformation in a violin mode. Indeed, one can verify this assumption using this asymmetric Young's moduli model. The total strain energy for this asymmetric model is 4151.2. If one takes mode 7, the violin Z mode of ribbon 1, it is clear to see the strain energy in ribbon 1 alone contributes to 99.2 % of the total strain energy in the suspension system, validating this assumption.

Now that the artificially stiffened model is verified, the no bond model was recreated with the top and bottom bonds attached to ribbon 1. Since the Young's moduli of ribbons 2, 3 and 4 varied by only 1 % from the symmetrical 4 ribbon, HC bonded model it was decided to directly compare the frequencies for the pendulum Z, pendulum X, bounce and violin Z frequencies to ensure model consistency. All frequencies between these two models matched to within 0.1 %.

One can also compare both the top and bottom bond energy ratios in these models for the pendulum Z, pendulum X and bounce modes. All top and bottom bond energies matched to within 0.7 %.

Finally, one can make a valid comparison between this artificially stiffened model's violin Z mode to the single ribbon suspension. The top bond energy ratios matched to within 0.2 % while the bottom bond energy ratio matched to within 0.8 %. Furthermore, the top and bottom bond energies closely matched each other for this model, as expected for the violin Z mode geometrical deformation.

The above work has created and validated a 4 ribbon, HC bonded, ET-like silicon suspension with an understanding of the energy distribution in the suspension system. From the information gained from the ANSYS energy distribution analysis, this allowed for the calculation of suspension thermal noise.

5.8 Silicon suspension thermal noise

This section determines the suspension thermal noise of the 4 ribbon, HC bonded silicon suspension, shown in the previous section, as well as applying some feasible suspension design parameters that may improve this.

For the remainder of this chapter, the pendulum Z mode is herein referred to as the pendulum mode since the pendulum X direction is not a main contributor to detector noise compared to the primary modes of interest. The are 3 primary modes of interest for suspension thermal noise. The most obvious is the direct length coupling from the pendulum Z mode into the interferometer. As discussed in section 2.4.3 of chapter 2, the bounce also contributes to detector displacement noise through the 0.1 % cross coupling. Finally, the violin Z mode is also a significant source of noise where the mass term in equation 2.13 contributes to detector displacement

in the primary Z direction.

5.8.1 4 ribbon, HC bonded model with ribbon to ear bonds

A MATLAB script was written, shown in appendix F, that calculated the thermal noise at a set temperature, T, over a user-defined frequency band. This script calculated the surface, bulk and thermoelastic loss. From the ANSYS models above, one can also calculate the bond loss contribution. The bond loss contributions were calculated for each mode.

For the pendulum mode, the top and bottom energy ratios were taken from table 5.10 and summed together. Since this represented the bond loss contribution from 2 ears and one ribbon, this was then multiplied by 4, producing the total bond energy ratio. This bond energy ratio was then multiplied by the temperature-dependent bond loss to determine the bond loss contribution. The energy ratio for the remaining loss mechanisms, which was called the "remainder energy ratio", was the total energy in the system minus the total bond energy ratio. Therefore, the total loss was calculated by summation of the surface, bulk and thermoelastic losses, multiplied by the remainder energy ratio.

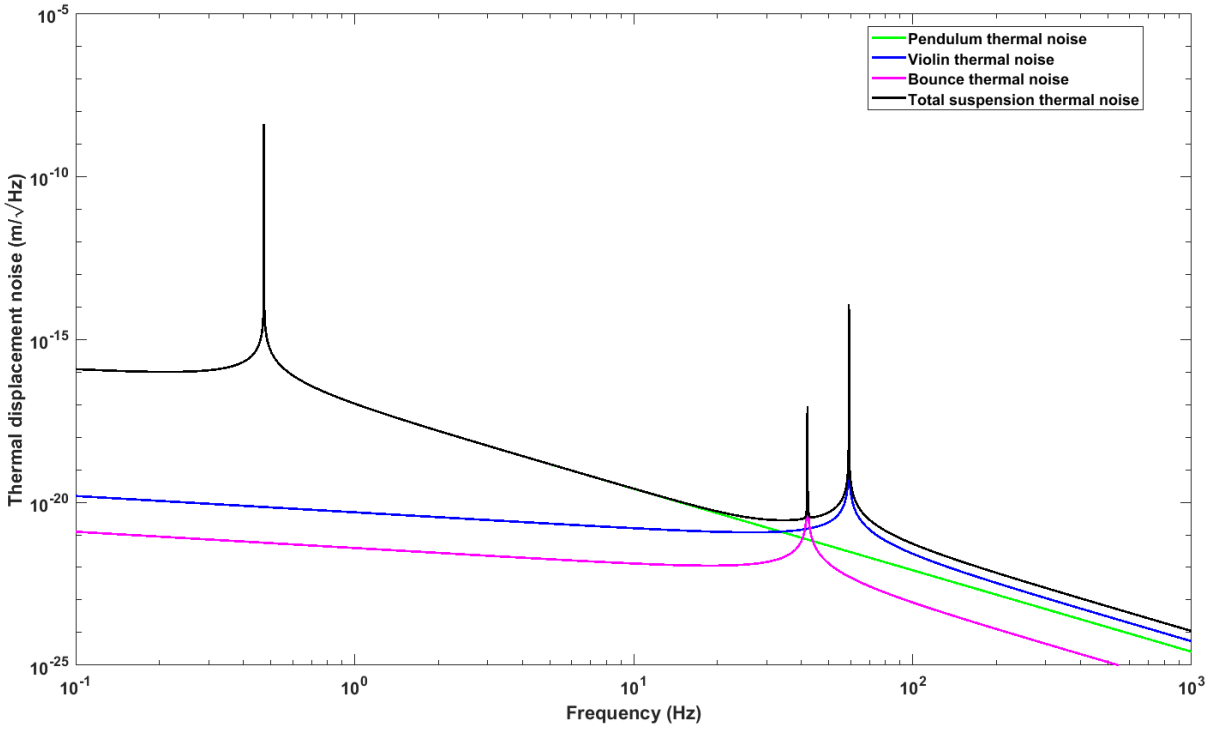

Figure 5.26: Suspension thermal noise of 4 ribbon, HC bonded silicon suspension showing modes of interest and their individual contribution to the total suspension thermal noise at 20 K.

For the violin modes, the process is the same to produce the violin mode bond loss contribution, except the summation of the other loss mechanisms is multiplied by 2 and then multiplied by the remainder energy ratio, consistent with equation 2.14 in chapter 2. The bounce mode

bond loss contribution is calculated in the same way as the pendulum mode, except the surface loss value is multiplied by 0.5, before summation of the other loss mechanisms and then multiplication by the remainder energy ratio as derived in [184]. Each violin mode from each fibre is treated separately and summed in the final noise calculation appropriately. One assumes identical ribbons and so the total sums at the same frequency, as shown in figure 5.26.

With this information, the total suspension loss and dilution for each of the modes that couple directly into detector noise was calculated. The script then uses uses equation 2.1 in chapter 2, derived from the Fluctuation-dissipation theorem, to calculate the suspension thermal noise for each mode. Total suspension thermal noise was determined via the quadrature sum of each independent resonant mode.

Figure 5.26 shows the suspension thermal noise at 20 K from each of the 3 modes of interest, including the total suspension thermal noise, which is the sum of the quadrature of the thermal noise of each mode. This is a proposed operating temperature of the low temperature regime for the ET-LF detector [127]. This is the first ever calculation of suspension thermal noise based on a physically assembled silicon detector and scaled to an ET-like detector size.

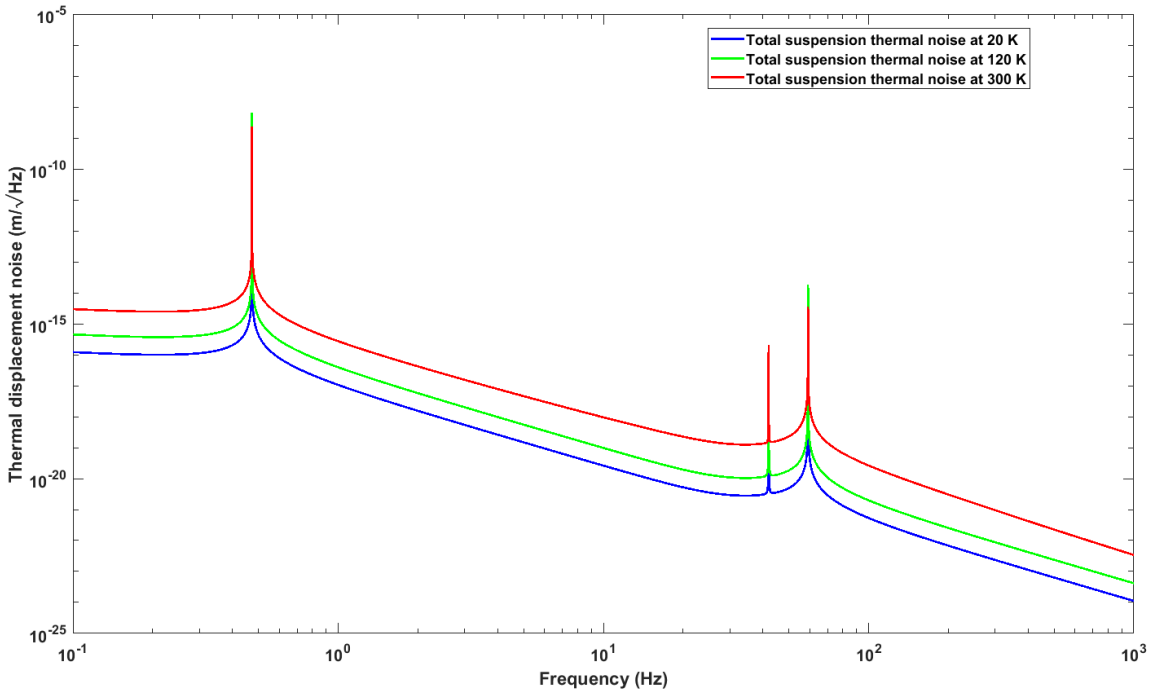

Figure 5.27: Suspension thermal noise of 4 ribbon, HC bonded silicon at 20 K, 120 K and 300 K.

Here one can see why the frequencies of the 3 primary modes are of interest. The pendulum mode dominates at the lowest of frequencies, acting as a limit to the low frequency sensitivity of the detector. Furthermore, the bounce and violin mode resonant peaks typically dominant over *all* detector noise sources, as can be seen for aLIGO in figure 2.1 of chapter 2. Fortunately, these frequencies are normally constrained to a very narrow frequency bands and the geometry of suspension ribbons is chosen to push the frequencies where they impinge less.

Figure 5.27 shows the advantage of operating such a detector at cryogenic temperatures. Here the ET-LF low temperature conduction-only regime of 20 K is compared to the higher temperature radiative and conductive regime of 120 K as well as at a room temperature of 300 K [127]. At 10 Hz, there is a factor of 36 reduction in displacement noise between 300 K and 20 K.

For the remainder of this work, all silicon suspension thermal noise plots are calculated at 20 K.

One can directly compare the modelled suspension thermal noise to the ET design sensitivity curve. Furthermore, with the ability to directly model HC bonds now demonstrated, this allows one to visualise the contribution from bond loss to the suspension thermal noise. These comparison are shown in figure 5.28.

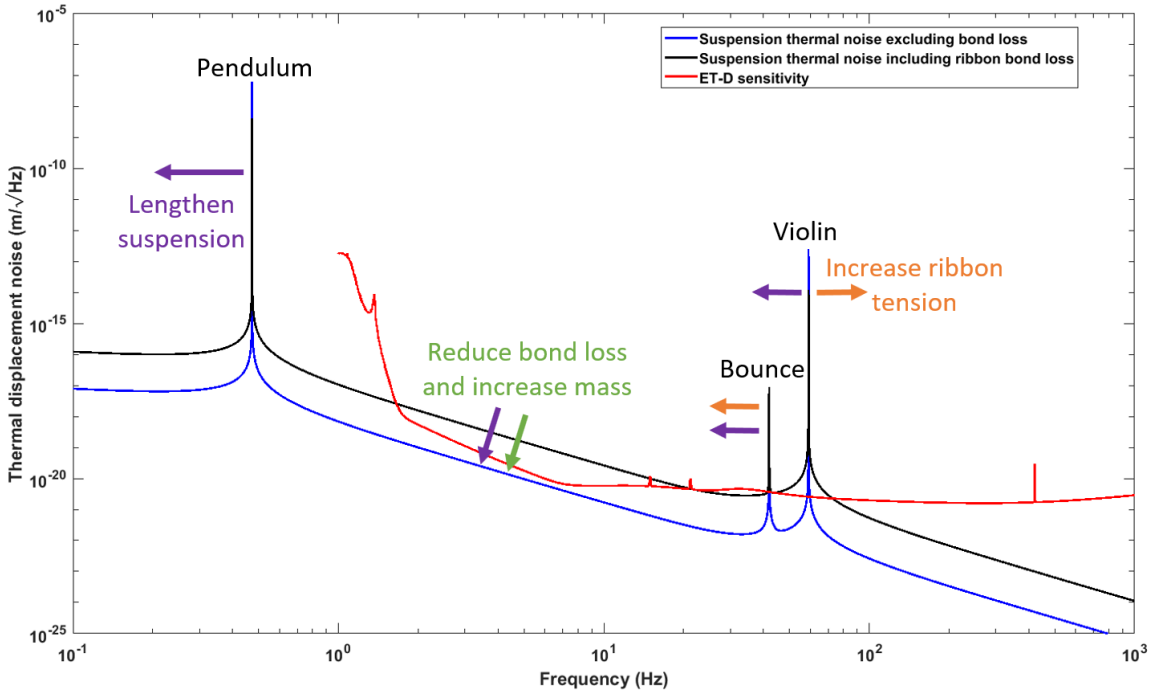

Figure 5.28: Suspension thermal noise of 4 ribbon, HC bonded silicon with and without HC bond losses compared to the ET design sensitivity curve provided from 1–1000 Hz [127]. Annotated arrows show effect of suspension design on thermal noise improvements.

Figure 5.28 shows the HC bonded model to be completed dominated by thermal noise arising from the HC bonds. The suspension thermal noise requirement for this 4 ribbon silicon suspension currently exceeds the design requirement by a factor of 4 at 10 Hz.

As discussed in section 2.3.3.4 of chapter 2, research is needed in order to significantly reduce this bonding loss while retaining sufficient bond strength in order to support the suspension. If one could produce theoretically lossless bonds, or truly monolithic silicon suspensions, the blue curve shown in figure 5.28 would represent this suspension thermal noise. At this frequency, one reduces the suspension thermal noise by a factor of 15 by removing the HC bond contribution. While this may seem promising by appearing to exceed the ET design sensitiv-

ity curve, it should be noted that this red curve represents *total* thermal noise budget from all contributing sources and not just suspension thermal noise.

However, this model uses silicon ribbons that are only 1.2 m long, double the aLIGO suspension length, thus providing a more conservative estimate of thermal noise compared to the ET design sensitivity curve, which is of a 1.5 m suspension [127]. As discussed in section 2.5.1 of chapter 2, while silicon ribbons or fibres are not currently available in this length, there is promising ongoing research investigating techniques that may yield longer suspensions. Lengthening the suspension ultimately allows for improved thermal noise due to the gain in dilution factor and reduction in pendulum mode resonant frequency.

5.8.2 4 ribbon, HC bonded model with ribbon to ear and ear to mass bonds

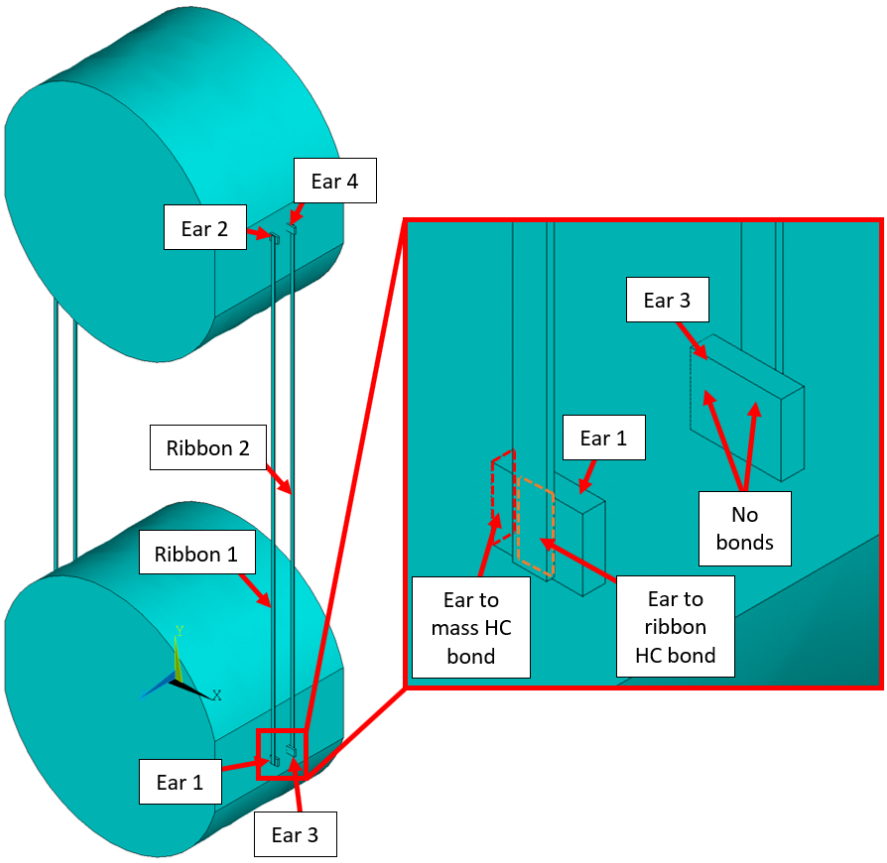

Figure 5.29: 4 ribbon, HC bonded, final stage silicon suspension with bonds between ears 1 and 2 and the test and penultimate masses, as well as ribbon 1.

The 4 ribbon, HC bonded silicon suspension only has bonds between the ribbons and ears and assumed the ears are part of the penultimate and test mass geometries. To investigate the effects of adding in another set of bonds, another model was built placing bonds between the ears and masses. These bonds were also 361 nm thick and had the same properties as the ribbon to ear

bonds with the exception of a larger surface area, 5 mm by 10 mm. Again, for computational efficiency ear to mass bonds were only implemented on ears 1 and 2. The model is shown in figure 5.29.

One noted the pendulum, bounce and violin mode frequencies remained unchanged with the introduction of this second pair of bonds. Due to the infinitesimal size of the ear to mass bonds, no significantly modal frequency change was expected. Similarly, there was no change in dilution between this model compared to the ribbon to ear bond only. To further check results, the bond energy ratios for the top and bottom bonds on the ribbon were compared to the previous ribbon to ear bond only model. For all 3 modes, these bond energy ratios were found to match to the 4th decimal place.

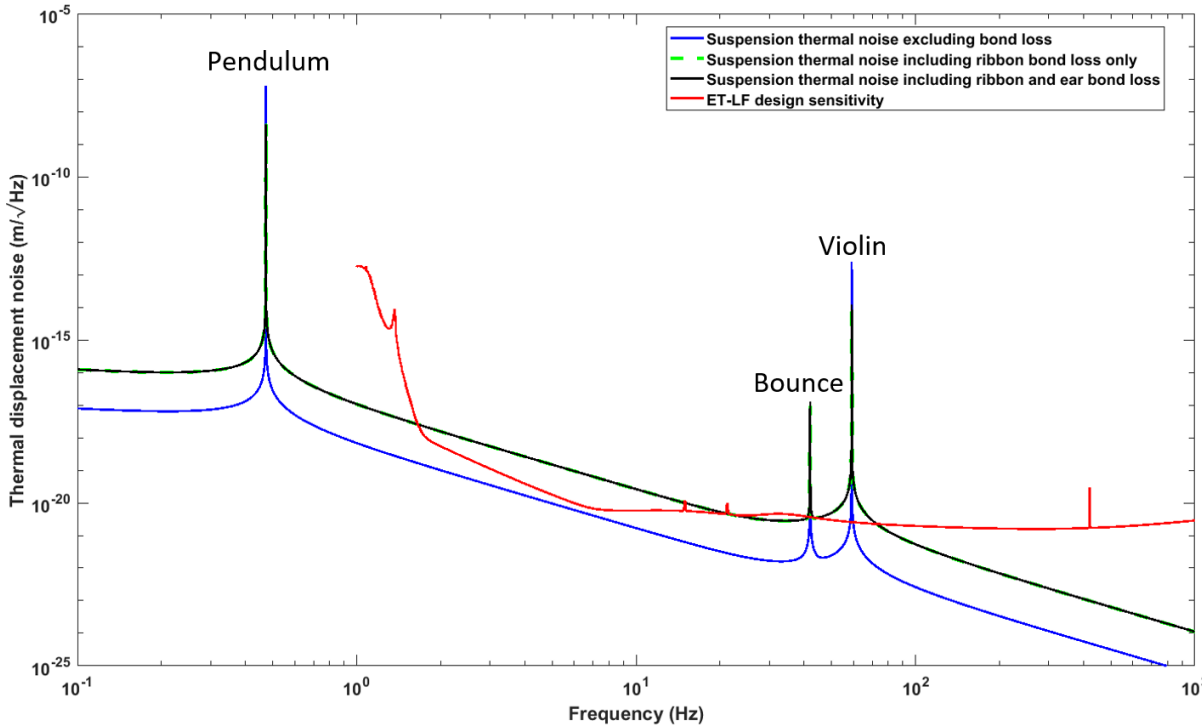

Figure 5.30: Suspension thermal noise of a 4 ribbon, HC bonded, final stage silicon suspension with bonds between ears 1 and 2 and the test and penultimate masses, as well as ribbon 1. The green dashed, made bold for clarity, represents the suspension thermal noise including ribbon only bond loss from figure 5.28.

The suspension thermal noise curve for this model with bonds between mass, ears and ribbons compared to the model with only ear to ribbon bonds and the model with no bond loss is shown in figure 5.30.

The energy in the ear to mass bonds is lower by a facor of ≈ 100 than the energy between the ear to ribbon bonds for the pendulum and violin modes which is to be expected. Conversely, the energy in the ear to mass bonds is higher, by a factor of ≈ 1.2, than the energy between the ear to ribbon bonds for the bounce mode. This increase in energy is due to the asymmetrical force applied from the test mass bounce mode, which produces a torque pulling on the 20 mm

wide ears attached onto the primary mass flats.

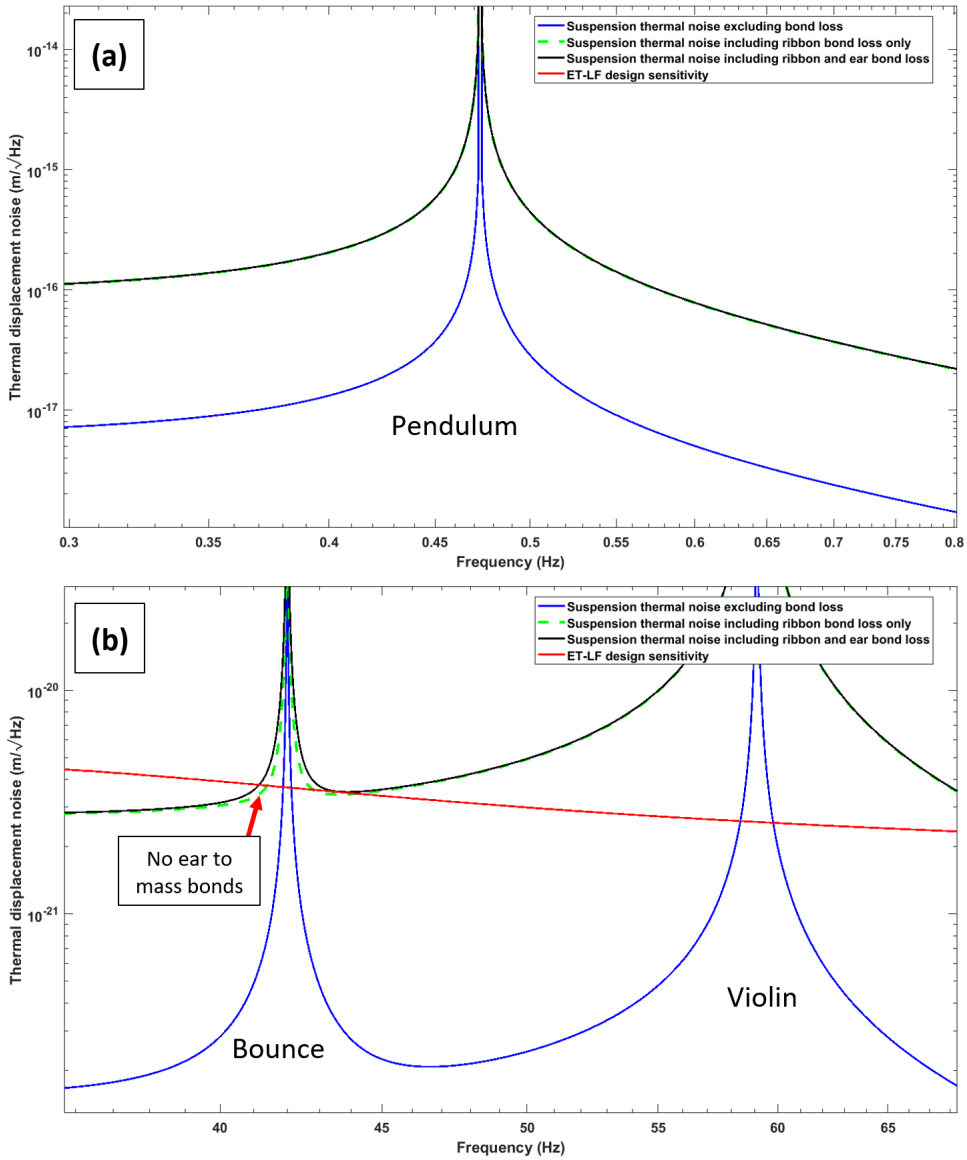

Figure 5.31: (a) Pendulum mode comparison for models with bonds between mass, ears and ribbons, ears and ribbons only and no bonds. (b) Bounce and violin mode comparison for models with bonds between mass, ears and ribbons, ears and ribbons only and no bonds.

From figure 5.31(a) and (b), there is an infinitesimal change in thermal noise in the pendulum and violin modes due to the minimal energy within the ear to mass bond regions from the geometrical deformation of these modes. However, as can be seen in figure 5.31(b), this translates to a noticeable change in the bounce mode. Indeed, due to the Fluctuation-dissipation theorem, this change in bounce mode energy should increase the thermal noise by a factor of $\sqrt{2}$.

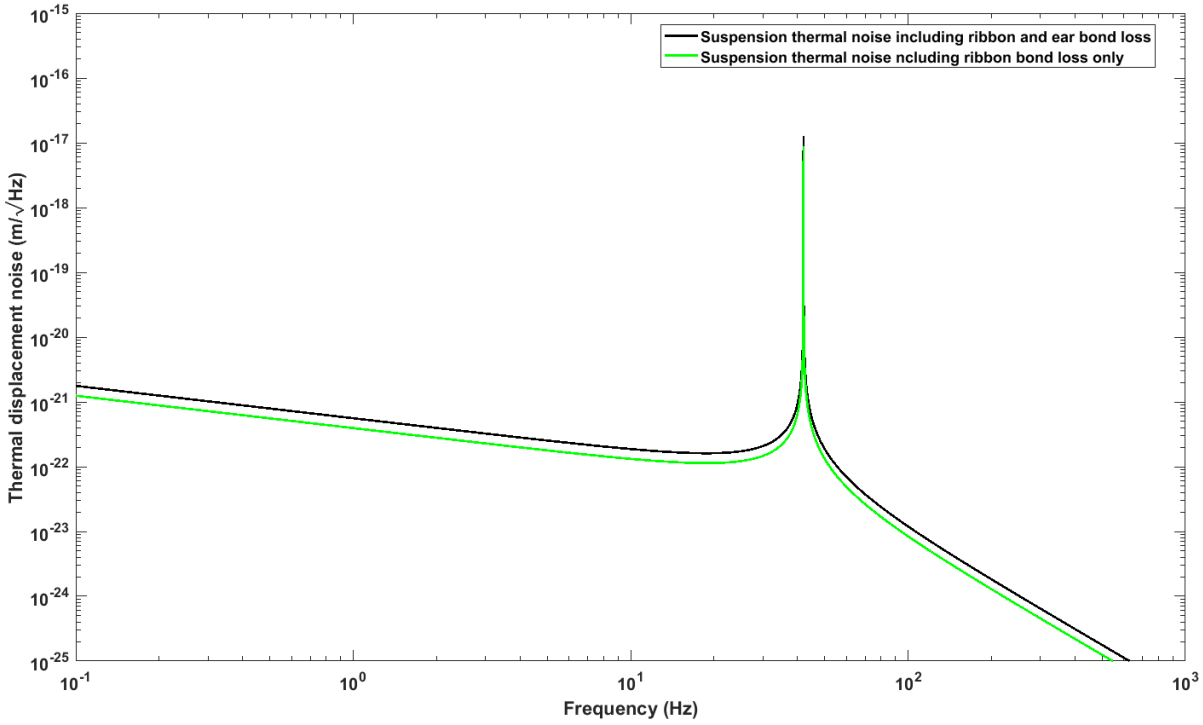

Figure 5.32: Bounce mode thermal noise contribution for 4 ribbon, HC bonded models with ribbon to ear bonds only compared to ribbon to ear and ear to mass bonds.

Figure 5.32 shows the bounce mode suspension thermal noise contribution for both models. The bounce mode thermal noise at 10 Hz for the suspension with ribbon and ear bonds is $1.88 \times 10^{-22}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$. The bounce mode thermal noise at 10 Hz for the suspension with ribbon and ear bonds is $1.32 \times 10^{-22}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$. This shows that, off resonance, there is indeed a factor of $\sqrt{2}$ higher thermal noise for the suspension that has bonds between both the ears and ribbon and the ears and masses.

Figures 5.30, 5.31 and 5.32 show that HC bonding of ears to the masses for this particular suspension model has a small effect on the total suspension thermal noise. At 10 Hz, the introduction of ear to mass bonds, compared to only ribbon to ear bonds increases the suspension thermal noise by 0.45 %.

This only highlights the comparatively significant energy stored in the ribbon to ear bonds which, when combined with the relatively high mechanical loss of HC bonds, are detrimental to the thermal noise performance of the detector as shown in figure 5.28.

5.8.3 4 ribbon, HC bonded model at double suspension stress

The standard suspension stress in the above models has been kept at 40 MPa. However, as highlighted in figure 5.28, one can reduce suspension thermal noise by increasing the tension in the ribbons, which should also push the violin mode frequency up and bounce mode frequency down. To first order, the pendulum frequency is affected by length changes only.

Looking at the results from figure 3.73 in section 3.9 of chapter 3, and with further surface treatment investigations, it is not unreasonable to assume one can feasibly double the operating stress of silicon from 40 MPa to 80 MPa.

As such, a new model was created by decreasing the cross-sectional area of the ribbons, in order to increase the stress in the suspension to 80 MPa. This also has the beneficial effect of reducing the bond area between the ribbon and the ears. However, this must be balanced with the bond stress itself, which for this model is now increased to 6.1 MPa, reducing the safety factor to 2. Scaling calculations provided a new ribbon cross sectional area of 5.26 mm by 1.17 mm. This resulted in a 29 % reduction in bond area between the ribbon and ears. Furthermore, thinner ribbons should lead to better dilution due to a smaller second area moment of inertia, I, which dominates as $t^{\frac{3}{2}}$, where t is the thickness of direction of bending.

5.8.3.1 40 MPa v 80 MPa suspension thermal noise without bonds

A model was created with no HC bonds in order to investigate the direct effects of doubling the stress in the suspension.

Figure 5.33 shows the comparison of the 80 MPa model to the 40 MPa model with no bonds. The 80 MPa model gains in thermal noise by a factor of 1.3. This is consistent with theoretical expectations due to the increase in dilution factor and reduction in surface loss for thinner ribbons, combined with improvement in thermoelastic loss due to higher ribbon stress.

The dilution for the 80 MPa model is also significantly improved, with a pendulum mode dilution of 70.5 an a violin mode dilution of 34.2. From equation 2.23 in chapter 2, these improvements in dilution, and hence thermal noise, come from the factor $t^{\frac{3}{2}}$, where t is the ribbon thickness in the bending direction as well as the factor of $\sqrt{T}$.

If one increases the suspension stress, the suspension thermal noise decreases. It is clear to see from the work in chapter 3, strengthening the working stress of suspension fibres or ribbons is both feasible and, from the work in this chapter, directly beneficial to the suspension thermal noise of cryogenic silicon suspensions.

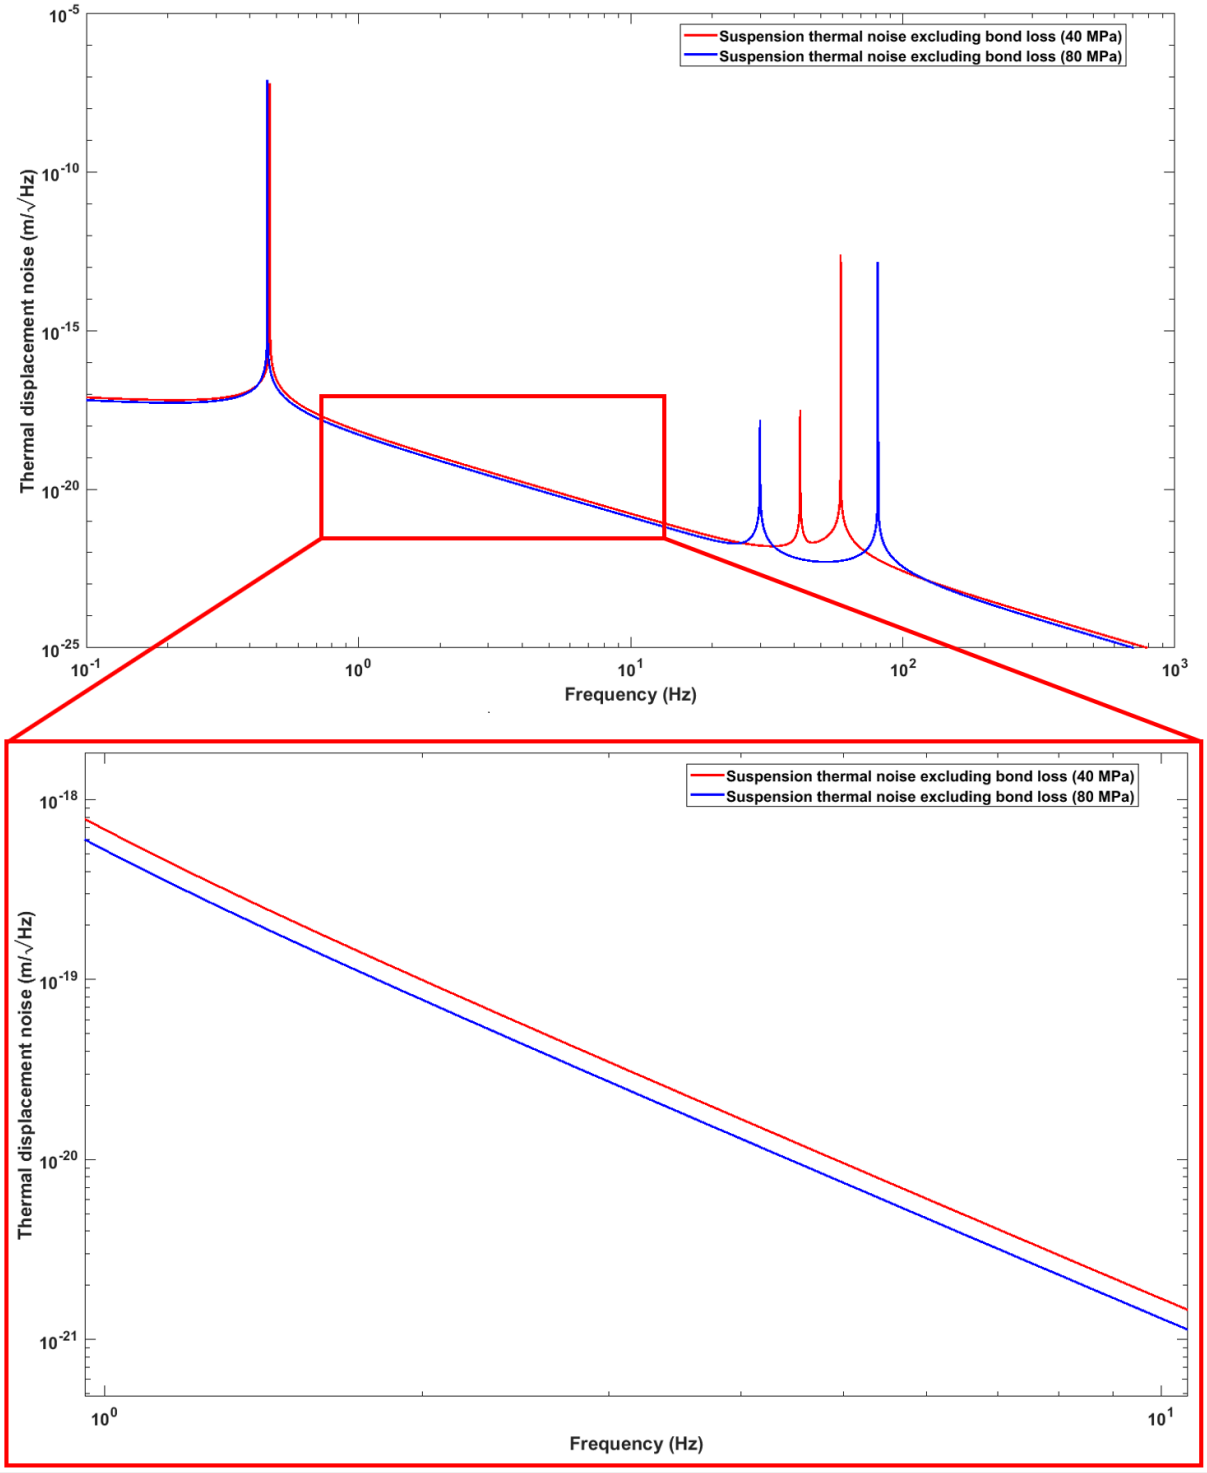

Figure 5.33: Suspension thermal noise of 40 MPa and 80 MPa silicon suspensions without HC bonds.

5.8.3.2 40 MPa v 80 MPa suspension thermal noise with bonds

The previous model was re-created with the introduction of HC bonds. For speed of computation, this model only included ribbon to ear bonds on ribbon 1 and did not include ear to mass

bonds. This allows for direct comparison to the 40 MPa bonded and unbonded models in figure 5.28.

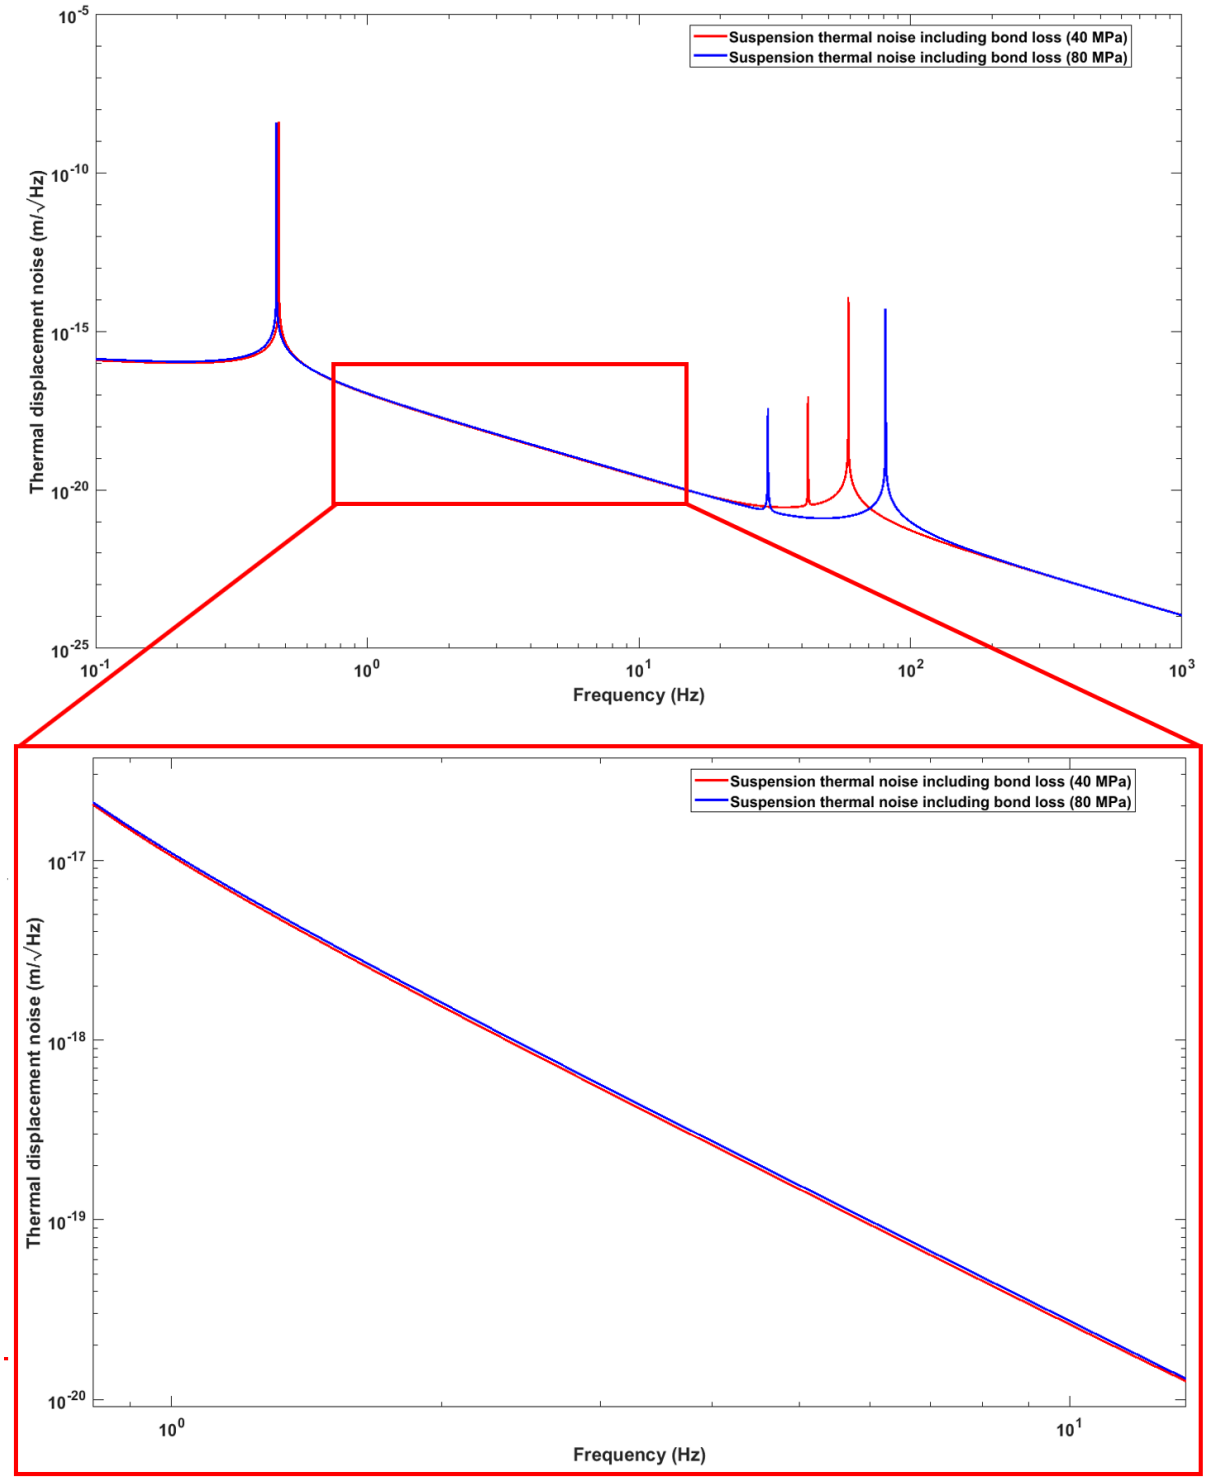

Figure 5.34: Suspension thermal noise of 40 MPa and 80 MPa silicon suspensions with HC ribbon to ear bonds.

Figure 5.34 shows the comparison of the 80 MPa model to the 40 MPa model with HC ribbon to ear bonds.

At 10 Hz, the 80 MPa model with bonds results in *higher* thermal noise compared to the 40 MPa model. This appears to defy simplistic assumptions that higher ribbon stresses lead to better thermal noise performance. When one considers the effect of the HC bonds in the suspension system, it is clear to see this relationship is more complicated.

The higher thermal noise for the 80 MPa bonded suspension is believed to arise from the fact that more energy is being put into the lossy HC bond regions. While one has reduced the bond area by thinning the ribbons, thus reducing the bond loss contribution to the thermal noise, this increases the static stress in the bond. Since the bond is softer than the silicon (by a factor of 10), when bending occurs, more strain energy is present in this region. Indeed, the energy ratio in the bounce mode was found to double between the 40 MPa and 80 MPa models. This leads to a 5 % increase in thermal noise at 10 Hz.

This highlights that bond loss improvements through new bonding techniques may not be sufficient alone in reducing suspension thermal noise to acceptable levels. As discussed in section 2.5.1 of chapter 2, research is required to investigate silicon manufacturing and shaping in order to have better control over where bending occurs in the suspension system. Section 1.5.4 of chapter 1 demonstrates the fact that, even with ultra low loss suspension materials, one still must have sufficient control of energy distribution throughout the final stage suspension in order to keep strain energy away from regions of high mechanical loss. This is the path that must be pursued in order to realised low thermal noise 3G cryogenic suspensions.

Comparing the 80 MPa models with and without bond loss, one notes a factor of 21 difference, at 10 Hz, in suspension thermal noise. This once again highlights the importance of developing new, low loss bonding techniques alongside higher stress suspensions.

5.8.3.3 40 MPa v 80 MPa suspension frequencies

One notes, for both figures 5.33 and 5.34, the expected shifts down in the bounce mode frequency and up in the violin mode frequency for the 80 MPa models, consistent with the increasing ribbon tensions. The bounce mode frequency has reduced from 42.0 Hz to 29.9 Hz and the violin mode frequency has increased from 59.1 Hz to 81.0 Hz. The increase in violin mode frequencies is less of a concern since the ET-HF detector is expected to be more sensitive at frequencies >30 Hz [126]. However, this reduction in bounce frequency is of greater concern as it begins to impact into the detection band where ET-LF is designed to be most sensitive.

Selection of fibre or ribbon stresses and cross-sectional geometries will need to be carefully considered for the purposes of a xylophone detector design. One must strive to increase the suspension stress while keeping the resultant bounce mode frequency as high as possible, out of the ET-LF sensitivity band and up into where the ET-HF band sensitivity can take over. This has not previously been a consideration for single detector fused silica suspensions, where the aim was always to push the vertical frequency down to below the low frequency detection band of the instrument.

5.9 Summary of cryogenic silicon suspension work

Work has been conducted in ANSYS FEA to determining the tolerable stress in the single ribbon suspension, built in chapter 4, arising from misalignment in supporting and assembly structures. Taking a conservative breaking stress of 33 MPa, the largest 250 µm misalignment in the bottom block to mass attachment would not increase the suspension stress more than 56 % of this breaking stress. Similarly, for the bottom bond to bottom block placement, a very high tolerance of ± 3 mm in the X direction is acceptable in the bond misalignment. Attachment blocks placed on the same side reduces stress in the ribbon and aids practical implementation of bonding and suspension design. While more tolerant to bonding misalignment in the X direction, the suspension is less tolerant to external forces or shocks applied through this thick ribbon direction, with only a 3 mm displacement at the bottom block likely to cause suspension failure. Furthermore, while not considered in this thesis but noted during the course of this work, future work should investigate cross coupling from the stiffer pendulum X direction into the softer pendulum Z direction arising from tooling tolerancing and silicon ribbon manufacturing mismatches.

Modelling of silicon's anistropic axis properties was conducted, noting the correct Young's modulus values must be selected for ribbon bending when comparing model to theory. However, it was noted that for longer suspensions, the anisotropic effects significantly reduced, producing similar levels of accuracy to isotropic models.

The first direct modelling of HC bonds between silicon substrates in a single and 4 ribbon suspension has been demonstrated. This modelling allows for visualisation of complex bond deformation in a variety of suspension modes and was shown to be an improvement on the work that utilised extrapolation techniques to determine the energy ratio in an HC bond.

By being able to determine energy ratios in HC bonds, one can then investigate the suspension thermal noise contribution of HC bonds compared to other thermal noise contributions. Iterative development of a single ribbon suspension, scaled from the as-built suspension in chapter 4, up to an ET-like mass, demonstrated that this suspension was dominated by bond loss.

Translating this to suspension thermal noise in a 4 ribbon suspension required careful selection of bending points and ear placement on the masses. Suspension thermal noise was significantly degraded by 1.5 orders of magnitude due to the introduction of HC bonds between the ribbons and ears of the suspension. This highlights the need to significantly reduce bond loss between silicon suspension elements.

The introduction of HC bonds between the ears and masses did not significantly increase the thermal noise, compared to the ribbon to ear bonds. Most of this additional thermal noise was attributed to the bounce mode and corresponded to a $\sqrt{2}$ increase.

By increasing the stress in a 4 ribbon suspension, one can reduce the suspension thermal noise. By doubling the 40 MPa suspension stress to 80 MPa, one can reduce the thermal noise by 22 % at 10 Hz. However, this model improvement relies on a bond safety factor of only 2 and a suspension safety factor of only 1.5. As discussed in section 2.5.4 of chapter 2, brittle

materials such as silicon generally require a slightly higher safety factor due to the lack of a well-defined yield point [228]. Increasing the stress of silicon suspensions may also require surface treatments similar to those discussed in chapter 3. Future analysis is required to ensure surface treatments do not detrimentally affect the surface loss of the suspension ribbons or fibres, since the calculations in this chapter have used previous high quality surface finish data.

By introducing HC bonds between the ribbon and ears for the 80 MPa suspension, it was found that the suspension thermal noise actually increased by 5 % at 10 Hz, compared to the equivalent 80 MPa. While counter-intuitive to the theory of suspension thermal noise and stress, this increase was attributed to the increase in energy within the HC bond region arising from the increase in ribbon stress. This highlights the need to be able to shape silicon ribbons or fibres, as well as ears, in order to control the strain energy distribution in the final stage suspension.

This body of modelling work, based on the experimental results from previous chapters, demonstrates that 3G detectors constructed using current manufacturing and bonding techniques, with currently envisaged surface treatments, do not meet design sensitivity requirements for suspension thermal noise. To improve this performance, increasing the strength of silicon while maintaining low mechanical loss, and developing low loss jointing methods, are the key challenges the field needs to overcome.

Being able to increase tension in the fibres by increasing their strength will permit improved thermal noise as demonstrated. Building on the work demonstrated in this thesis, lowering the loss of jointing techniques in the construction of silicon suspensions remains a crucial ongoing and future research area. The materials, geometries and construction methodology used for 3G instruments may ultimately completely define the mechanical design of the final stage suspension.

Chapter 6

Conclusions

Significant advances in physics and experimental techniques have allowed for the direct detection of gravitational waves through laser interferometry. This technique allows for detection over a wide frequency band, ensuring a plethora of astrophysical sources can be detected and better understood, particularly when used in conjunction with more traditional electromagnetic techniques. By building bespoke high and low frequency, or "xylophone", next generation detectors one can further improve the detection sensitivity across the full bandwidth of interest. By continuing to develop a range of gravitational wave detectors in operation and in proposed design, this will allow humanity to better understand the universe in which we live.

The fundamental sensitivities of these detectors are limited by an array of physical noise sources that must be understood and mitigated. This thesis focuses particularly on suspension thermal noise and its improvement by the use of novel low mechanical loss materials alongside direct thermal noise reduction through cryogenic cooling. By reducing the temperature through cryogenic cooling, one can directly improve suspension thermal noise by removing thermal energy from the suspension system. Cryogenic cooling still presents challenges for practically cooling suspension structures and test masses.

Reduction in temperature then also exploits the multiple advantageous properties of silicon; the coefficient of thermal expansion nulls, the high thermal conductivity and the low mechanical loss. Manufacturing technologies will need to be developed in order to allow suspension resonant frequencies to be controlled to place them outside the gravitational wave detection frequency bands of interest.

Development of silicon manufacturing technologies will also allow for lengthening of the suspension chain, which will increase dissipation dilution, improving suspension thermal noise. The geometries available from laser cutting were shown to be limited and restricted to silicon wafers of certain thickness. However, alternative production techniques, still to be demonstrated at scale, provide promising abilities to produce silicon ribbons or fibres required 3G detectors. Furthermore, increasing the suspended test mass will also increase dissipation dilution, lowering thermal noise further and helping to push suspension resonant frequencies out of the detection

band. However, it is vital that the strength of silicon is increased in order to support the proposed large-scale test masses for ET-LF with low-noise suspension fibres or ribbons. Furthermore, reducing the cross-sectional area of the suspension geometry in order to improve dilution and reduce thermoelastic loss also requires higher ultimate tensile strengths of silicon.

Work was conducted to investigate the ultimate tensile strength of silicon with an average median value of 60.3 MPa measured for the *control new* laser cut silicon sample set. Experiments were carried out to determine if strengthening of silicon was possible through three techniques; chemomechanical polishing, argon ion etching and wet chemical etching. All surface treatments were found to increase the strength of the silicon samples, with most approximately doubling the median strength of the silicon sets, due to the removal of surface damaged induced by laser cutting. However, chemomechanical polishing will rely on substantial physical contact with suspension fibres or ribbons which may not be practical or scalable for circular fibre geometries. By comparison, argon ion etching shows promising results for silicon surface and strength enhancement. This surface treatment has the advantages of being non-contact, with the potential for scalability, the treatment of circular fibre geometries and shaping of suspension element geometries. However, the surface finish may result in undesirable surface loss, although this will have to be investigated. Likewise, wet chemical etching shows promising results for silicon surface and strength enhancement. The most promising chemical etch set was the ITM KOH etch, with the strongest sample a factor of 3 greater than the relevant control. Furthermore, this sample set had the smallest IQR meaning fewer strength outliers, increasing confidence in the ability to run 3G detector suspensions at higher stress or lower safety factors. Chemical etching surface treatment is the most practical and scalable treatment for silicon ribbons or fibres. Wet etching allows all surfaces of a ribbon or fibre to be accessed and treated evenly across the sample length, regardless of geometry, provided etchant can sufficiently access and flow. Also, chemical baths can be and are built to dimensions scaling to 10s of metres. It is recommended chemical etching be pursued as a surface treatment method for strengthening silicon ribbons or fibres. Chapter 3 and Appendix A detail how to proceed specifically with KOH etching. It is also recommended that further investigations be made into practicality and scalability of argon ion etching, or similar ion source experiments to find ways of treating suspension-scale fibres or ribbons. Future work should also investigate strength testing within cryogenic regimes to investigate any potential effects lower temperature has on ultimate tensile strength.

This thesis has also demonstrated what is believed to be the world's first HC bonded silicon suspension. This work has demonstrated the successful hang of this suspension down to cryogenic temperatures, stabilised at proposed 3G detector operating temperatures of 20 K and 120 K and back up to room temperature without failure, surviving for over 8 months, with the suspension remaining intact at the time of writing. Iterations of each suspension have shown the importance of the coefficient of thermal expansion in jointing prototype silicon suspensions and highlights the significant challenge posed by cryogenics in the context of sensing, with un-

wanted and poorly characterised heat sources present throughout the system. Further work is required to better characterise the cryostat and improve any cryogenic wiring techniques for future iterations. This work has also demonstrated the ability to extract thermal conductivity values from this cryogenic experiment. While the results are not in agreement with other experimentally measured thermal conductivities of similar silicon, a number of factors inhibiting phonon transport and minimising the cross-sectional conductive area are believed to be contributing to these low measured thermal conductivities. Future work should look to better characterise these silicon ribbons, particularly to understand the crystal structure across the full cross-section to provide insight on the surface damage and, in particular, to provide a better understanding of thermal conductivity for long, thin suspension samples with a finite surface roughness.

Modelling work has been conducted in ANSYS FEA to determine the tolerable stress in the HC bonded ribbon suspension, arising from misalignment in supporting and assembly structures, which led to the first successful prototype suspension without failure. While not considered in this thesis but noted during the course of this work, future work should investigate cross coupling from the stiffer pendulum X direction into the softer pendulum Z direction arising from tooling tolerancing and silicon ribbon manufacturing mismatches. It was noted that for longer suspensions, the anisotropic effects of silicon were significantly reduced, producing similar levels of accuracy to isotropic models which may make future modelling of suspension more efficient. This work demonstrated the first direct modelling of HC bonds between silicon substrates in a single and 4 ribbon suspension. This modelling allows for visualisation of complex bond deformation in a variety of suspension modes and was shown to be an improvement on the work that utilised extrapolation techniques to determine the energy ratio in an HC bond. This work then allowed for the determination of the suspension thermal noise contribution of HC bonds compared to other thermal noise contributions such as surface, bulk and thermoelastic loss. Iterative development of a single ribbon suspension, scaled from the real HC bonded silicon suspension, up to an ET-like mass, demonstrated that the resultant suspension thermal noise was dominated by bond loss. Translating this to suspension thermal noise in a 4 ribbon suspension demonstrated that the suspension thermal noise was significantly degraded by 1.5 orders of magnitude due to the introduction of HC bonds between the ribbons and ears of the suspension. It was found that the suspension thermal noise of this design of silicon suspension, inclusive of HC bonds to joint the suspension together, exceeded the design sensitivity target. If one was to remove the bond losses, the suspension thermal noise was reduced by a factor of 15 and found to be below the target, however the design target includes the entire thermal noise budget of the ET-LF detector, whereas the work in this thesis only considers suspension thermal noise arising from the test mass mirror and final stage suspension elements. Furthermore, this work includes the best case mechanical loss values for silicon, where these values may be affected by any surface treatments deemed necessary for strength improvements in future. The introduction of HC bonds between the ears and masses did not significantly increase the thermal noise, compared to

the ribbon to ear bonds. Most of this additional thermal noise was attributed to the bounce mode and corresponded to a $\sqrt{2}$ increase. During the course of this work, it was noted that one has to consider higher stresses pushing the bounce mode further down into the frequency band where these cryogenic low frequency detectors are meant to be most sensitive. The high frequency part of these next generation detectors are designed to have superior sensitivity >30 Hz. This has not been a previous consideration for broadband detectors utilising fused silica suspensions, where the aim was always to push the bounce mode frequency as low as possible out of the detection frequency of interest. By increasing the stress in a 4 ribbon suspension, one can reduce the suspension thermal noise. By doubling the 40 MPa suspension stress to 80 MPa, one can reduce the thermal noise by 22 % at 10 Hz. Increasing the stress of silicon suspensions may also require surface treatments similar to those discussed above.

This work highlights the fact that 3G detectors, using current manufacturing and bonding techniques, and with available surface strength treatments, do not currently meet design sensitivity requirements for suspension thermal noise. Therefore, to obtain the desired performance, it is clear the key points of increasing the strength of silicon while maintaining low mechanical loss, and developing low loss jointing methods, are the technical challenges the field needs to overcome.

Strength in the fibres will permit improved thermal noise as demonstrated, and jointing techniques in the construction of silicon suspensions remains a crucial ongoing research area, building on the work demonstrated in this thesis. This constructional challenge may ultimately define the materials, geometries and construction methodology used in these 3G instruments.

Appendix A

Future chemical etching work

This appendix describes work carried out by the author to further investigate KOH etching to strengthen silicon, as discussed in chapter 3, along with future experimental proposals.

Advice and recommendations were sought from Dr Phillip Dobson, School of Engineering, University of Glasgow and Mr James McIver, School of Chemistry, University of Glasgow, both of whom have extensive experience with KOH etching. A commonly used KOH etch was advised by both experts of 3 parts of 40 % KOH concentration to 1 part IPA at 55 °C. The etch rate for this recipe should be $0.3\,\mu m\,min^{-1}$ [313] [410]. A concentration of 40 % KOH should minimise hillocks that may create stress concentrators on the vertices of the samples.

Further improvements to the experiment would include Piranha cleaning of the samples, rather than standard solvent cleaning prior to etching. A standard Piranha solution of 7 parts sulphuric acid (H_2SO_4) to 1 part hydrogen peroxide (H_2O_2) would be used. Piranha cleaning will provide a more even etch rate on all surfaces and minimise further contamination of the etch solution [313].

Furthermore, sufficient agitation of the etchant solution is important to ensure a flow of fresh etchant solution over the surface of the samples, minimise any areas of uneven etching and move hydrogen gas away from the etching surfaces. This etch should also have a suitable HF pre-dip to remove any native or thermal oxide on the laser cut silicon samples.

The experiment would be split into two. The first experiment would be in the form of a test etch to investigate the etch rate using two etching baths in parallel. This initial etch test is required in order to understand how samples etch, due to the fact there may be an initial layer of polycrystalline silicon on the laser cut edges, combined with the potential for unusual etch rates where the (100) silicon plane polished faces and laser cut edges meet.

This initial etch would consist of 3 batches. One batch would consist of random crystal orientated samples and two batches would consist of crystal orientated samples.

All 3 batches would be Piranha cleaned followed by a pre-dip in a 2 % HF dip solution, consisting of a ratio of 20 ml of HF at 49 % concentration mixed with 480 ml of DI water. These samples would be left in for a couple of hours with occasional agitation to ensure all oxide was

removed. The samples would then be placed into a succession of DI water baths to rinse any residual HF away, both for safety and to minimise any HF contamination in the KOH etch.

One random and one crystal orientated batch of samples would then be placed in the first bath and etched in the recipe for 2 hours, which should yield an etch amount of 36 µm if etched at the theoretical etching rate and would be measurable on the thin fibre profiler in both dimensions. This bath would have some form of condenser to ensure IPA did not boil off during the etch. The second bath would be identical to the first and contain the second batch of crystal orientated sample etched in the same recipe, for the same time, except this apparatus would contain no condenser.

After etching, the samples would again be placed into a succession of DI water baths to rinse any residual KOH away both for safety and to cease any etching processes that may remain. pH testing of the water would confirm the presence of any residual KOH.

On top of determining the etch rate, this experiment would allow for the effect of the IPA boiling off to be investigated by comparing both sets of crystal orientated samples in both surface characterisation post-etch and strength. Furthermore, one can investigate any potential effects the etch has on randomly orientated samples by comparing them to the crystal orientated samples in the same bath after etching.

The second experiment would seek to correlate strength to surface damage removal. This etch would involve 5 batches of 20 samples, all 775 µm thick and all crystal orientated.

All 5 batches would undergo the same Piranha cleaning and HF pre-dip as the initial experiment. All 5 batches would then be placed into the first etching bath together complete with a condenser to stop any IPA boiling off during the etch (unless the initial etch test shows this is not necessary). After 2 hours, the first batch of samples would be removed into successive DI water baths and stored in the final water bath to ensure any etching processes has stopped. Every 2 hours after this, another batch of samples would be removed and this process repeated until the final batch, after 10 hours of etching, was removed. These samples would then be removed from the final water bath ready for characterisation and strength testing.

One advantage of this experiment is that it would allow one to compare the first batch of samples removed after 2 hours to the crystal orientated samples etched in the first bath to ensure consistency in etch rate and surface quality between the two experiments. It also ensures all samples undergo the exact same etching process since they are all contained within the same bath and hence the same etchant solution.

Most importantly however, one could correlate surface damage removal to tensile strength. By etching 5 batches of samples by 36 µm, 72 µm, 108 µm, 144 µm and 180 µm respectively one could plot the strengths of each batch against surface damage removal. If a plateau is observed, one could then deduce that surface damage removal beyond this amount does not aid in increasing strength further. This would allow for quantitative data of surface damage depth on laser cut samples, something which is not currently well understood or published in

literature [239]. The etch times also allow one to probe past the 150 µm surface removal standard applied in previous surface treatments to discover if this material removal amount was sufficient or excessive. It should be noted, this experiment does work on the assumption that the surface finish from the KOH etch itself does not leave further surface damage from etching process itself.

A.0.0.1 Experimental design

Due to the COVID-19 pandemic, ITM could no longer support further KOH work. A substantial amount of work went into developing in-house chemical etching baths and techniques. Unfortunately, this work was not completed and the etches were never carried out due to the COVID-19 pandemic disruption and similar issues faced by ITM. Work that was carried out is described below.

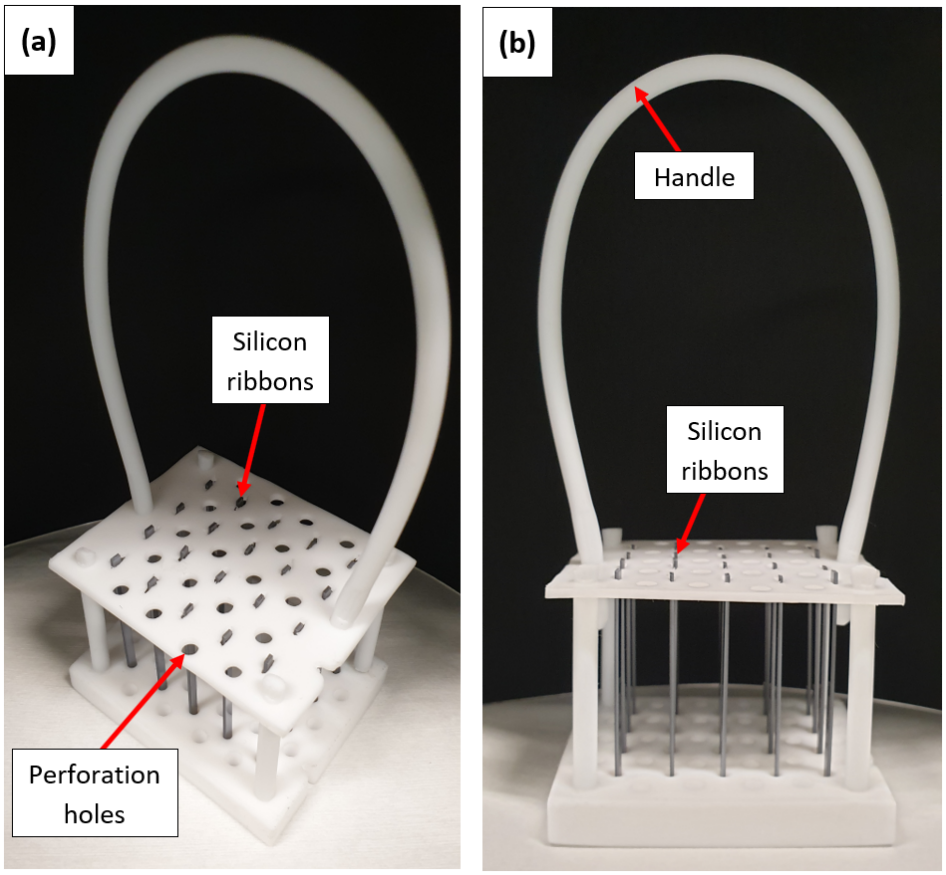

Figure A.1: Example of a PTFE KOH jig designed to hold a batch of 20 silicon samples for KOH etching.

For the initial etch, all 60 samples were individually profiled on the thin fibre profiler to measure the sample width and thickness before etching. A subset of 5 samples from each batch underwent SEM imaging on each face and edge. Surface profiling was also conducted on the polished faces.

For the second etch, all 100 samples were individually profiled on the thin fibre profiler to measure the sample width and thickness before etching. A subset of 5 samples from each batch underwent SEM imaging on each face and edge. Surface profiling was also conducted on the polished faces.

In order to hold the samples in batches and track their orientation, an etching jig was designed, as shown in figure A.1. This jig had to be made from KOH compatible materials, so PTFE was selected. KOH etches metals, so no screws or metal components could be used to build or secure the jig sections. The jig was designed to be press-fitted together to avoid the use of incompatible adhesives. A handle was incorporated for safety to allow the jig to be hooked from the Piranha, HF pre-dip and KOH baths.

As can be seen in figure A.1(a), perforation holes were drilled through the horizontal plate at the top. This was to allow sufficient flow in and around all samples, to allow hydrogen gas to escape to the surface and to minimise any Piranha, HF or KOH solution that could be trapped on surfaces when the jigs were removed from the respective baths. At the bottom of each sample slot, small holes were also drilled to allow further drainage of these solutions; these are not seen in figure A.1 but can be seen in figure A.3(b) below.

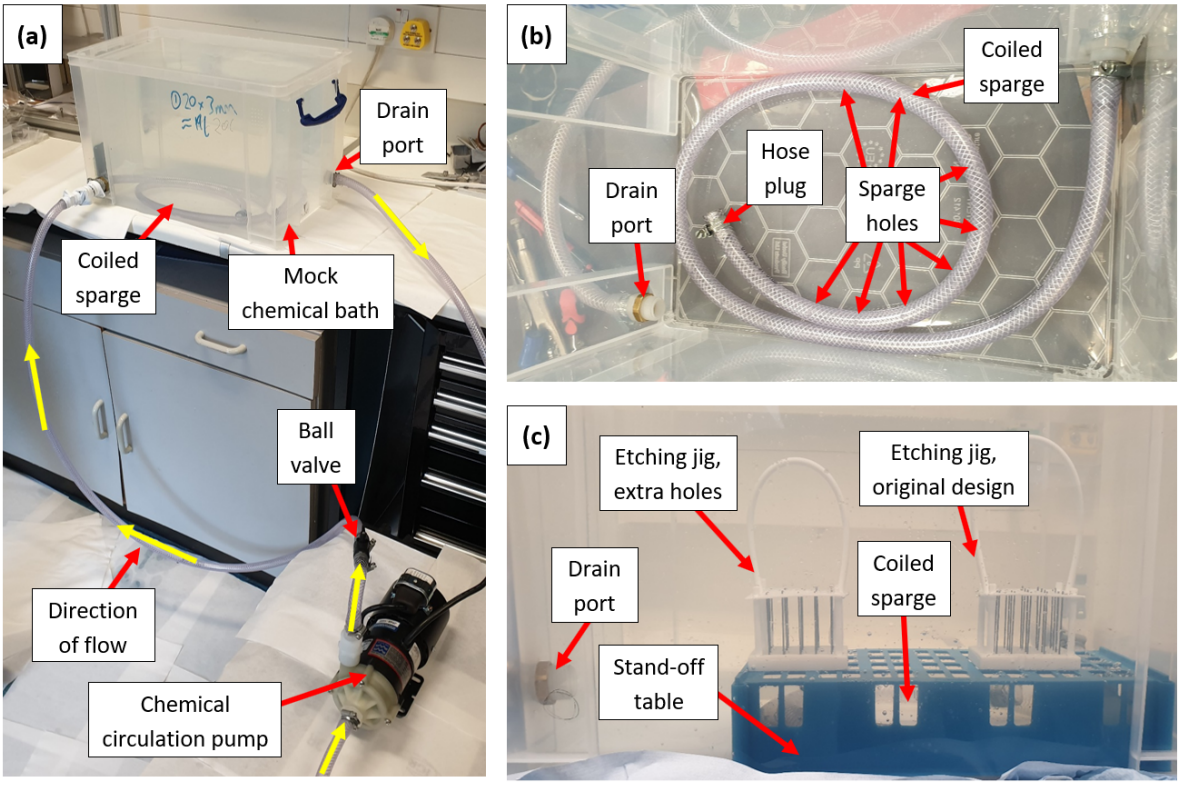

Figure A.2: (a) Mock up of the KOH chemical bath and external circulation system. (b) Inside of the chemical bath detailing the sparge system. (c) Two etching jigs with samples sitting on a stand-off table above the sparge, submersed in water.

For the chemical bath design, it was decided a fluid sparge would be employed. This effec-

tively operates like placing a shower head under the water of a bath facing upwards, providing a gentle upward flow of solution fed from the bath itself and re-circulated via a KOH compatible fluid pump. The KOH fluid pump was a March May AC-3P-MD [411]. A mock chemical bath set-up, using DI water and food colouring, was created in order to test fluid re-circulation rates and general operation, as shown in figure A.2(a). The aim was to provide a satisfactory flow of solution over the samples without being too forceful to blow the samples out of the etching jigs. Since only DI water and food colouring was used for re-circulation testing, normal plastics and metals were used. The sparge was made by blocking one end of the hose with a metal plug, coiling the hose up and drilling twenty 3 mm holes into the tubing, as shown in figure A.2(b). This hose was fed by the output of the chemical pump. A drain port was installed in the wall of the bath to supply the chemical pump with water from the bath. A ball valve was inserted onto the output line from the chemical pump for rough control of the fluid flow rate. A plastic stand-off table was used so that the etching jigs could sit over the sparge without blocking any of the sparge holes, as shown in figure A.2(c).

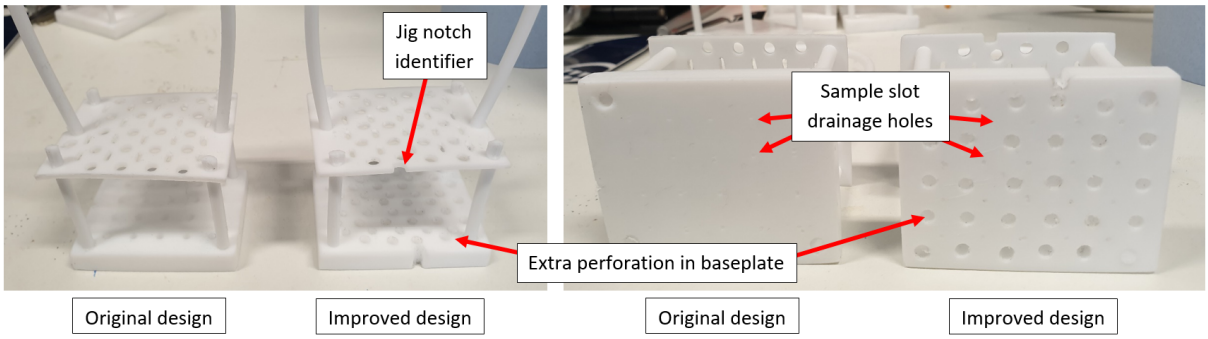

Figure A.3: Images showing the improved PTFE etching jig with extra perforation on the base plate compared to the original jig design.

It was suspected the original jig design, shown in figure A.1, may suffer from poor flow rate over the samples due to a lack of perforation holes in the bottom horizontal plate. A second jig had an extra set of perforation holes drilled in its base plate of the jig, as shown in figure A.3, to improve flow over the samples.

With the pump running, food colouring could then be injected via a syringe directly into the drain port. This flowed through the pump and sparge set-up, being expelled through the sparge to mimic an etching solution, as shown in figure A.4.

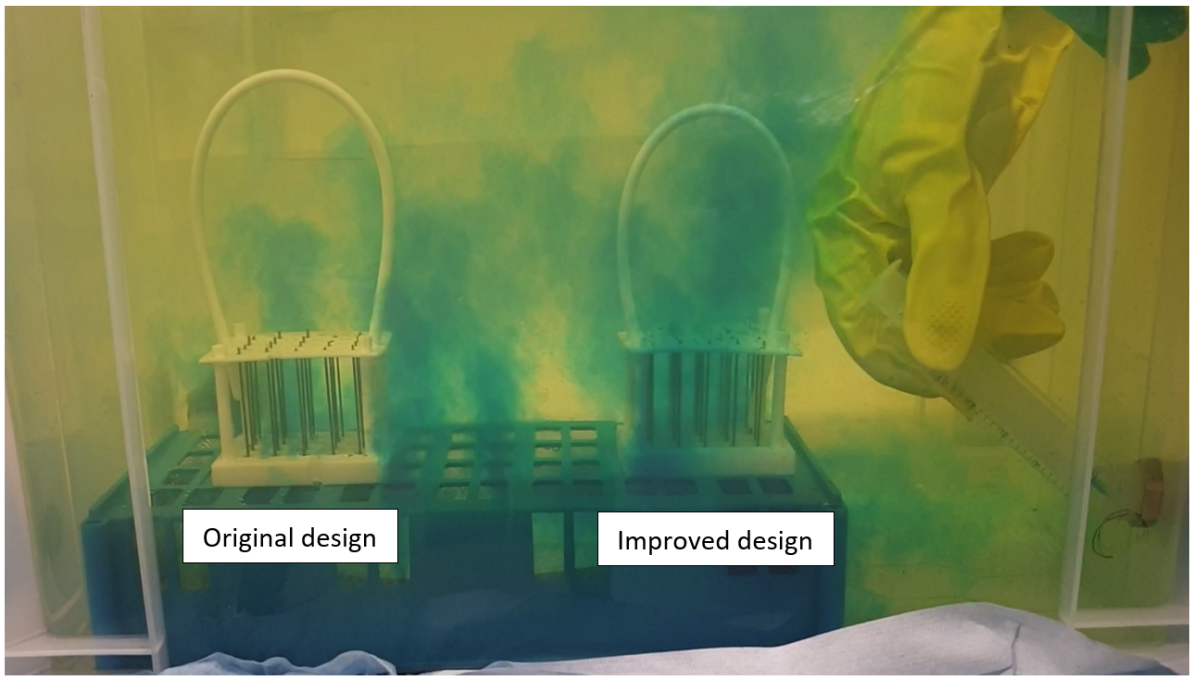

Figure A.4: Food colouring mock test visualising flow over samples comparing the two etching jig designs.

It can be seen in figure A.4, that flow over the right hand jig with extra perforation is superior, with observable fluid flow over all samples. The original jig design shows the base blocking the majority of fluid expelled from the sparge underneath meaning the interior silicon samples would likely etch at a slower or uneven rate compared to the exterior samples which will experience a fresh flow of solution over their surfaces. To hold 100 samples, 5 jigs were made with each having a specific notch identifier in it in order to identify the batch contained within the jig, as shown in figure A.3(b).

A.0.1 Agitation set-up

This section covers the agitation set-up for the real chemical bath. Figure A.5(a) shows the basic working principle of the chemical etch bath itself, which consists of a polypropylene (PP) bucket with a coiled sparge made from Norprene. The ball valve is made from PP with EPDM seals. Metal jubilee clips are used to secure any external tubing as they do not come into contact with the KOH. For internal tube and fixings PTFE securing straps are used, as shown in figure A.5(b). Figure A.5(c) shows the sparge placed at the bottom of the inner PP bucket with the output flow line from the KOH pump coming in vertically from the top. A perforated mesh has been cut and placed over the sparge to act as a standoff table. The final design incorporated an external water bath in order to modulate the temperature of the etching process.

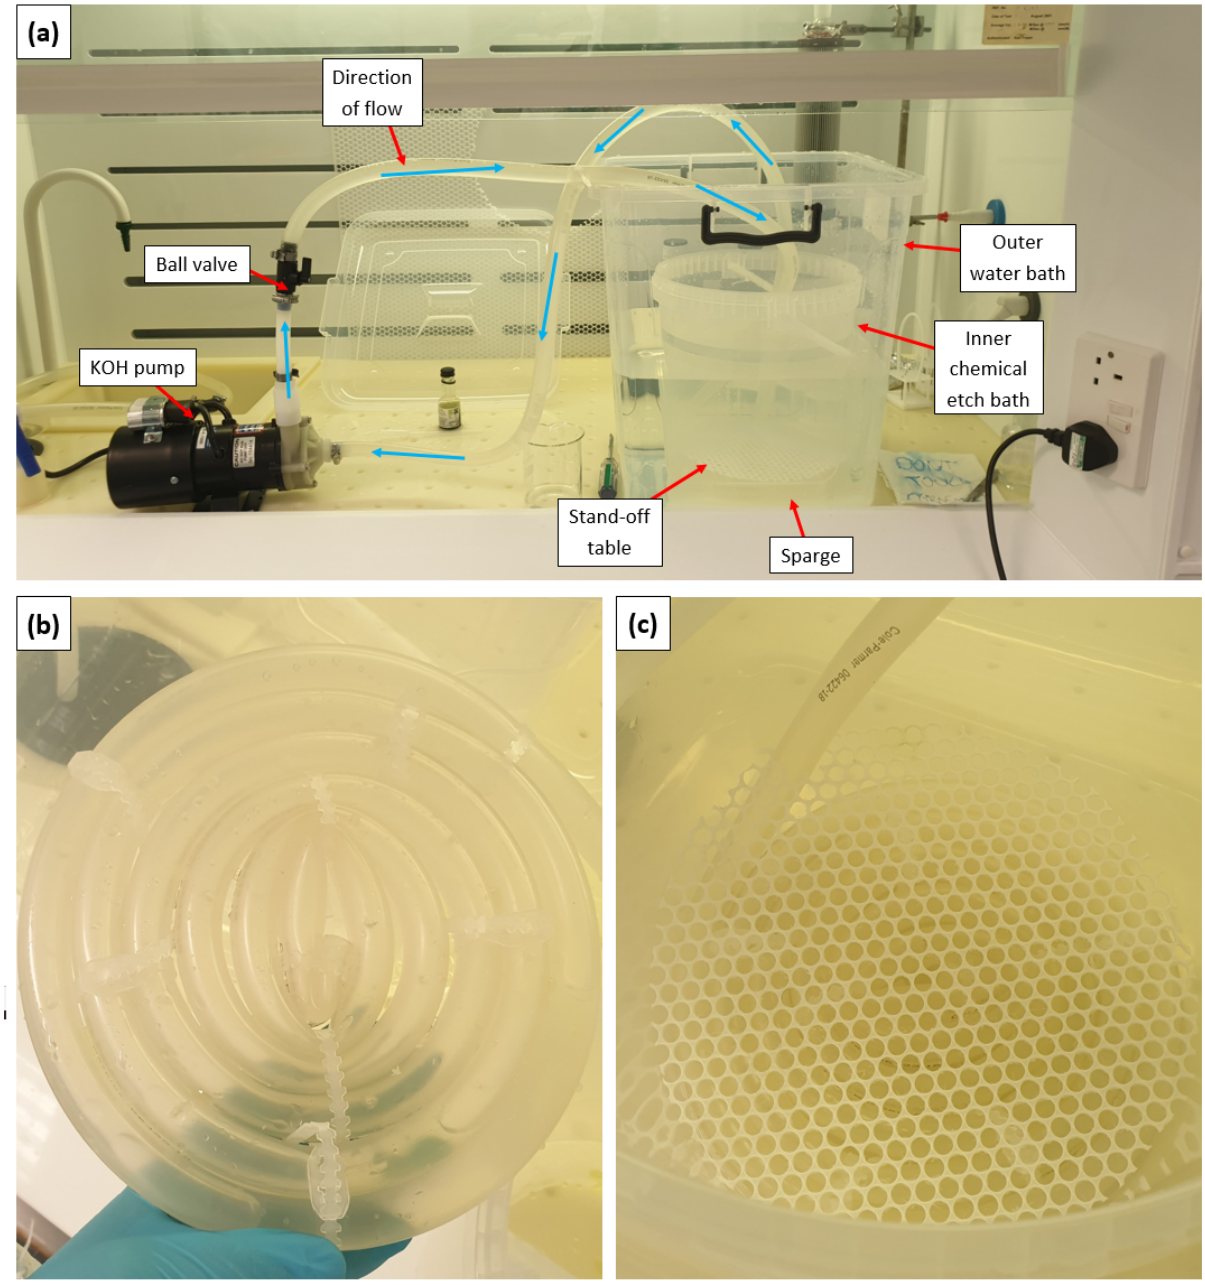

Figure A.5: Basic working principle of in-house chemical etching bath. (a) Chemical flow part of the set-up. (b) Sparge made from Norprene tubing. (c) Sparge placed at the bottom of the etching bath with a stand-off table on top.

A.0.2 Complete set-up

This section covers the remaining set-up for the real chemical bath, as shown in figure A.6. This set-up includes a submersible heater, suitable for heating the water bath to up to 80 °C. A mesh stand-off table is used within the water bath to stand the chemical etching bath off the heater element below. A heater temperature probe is used to monitor the water bath temperature, complete with a feedback solenoid to stabilise the system as shown in the inset of figure A.6. A

simple water re-circulation pump is shown on the right hand side of figure A.6. This circulates water in the water bath in order to remove any hot spots and provide even heating to the internal chemical etch bath.

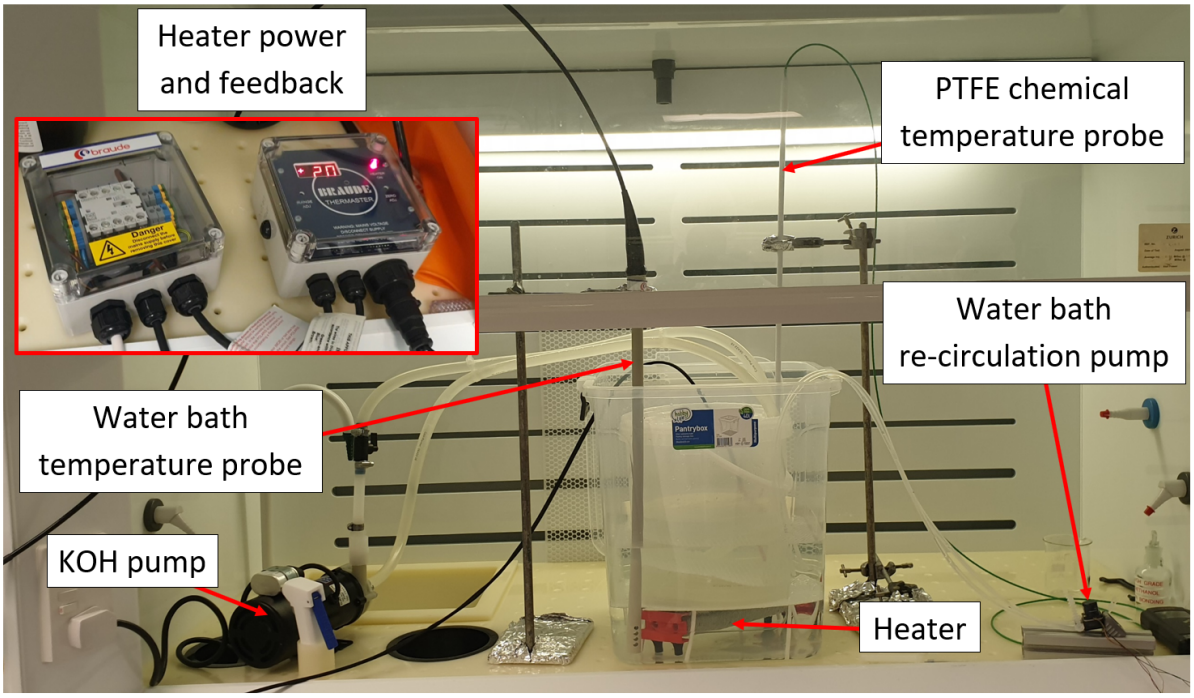

Figure A.6: Finalised set-up of KOH chemical etching bath complete with agitation and heating system. The only item not shown in this set-up is a findenser which will be placed through the centre of the internal chemical etch bath lid.

The internal chemical etch bath has a sealed lid, to minmise boil-off of any IPA or etchant solution. A PTFE chemical temperature probe is passed through this sealed lid in order to monitor the temperature in the bath. Any penetration of the chemical etching bath lid is to be sealed with PTFE tape around the openings to minimise evaporation.

Not shown in the set-up was a findenser, or air condenser. The findenser is placed through the centre of the internal chemical bath lid. This both allows for any gases from the etch process to escape and condenses any etchant solution or IPA, allowing it to drip back into the etch bath.

A.0.3 Future work

This is as far as this chemical etching set-up was taken before work had to be ceased due to the COVID-19 pandemic. For future work, before any etching of samples occurs within chemical baths, one needs to season the baths in order to remove any remaining contaminants and to season all surfaces exposed to the etchant; typically two etching runs with no samples, draining each bath at the end of each run, should be sufficient [313] [410]. This opportunity should also be used to move the PTFE chemical temperature probe to different points in the chemical etching bath to measure the temperature profile of the bath. This would allow for the adjustment

of chemical fluid flow and water bath heating and re-circulation to minimise any hot spots that would etch samples at a different rate.

Appendix B

Thermal conductivity derivations

For the low temperature regime of silicon, one assumes specific heat dominates as shown in figure 2.14 of chapter 2, therefore one can assume:

$$k = \alpha T^3 \tag{B.1}$$

Taking the equation for heat current, P, in thermal equilibrium [369]:

$$P = \frac{dQ}{dt} = kA\frac{dT}{dL} \tag{B.2}$$

where dQ is some quantity of heat energy transferred in time dt, k is the thermal conductivity of the material, A is the cross-sectional area of the path of conduction and dT is the change in temperature measured over distance dL.

Taking the initial specific heat assumption, equation B.2 becomes:

$$P = \alpha T^3 A\frac{dT}{dL} \tag{B.3}$$

which can be rearranged to give:

$$\int_{L1}^{L2} P\,dL = \int_{T1}^{T2} \alpha A T^3\,dT \tag{B.4}$$

This defines a new equation for heat current:

$$PL = \frac{1}{4}A\alpha\left(T_2^4 - T_1^4\right) \tag{B.5}$$

If one applies a small heat step, ΔP, to this system, this equation becomes:

$$(P+\Delta P)L = \frac{1}{4}A\alpha\left[(T_2+\Delta T_2)^4 - (T_1+\Delta T_1)^4\right] \tag{B.6}$$

If one assumes, $\Delta T << T$, this can be simplified, by taking the two highest-order terms from the expansion and assuming negligible ΔT terms, to give:

$$\begin{aligned}(P+\Delta P)L &= \frac{1}{4}A\alpha\left(T_2^4+4T_2^3\Delta T_2-4T_1^3\Delta T_1-T_1^4\right)\\ PL+\Delta PL &= \frac{1}{4}A\alpha\left(T_2^4-T_1^4\right)+\frac{1}{4}A\alpha\left(4T_2^3\Delta T_2-4T_1^3\Delta T_1\right)\\ \Delta PL &= A\alpha\left(T_2^3\Delta T_2-T_1^3\Delta T_1\right)\end{aligned} \tag{B.7}$$

However, this assumption was found to be incorrect since recorded temperature data from suspension experiments showed $\Delta T \approx T$ and therefore the full expansion of equation B.6 becomes:

$$\begin{aligned}(P+\Delta P)L = \frac{1}{4}A\alpha\Big[&\left(T_2^4+4T_2^3\Delta T_2+6T_2^2\Delta T_2^2+4T_2\Delta T_2^3+\Delta T_2^4\right)\\ &-\left(T_1^4+4T_1^3\Delta T_1+6T_1^2\Delta T_1^2+4T_1\Delta T_1^3+\Delta T_1^4\right)\Big]\end{aligned} \tag{B.8}$$

This simplifies to:

$$\begin{aligned}\Delta PL = A\alpha\Big[&\left(T_2^3\Delta T_2+\frac{3}{2}T_2^2\Delta T_2^2+T_2\Delta T_2^3+\frac{1}{4}\Delta T_2^4\right)\\ &-\left(T_1^3\Delta T_1+\frac{3}{2}T_1^2\Delta T_1^2+T_1\Delta T_1^3+\frac{1}{4}\Delta T_1^4\right)\Big]\end{aligned} \tag{B.9}$$

For the analysis conducted in section 4.5 of chapter 4, T_1 and T_2 in equation B.9 are fixed temperatures, taken from the top and bottom ribbon sensors at the cold thermal equilibrium (also referred to as "ambient").

Error analysis With the errors on T_1 and T_2, defined as δT_1 and δT_2, taken as the absolute temperature sensor error, one must determine the ΔT_1 and ΔT_2 errors, $\delta\Delta T_1$ and $\delta\Delta T_2$. Since these absolute temperature sensor errors are primarily systematic errors, one must therefore subtract the absolute temperature error of the sensor at cold thermal equilibrium, $\delta T_{ambient}$, from the absolute temperature error of the same sensor at some elevated heat step n, $\delta T_{heat\,step\,n}$, to give:

$$\delta\Delta T_n = \delta T_{heat\,step\,n} - \delta T_{ambient} \tag{B.10}$$

where $\delta T_{heat\,step\,n}$ is the stabilised temperature error of the top or bottom sensor at heat step n and $\delta T_{ambient}$ is the corresponding top or bottom sensor stabilised temperature error with no heating applied. T_2 is always the hotter sensor and T_1 is always the colder sensor, so in all cases these are the bottom and top ribbon sensors respectively.

To determine the error on the thermal conductivity, k, one must determine the error on α in equation B.6.

Firstly, re-arranging, simplifying and assuming P is some unknown but constant ambient heat load on the system, equation B.6 gives:

$$\alpha = \frac{4\Delta PL}{AT_{term}} \tag{B.11}$$

where:

$$T_{term} = (T_2 + \Delta T_2)^4 - (T_1 + \Delta T_1)^4 \tag{B.12}$$

One must now split the variables into correlated and uncorrelated values. For each step, L and A are correlated between heat steps used to calculate α. P is uncorrelated between heat steps. There is however a subtlety in the T_{term}. Individually, the values of $T_2 + \Delta T_2$ and $T_1 + \Delta T_1$ are uncorrelated, however when combined and applied to equation B.11, the values become correlated due to the underlying physics of the system, where the applied power at the bottom of the mass will raise or lower both ΔT_2 and ΔT_1 by some correlated amount. This is addressed in the errors below.

To approach the combination of uncorrelated and correlated errors analysis one starts by defining each variable's error:

$$\begin{aligned} \delta\alpha_1 &= \frac{4\Delta P}{AT_{term}}\delta L \\ \delta\alpha_2 &= \frac{4L}{AT_{term}}\delta\Delta P \\ \delta\alpha_3 &= \frac{4\Delta PL}{T_{term}}\delta\left(\frac{1}{A}\right) = \frac{-4\Delta PL}{T_{term}}\frac{1}{A^2}\delta A \end{aligned} \tag{B.13}$$

where δL is the error in the measured length as defined in section 4.5.2.2 and $\delta\Delta P$ is the error in the measured heater power as defined in section 4.5.2.1 of chapter 4.

δA is the error in the measured area using the errors defined in section 4.5.2.2 of chapter 4, and using the following equation:

$$\delta A = \sqrt{t^2(\delta w)^2 + w^2(\delta t^2)} \tag{B.14}$$

where t is the measured thickness of the ribbon with measurement error, δt and w is the measured width of the ribbon with measurement error, δw.

Since the δT_{term} consists of two uncorrelated values, $T_2 + \Delta T_2$ and $T_1 + \Delta T_1$, one can write:

$$\begin{aligned} \delta x &= 4(T_2 + \Delta T_2)^3\delta(T_2 + \Delta T_2) \\ \delta y &= 4(T_1 + \Delta T_1)^3\delta(T_1 + \Delta T_1) \end{aligned} \tag{B.15}$$

where $\delta(T_n + \Delta T_n)$ is the absolute temperature error at the elevated $T_n + \Delta T_n$ temperature interpolated from tables 4.1 and 4.2.

One can then define:

$$\delta\alpha_4 = \frac{4\Delta PL}{A}\delta\left(\frac{1}{T_{term}}\right) = \frac{-4\Delta PL}{A}\frac{1}{T_{term}^2}\delta T_{term} \tag{B.16}$$

where:

$$\delta T_{term} = \sqrt{(\delta x)^2 + (\delta y)^2} \tag{B.17}$$

Finally, one can now gather correlated and uncorrelated errors together, noting the underlying correlation of the $\delta\alpha_4$ arising from the δT_{term} term, into terms $\delta\beta$ and $\delta\gamma$ respectively.

Correlated errors:

$$\delta\beta = \sqrt{(\delta\alpha_1)^2 + (\delta\alpha_3)^2 + (\delta\alpha_4)^2} \tag{B.18}$$

Uncorrelated errors:

$$\delta\gamma = \sqrt{(\delta\alpha_2)^2} \tag{B.19}$$

One can then combine the correlated, $\delta\beta$, and uncorrelated, $\delta\gamma$, errors to attain the final error on α, $\delta\alpha$ from:

$$\delta\alpha = \pm\delta\beta \pm \sqrt{(\delta\gamma)^2} \tag{B.20}$$

By calculating and taking the largest upper and lower errors for each calculation of α, which should be close to a constant across any one suspension (or indeed between any two suspension experiencing the same cryogenic environment), one then provides the most conservative error estimate on the mean α value used to calculate the thermal conductivity, k, at temperature, T.

Appendix C

Cryogenic suspension temperature data

This appendix lists the temperature data for suspension experiments 8, 9 and 10, taken for each stabilised heat step. Data in red highlights a temperature step where the average temperature of the silicon ribbon between the top and bottom ribbon sensors exceeds 25 K, where the specific heat-dominated $k = \alpha T^3$ relationship is believed to break down.

This appendix also lists the absolute error on an uncalibrated DT-670A SD sensor.

C.1 Suspension 8 temperature data

Heat step (Lake Shore power applied)	Top fuse end (K)	Top ribbon (K)	Bottom ribbon (K)	Bottom fuse end (K)	Bottom of mass (K)	Bottom of suspension structure (K)	IVC wall (K)
Ambient (0 %)	10.159	13.767	15.370	16.599	16.658	13.630	6.292
Heat step 1 (5.12 %)	10.251	14.399	16.190	17.828	17.914	13.567	6.282
Heat step 2 (10.24 %)	10.658	16.228	18.328	20.980	21.137	13.609	6.271
Heat step 3 (8.85 %)	10.418	15.556	17.598	19.943	20.077	13.440	6.242
Heat step 4 (15.36 %)	11.241	18.785	21.069	25.098	25.330	13.604	6.239
Heat step 5 (14.71 %)	11.127	18.406	20.671	24.507	24.736	13.560	6.234
Heat step 6 (12.56 %)	Heat step never stabilised due to compressor failure						

Table C.1: Suspension 8 stabilised heat steps with corresponding cryogenic temperature sensor readings.

C.2 Suspension 9 temperature data

Heat step (Lake Shore power applied)	Top fuse end (K)	Top ribbon (K)	Bottom ribbon (K)	Bottom fuse end (K)	Bottom of mass (K)	Bottom of suspension structure (K)	IVC wall (K)
Ambient (0 %)	6.232	8.313	10.022	10.452	10.465	7.626	11.896
Heat step 1 (15.36 %)	8.314	16.637	19.363	22.458	22.679	8.171	11.879
Heat step 2 (14.91 %)	8.212	16.329	19.068	22.065	22.292	8.129	11.890
Heat step 3 (12.19 %)	7.703	14.427	17.211	19.625	19.797	8.107	12.066
Heat step 4 (7.35 %)	6.745	11.306	13.985	15.443	15.530	7.713	11.861
Heat step 5 (3.68 %)	6.318	9.295	11.536	12.348	12.386	7.584	11.800
Heat step 6 (30.72 %)	12.143	28.454	31.047	37.987	38.478	9.563	11.890
Heat step 7 (1.84 %)	6.198	8.550	10.446	10.989	11.007	7.526	11.769
Heat step 8 (60.00 %)	19.012	52.087	55.420	71.056	72.024	13.794	12.858

Table C.2: Suspension 9 stabilised heat steps with corresponding cryogenic temperature sensor readings.

C.3 Suspension 10 temperature data

Heat step (Lake Shore power applied)	Top fuse end (K)	Top ribbon (K)	Bottom ribbon (K)	Bottom fuse end (K)	Bottom of mass (K)	Bottom of suspension structure (K)	IVC wall (K)
Ambient (0 %)	5.661	8.492	9.369	9.848	9.861	5.768	6.245
Heat step 1 (1.84 %)	5.639	8.814	9.823	10.422	10.464	5.744	6.225
Heat step 2 (3.68 %)	5.708	9.639	10.882	11.681	11.729	5.818	6.267
Heat step 3 (7.35 %)	5.830	11.483	13.149	14.198	14.273	5.930	6.296
Heat step 4 (12.19 %)	6.034	13.513	15.683	16.903	17.019	6.115	6.306
Heat step 5 (14.91 %)	6.200	14.490	16.931	18.226	18.350	6.267	6.326
Heat step 6 (15.36 %)	6.218	14.641	17.128	18.435	18.565	6.280	6.319
Heat step 7 (30.72 %)	7.405	19.009	22.790	24.612	24.839	7.361	6.325
Heat step 8 (60.00 %)	10.508	25.241	32.618	34.479	34.881	10.280	6.543
Heat step 9 (19.06 %)	6.473	15.828	18.666	20.069	20.219	6.511	6.341
Heat step 10 (50 %) + 100 W plate heaters	118.021	118.843	124.054	121.201	121.494	118.086	118.988

Table C.3: Suspension 10 stabilised heat steps with corresponding cryogenic temperature sensor readings.

C.4 Uncalibrated DT-670A SD sensor error

Temperature (K)	Typical sensor accuracy (mK)	Typical sensor sensitivity ($mV\,K^{-1}$) [347]	Instrument measurement resolution (μV) [348]
4.2	250	-31.6	10.0
10	250	-26.8	
20	250	-15.6	
30	250	-2.0	

Table C.4: Sensor accuracy and sensitivity for an uncalibrated DT-670A silicon diode sensor and Lake Shore 340 temperature controller measurement resolution.

Temperature (K)	Measurement resolution converted (mK)	Electronic accuracy (V) [343]	Electronic accuracy converted (mK)	Absolute temperature error (mK) [345]
4.2	0.3	1.6	5.0	255.0
10	0.4	1.4	5.6	255.6
20	0.6	1.2	9.0	258.9
30	5.1	1.1	69.0	319.0

Table C.5: Electronic accuracy of Lake Shore 340 temperature controller showing absolute combined temperature error calculated from [345].

Appendix D

Silicon thermal conductivity data from Touloukian

Figure D.1 shows the thermal conductivity recorded for a variety of silicon samples. The details of the silicon sample for each curve can be found on pages 327-332, with the corresponding numerical data found on pages 333-338, and Touloukian's recommended curve with numerical data on page 339 of [376].

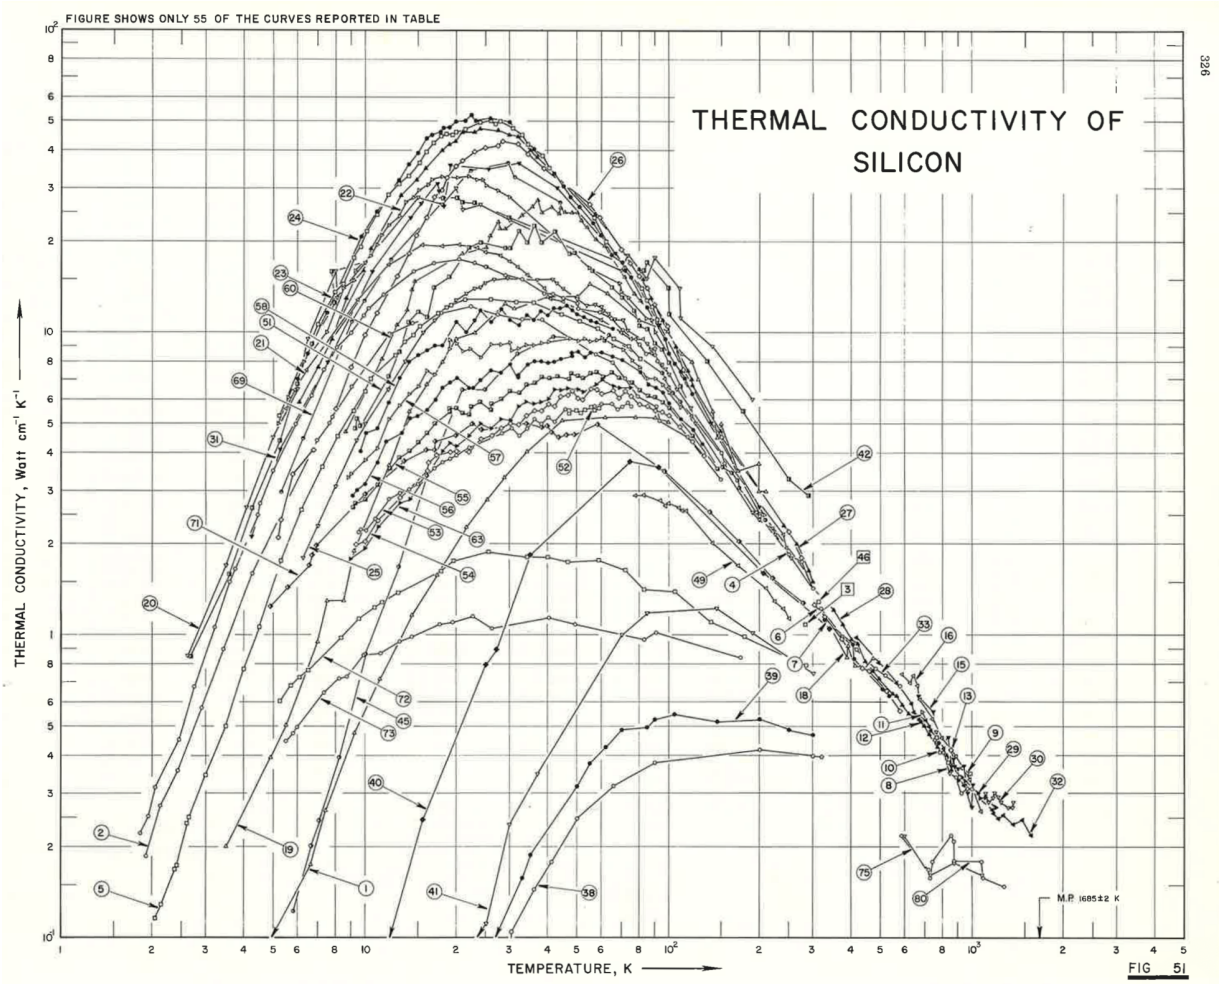

Figure D.1: Thermal conductivity of silicon samples with corresponding curves detailed in [376].

Appendix E

ANSYS APDL script to build 4 ribbon model with HC bonds on one ribbon

This appendix references ANSYS APDL code written by the author that was used and modified to obtain the results in chapter 5. This APDL code consists of 3 files and produces 1 file;

- the main file which builds the 4-ribbon, ET-like model with 361 nm bonds between ear 1, ear 2 and ribbon 1
- the earelemcount.txt file which counts the number of elements in ear 1 and ear 2
- the ribbonelemcount.txt file which counts the number of elements in ribbon 1
- the parameters.txt file which is generated from the main file to produce a series of useful parameters specific to the test mass geometry.

The code has been uploaded to the LIGO Document Control Centre, DDC number LIGO-T2200329.

Appendix F

MATLAB suspension thermal noise script

This appendix references MATLAB code written by the author that was used and modified to obtain the results in chapter 5. This MATLAB code is used to calculate suspension thermal noise. The main code calls on a series of functions described below. It should be noted this code requires silicon experimental data and material property files in order to run.

- ThermalNoiseFDT.m - the main MATLAB file that calls on each function
- interpolatedata.m - this function imports and interpolates variable data over the temperature range of interest
- surfacel.m - this function calculates the surface loss of the ribbon
- bulkl.m - this function imports the bulk loss data
- bondl.m - this function imports the bond loss data
- TEloss.m - this function calculates the thermoelastic loss over the chosen frequency range at the defined temperature
- totalNexbond.m - this function calculates the total suspension loss and dilution for the pendulum, bounce and violin modes, assuming no HC bonds exist within the suspension, i.e. it is fully monolithic
- totalNinbond.m - this function calculates the total suspension loss and dilution for the pendulum, bounce and violin modes, with HC bonds included. This bond energy data is produced from the strain energy ratios in the bonds of the 4-ribbon suspensions from chapter 5
- FDT.m - this function takes the losses excluding and including bond loss. The Fluctuation-dissipation theorem is then used to calculate the corresponding displacement noise for each mode, with and without bond loss, to show bond loss contribution to suspension

thermal displacement noise. One notes the violin mode here is included 4 times, to represent each violin mode contribution, however due to the methodology used to determine violin mode energy distribution (see section 5.7.3 of chapter 5), instead of 4 distinct violin mode frequencies, they are combined into one frequency.

The code has been uploaded to the LIGO Document Control Centre, DDC number LIGO-T2200330.

Bibliography

[1] A. Einstein. Die grundlage der allgemeinen relativitätstheorie. In *Annalen der Physik, 49 (7), pp 769-822*, 1916.

[2] A. Einstein. The foundation of the general theory of relativity (translated). In *The Collected Papers of Albert Einstien, 6, pp 146-200*. Princeton University Press, 1997.

[3] A. Cho. Gravitational waves, Einstein's ripples in spacetime, spotted for first time. Online, accessed 29th June 2022. https://www.science.org/content/article/gravitational-waves-einstein-s-ripples-spacetime-spotted-first-time.

[4] T. Creighton. Gravitational waves demystified. Online, accessed 30th June 2022. http://www.tapir.caltech.edu/~teviet/Waves/index.html.

[5] M. Pitkin, S. Reid, S. Rowan, and J. Hough. Gravitational wave detection by interferometry (ground and space). *Living Reviews in Relativity, 14, 5*, 2011.

[6] LIGO Scientific Collaboration. Sources and types of gravitational waves. Online, accessed 25th August 2022. https://www.ligo.caltech.edu/page/gw-sources.

[7] K. T. McDonald. What is the stiffness of spacetime? Online, accessed 30th June 2022. https://physics.princeton.edu/~mcdonald/examples/stiffness.pdf.

[8] B. S. Sathyaprakash and B. F. Schutz. Physics, astrophysics and cosmology with gravitational waves. *Living Reviews in Relativity, 12, 2*, 2009.

[9] B. F. Schutz. Determining the Hubble constantfrom gravitational wave observations. *Nature, 323, pp 310-311*, 1986.

[10] T. J. Massinger. *Detector characterization for Advanced LIGO*. PhD thesis, Syracuse University, 2016.

[11] LIGO Scientific Collaboration. The sound of two black holes colliding. Online, accessed 25th August 2022. https://www.ligo.caltech.edu/video/ligo20160211v2#:~:text=As%20the%20black%20holes%20spiral,sound%20like%20a%20bird's%20chirp.

[12] M. Cabero et al. Observational tests of the black hole area increase law. *Physical Review D, 97, 124069*, 2018.

[13] C. Knox. Artistic image inspired by a BHNS merger. Online, accessed 30th June 2022. https://gravitational-waves.phas.ubc.ca/files/2021/11/Carl_Knox_graphic-768x229.png.

[14] J. Hester and A. Loll. Core-collapse supernovae: fascinating cosmic fireworks. Online, accessed 30th June 2022. https://compstar.uni-frankfurt.de/outreach/short-articles/core-collapse-supernovae-fascinating-cosmic-fireworks/.

[15] F. Mirabel. Artist's view of a black hole and companion star GRO J1655-40. Online, accessed 5th July 2022. https://hubblesite.org/contents/media/images/2002/30/1273-Image.html.

[16] C. J. Moore, R. H. Cole, and C. P. L. Berry. Gravitational-wave sensitivity curves. *Classical and Quantum Gravity, 32 (1), 015014*, 2014.

[17] N. Anderson et al. The transient gravitational-wave sky. *Classical and Quantum Gravity, 30 (19), 193002*, 2013.

[18] S. J. Bell et al. Observation of a rapidly pulsating radio source. *Nature, 217, pp 709-713*, 1968.

[19] C. H. Agar et al. A broad-band radio study of PSR J0250+5854: The slowest spinning radio pulsar known. *Monthly Notices of the Royal Astronomical Society, 508, pp 1102-1114*, 2021.

[20] J. W. T. Hessels et al. A radio pulsar spinning at 716 Hz. *Science, 311 (5769), pp 1901-1904*, 2006.

[21] V. Kalogera and G. Baym. The maximum mass of a neutron star. *The Astrophysical Journal, 470, L61-L64*, 1996.

[22] B. P. Abbott et al. Searches for periodic gravitational waves from unknown isolated sources and Scorpius X-1: Results from the second LIGO science run. *Physical Review D, 76 (8), 082001*, 2007.

[23] B. P. Abbott et al. Upper limits on gravitational wave emission from 78 radio pulsars. *Physical Review D, 76, 042001*, 2007.

[24] D. Payne and A. Melatos. Frequency spectrum of gravitational radiation from global hydromagnetic oscillations of a magnetically confined mountain on an accreting neutron star. *The Astrophysical Journal, 641, pp 471-478*, 2006.

[25] R. Barnard et al. New XMM-Newton analysis of three bright x-ray sources in M31 globular clusters, including a new black hole candidate. *Astrophysical Journal, 689, pp 1215-1221*, 2008.

[26] R. V. Wagoner. Gravitational radiation from accreting neutron stars. *Astrophysical Journal, 278, pp 345-348*, 1984.

[27] M. Sieniawska and M. Bejger. Continuous gravitational waves from neutron stars: current status and prospects. *Universe, 5 (11), 217*, 2019.

[28] B. P. Abbott et al. Upper limits on the stochastic gravitational-wave background from Advanced LIGO's first observing run. *Physical Review Letters, 118, 121101*, 2017.

[29] N. J. Cornish and J. D. Romano. When is a gravitational-wave signal stochastic? *Physical Review D, 92, 042001*, 2015.

[30] P. R. Saulson. *Fundamentals of interferometric gravitational wave detectors*. World Scientific, 2017.

[31] J. M. Weisberg and J. H. Taylor. Relativistic binary pulsar B1913+16: Thirty years of observations and analysis. *ASP Conference Series, 328, pp 25*, 2005.

[32] R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. *Astrophysical Journal, 195, pp L51-L53*, 1975.

[33] J. H. Taylor, L. A. Fowler, and P. M. McCulloch. Measurements of general relativistic effects in the binary pulsar PSR1913+16. *Nature, 277, pp 437-440*, 1979.

[34] The Royal Swedish Academiy of Sciences. Press release of the Nobel prize in physics 1993. Online, accessed 5th July 2022. https://www.nobelprize.org/prizes/physics/1993/press-release/.

[35] J. Weber. Detection and generation of gravitational waves. *Physical Review, 117, pp 306-313*, 1960.

[36] J. Weber. Evidence for discovery of gravitational radiation. *Physical Review Letters, 22 (24), pp 1320-1324*, 1969.

[37] J. Weber. Gravitational-wave detector events. *Physical Review Letters, 20 (23), pp 1307-1308*, 1968.

[38] J. Weber. Observation of the thermal fluctuations of a gravitational wave detector. *Physical Review Letters, 17 (24), pp 1228-1230*, 1966.

[39] J. Weber. Anisotropy and polarization in the gravitational-radiation experiments. *Physical Review Letters, 25 (3), pp 180-184*, 1970.

[40] J. Hough, S. Rowan, and B. S. Sathyaprakash. The search for gravitational waves. *Journal of Physics B: Atomic, Molecular and Optical Physics, 38 (9), S497-S519*, 2005.

[41] R. W. P. Drever, J. Hough, R. Bland, and G. W. Lessnoff. Search for short bursts of gravitational radiation. *Nature, 246, pp 340-344*, 1973.

[42] J. Tyson and R. Giffard. Gravitational-wave astronomy. *Annual Review of Astronomy and Astrophysics, 16, pp 521-554*, 1978.

[43] G. W. Gibbons and S. W. Hawking. Theory of the detection of short bursts of gravitational radiation. *Physical Review D, 4 (8), pp 2191-2197*, 1971.

[44] E. Mauceli et al. The Allegro gravitational wave detector: data acquisition and analysis. *Physical Review D, 54, pp 1264-1275*, 1996.

[45] A. Vinante et al. Present performance and future upgrades of the AURIGA capacitive readout. *Classical and Quantum Gravity, 23 (8), S102-S110*, 2006.

[46] P. Astone et al. Long-term operation of the Rome "Explorer" cryogenic gravitational wave detector. *Physical Review D, 47 (2), pp 362-376*, 1993.

[47] P. Astone et al. The gravitational wave detector NAUTILUS operating at T = 0.1 K. *Astroparticle Physics 7, pp 231-243*, 1997.

[48] I. S. Heng et al. Long term operation of a niobium resonant bar gravitational wave antenna. *Physics Letters A, 218 (6), pp 190-196*, 1996.

[49] A. de Waard et al. MiniGRAIL, the first spherical detector. *Classical and Quantum Gravity, 20 (10), S143-S151*, 2003.

[50] O. D. Aguiar et al. The status of the Brazilian spherical detector. *Classical and Quantum Gravity, 19 (7), pp 1949-1953*, 2002.

[51] K. Riles. Gravitational waves: Sources, detectors and searches. *Progress in Particle and Nuclear Physics, 68, pp 1-54*, 2013.

[52] G. E. Moss, L. R. Miller, and R. L. Forward. Photon-noise-limited laser transducer for gravitational antenna. *Applied Optics, 10 (11), pp 2495-2498*, 1971.

[53] A. A. Michelson and E. W. Morley. On the relative motion of the Earth and the luminiferous ether. *American Journal of Science, 34 (203), pp 333-345*, 1887.

[54] C. Bond, D. Brown, A. Freise, and K. Strain. Interferometer techniques for gravitational-wave detection. *Living Reviews in Relativity, 13, 1*, 2010.

[55] R. Weiss. Electromagnetically coupled broadband gravitational antenna. *Quarterly Progress Report of the MIT Research Laboratory of Electronics, 105, pp 54-76*, 1972.

[56] D. Herriott, H. Kogelnik, and R. Kompfner. Off-axis paths in spherical mirror interferometers. *Applied Optics, 3 (4), pp 523-526*, 1964.

[57] C. Fabry and A. Pérot. Théorie et applications d'une nouvelle méthode de spectroscopie interférentielle. *Annales de chimie et de physique, 16 (7), pp 115-144*, 1899.

[58] R. W. P. Drever et al. A gravity-wave detector using optical cavity sensing. *Proceedings, 9th International Conference on General Relativity and Gravitation, pp 265-267*, 1980.

[59] B. Meers. Recycling in laser-interferometric gravitational-wave detectors. *Physical Review D, 38 (8), pp 2317-2326*, 1988.

[60] K. A. Strain and B. J. Meers. Experimental demonstration of dual recycling for interferometric gravitational-wave detectors. *Physical Review Letters, 66 (11), pp 1391*, 1991.

[61] P. Nguyen et al. Environmental noise in Advanced LIGO detectors. *Classic and Quantum Gravity, 38 (14), 145001*, 2021.

[62] D. M. Macleod et al. Reducing the effect of seismic noise in LIGO searches by targeted veto generation. *Classical and Quantum Gravity, 29 (5), 055006*, 2012.

[63] G. Hammond et al. Advanced technologies for future ground-based, laser-interferometric gravitational wave detectors. *Journal of Modern Optics, 61 (S1), pp S10-S45*, 2014.

[64] N. A. Robertson et al. Quadruple suspension design for Advanced LIGO. *Classical and Quantum Gravity, 19, pp 4043-4058*, 2002.

[65] M. V. Plissi et al. Aspects of the suspension system for GEO 600. *Review of Scientific Instruments, 69 (8) , pp 3055-3061*, 1998.

[66] F. Matichard et al. Seismic isolation of Advanced LIGO : Review of strategy, instrumentation and performance. *Classic and Quantum Gravity, 32, 185003*, 2015.

[67] L. Somlai et al. Seismic noise measures for underground gravitational wave detectors. *Acta Geodaetica et Geophysica, 54, pp 301-313*, 2019.

[68] J. Driggers et al. Subtraction of Newtonian noise using optimized sensor arrays. *Physical Review D, 86, 102001*, 2012.

[69] M. Coughlin et al. Towards a first design of a Newtonian-noise cancellation system for Advanced LIGO. *Classical and Quantum Gravity, 33 (24), 244001*, 2016.

[70] M. G. Beker et al. Improving the sensitivity of future GW observatories in the 1–10 Hz band: Newtonian and seismic noise. *General Relativity and Gravitation, 43, pp 623-656*, 2011.

[71] S. Hild. *A basic introduction to quantum noise and quantum-non-demolition techniques*. Springer, 2014.

[72] B. Picinbono and C. Bendjaballah. Photoelectron shot noise. *Journal of Mathematical Physics, 11 (7), pp 2166-2176*, 1970.

[73] C. M. Caves. Quantum-mechanical noise in an interferometer. *Physical Review D, 23 (8), pp 1693-1708*, 1981.

[74] H. Miao. *Exploring macroscopic quantum mechanics in optomechanical devices*. Springer, 2012.

[75] V. B. Braginsky and F. J. Khalili. Gravitational wave antenna with QND speed meter. *Physics Letters A, 147 (5-6), pp 251-256*, 1990.

[76] 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; and on the general existence of active molecules in organic and inorganic bodies. *Philosophical Magazine and Annals of Philosophy, 4 (21), pp 161-173*, 1828.

[77] A. V. Cumming et al. Design and development of the advanced LIGO monolithic fused silica suspension. *Classical and Quantum Gravity, 29, 035003*, 2012.

[78] I. Martin. *Studies of materials for use in futureinterferometric gravitational wave detectors*. PhD thesis, School of Physics and Astronomy, University of Glasgow, 2009.

[79] V. B. Braginsky, M. L. Gorodetsky, and S. P. Vyatchanin. Thermo-refractive noise in gravitational wave antennae. *Physics Letters A, 271 (5-6), pp 303-307*, 2000.

[80] LIGO Scientific Collaboration. LIGO Hanford and LIGO Livingston. Online, accessed 8th July 2022, Courtesy Caltech/MIT/LIGO Laboratory. https://www.ligo.caltech.edu/images?category=photograph.

[81] Virgo collaboration. Virgo detector. Online, accessed 13th September 2022, Courtesy Caltech/MIT/LIGO Laboratory. https://www.ligo.caltech.edu/image/ligo20170927b.

[82] B. P. Abbott et al. LIGO: the laser interferometer gravitational-wave observatory. *Reports on Progress in Physics, 72 (7), 076901*, 2009.

[83] J. R. Smith. The path to the enhanced and advanced LIGO gravitational-wave detectors. *Classical and Quantum Gravity, 26 (11), 114013*, 2009.

[84] LIGO Scientific Collaboration. Advanced LIGO. *Classical and Quantum Gravity, 32 (7), 074001*, 2015.

[85] T. Accadia et al. Virgo: A laser interferometer to detect gravitational waves. *Journal of Instrumentation, 7, P03012*, 2012.

[86] F. Acernese et al. Advanced Virgo: A second-generation interferometric gravitational wave detector. *Classical and Quantum Gravity, 32 (2), 024001*, 2014.

[87] B. Willke et al. The GEO 600 gravitational wave detector. *Classical and Quantum Gravity, 19, pp 1377-1387*, 2002.

[88] GEO600. A brief history of GEO600. Online, accessed 8th July 2022. https://www.geo600.org/23920/a-brief-history-of-geo600.

[89] M. Saleem et al. The science case for LIGO-India. *Classical and Quantum Gravity, 39 (2), 025004*, 2022.

[90] B. Schutz. Networks of gravitational wave detectors and three figures of merit. *Classical and Quantum Gravity, 28 (12), 125023*, 2011.

[91] M. G. Beker et al. Seismic attenuation technology for the Advanced Virgo gravitational wave detector. *Physics Procedia 37, pp 1389-1397*, 2012.

[92] T. Accadia et al. The seismic superattenuators of the Virgo gravitational waves interferometer. *Journal of Low Frequency Noise Vibration and Active Control, 30, pp 63-79*, 2011.

[93] S. A. Hughes and K. S. Thorne. Seismic gravity-gradient noise in interferometric gravitational-wave detectors. *Physical Review D, 58, 122002*, 1998.

[94] F. Thies et al. Nd:YVO_4 high-power master oscillator power amplifier laser system for second-generation gravitational wave detectors. *Optics Letters, 44 (3), pp 719-722*, 2019.

[95] L. McCuller et al. Frequency-dependent squeezing for Advanced LIGO. *Physical Review Letters, 124, 171102*, 2020.

[96] LIGO Scientific Collaboration. Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light. *Nature Photonics, 7, pp 613–619*, 2013.

[97] B. P. Abbott et al. Observation of gravitational waves from a binary black hole merger. *Physical Review Letters, 116 (6), 061102*, 2016. https://doi.org/10.1103/PhysRevLett.116.061102.

[98] The Royal Swedish Academiy of Sciences. Press release: the Nobel prize in physics 2017. Online, accessed 7th July 2022. https://www.nobelprize.org/prizes/physics/2017/press-release/.

[99] Fundamental Physics Breakthrough Prize. Laureates of the Special Breakthrough Prize. Online, accessed 28th August 2022. https://breakthroughprize.org/Laureates/1/P4/Y2016.

[100] R. Abbott et al. GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run. *arXiV preprint, arXiV:2111.03606*, 2021.

[101] B. P. Abbott et al. Localization and broadband follow-up of the gravitational-wave transient GW150914. *Astrophysical Journal Letters, 826, L13*, 2016.

[102] B. P. Abbott et al. Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA. *Living Reviews in Relativity, 23, 3*, 2020.

[103] N. Arnaud. Status of the Advanced LIGO and Advanced Virgo detectors. In *EPS Conference on High Energy Physics*, 2017.

[104] B. P. Abbott et al. GW170817: Observation of gravitational waves from a binary neutron star inspiral. *Physcial Review Letters, 119, 161101*, 2017.

[105] B. P. Abbott et al. Multi-messenger observations of a binary neutron star merger. *Astrophysical Journal Letters, 848 (2), L12*, 2017. https://doi.org/10.3847/2041-8213/aa91c9.

[106] B. Zhang. The delay time of gravitational wave — gamma-ray burst associations. *Frontiers of Physics, 14, 64402*, 2018.

[107] R. Weiss. LIGO and Gravitational Waves I. Online, accessed 14th November 2022, 2017. https://www.nobelprize.org/prizes/physics/2017/weiss/lecture/.

[108] N. A. Robertson. Suspension design for Advanced LIGO. In *Aspen Winter Conference on Gravitational Waves*, 2004.

[109] A. Cumming et al. Finite element modelling of the mechanical loss of silica suspension fibres for advanced gravitational wave detectors. *Classical and Quantum Gravity, 26, 215012*, 2009.

[110] A. A. van Veggel. Cryogenic suspensions for future gravitational wave detectors. In *Gravitational Wave Advanced Detector Workshop*, 2018.

[111] L. G. Prokhorov et al. Upper limits on the mechanical loss of silicate bonds in a silicon tuning fork oscillator. *Physics Letters A, 382 (33), pp 2186-2191*, 2018.

[112] J. Miller et al. Prospects for doubling the range of Advanced LIGO. *Physical Review D, 91, 062005*, 2015.

[113] T. Isogai et al. Loss in long-storage-time optical cavities. *Optics Express, 21 (24), pp 30113-30125*, 2013.

[114] P. Kwee et al. Decoherence and degradation of squeezed states in quantum filter cavities. *Physical Review D, 90, 062006*, 2014.

[115] G. D. Hammond et al. Reducing the suspension thermal noise of advanced gravitational wave detectors. *Classical and Quantum Gravity, 29, 124009*, 2012.

[116] LIGO Scientific Collaboration. Report from the LSC post-O5 study group, LIGO-T2200287. Technical report, LIGO Scientific Collaboration, 2022.

[117] KAGRA Collaboration. Overview of KAGRA: Detector design and construction history. *Progress of Theoretical and Experimental Physics, 2021 (5), 05A101*, 2021.

[118] KAGRA Collaboration. KAGRA: 2.5 generation interferometric gravitational wave detector. *Nature Astronomy, 3, pp 35–40*, 2019.

[119] T. Ushiba et al. Cryogenic suspension design for a kilometer-scale gravitational-wave detector. *Classical and Quantum Gravity, 38 (8), 085013*, 2021.

[120] ICRR, University of Tokyo. An illustration of the underground KAGRA gravitational-wave detector in Japan. Online, accessed 22nd July 2022, Photo courtesy: KAGRA Observatory, ICRR, The University of Tokyo. https://www.ligo.caltech.edu/LA/image/ligo20191004b.

[121] LVK Collaboration. Results from the O3 GEO-KAGRA observing run. *Progress of Theoretical and Experimental Physics, 6, 063F01*, 2022. https://www.ligo.org/science/Publication-O3GEO-KAGRA/flyer.pdf.

[122] R. Adhikari et al. A cryogenic silicon interferometer for gravitational-wave detection. *Classical and Quantum Gravity, 37, 165003*, 2020.

[123] M. Evans et al. A horizon study for Cosmic Explorer: Science, observatories, and community. Technical report, Cosmic Explorer Consortium, 2021.

[124] D. Reitze. Cosmic Explorer: the U.S. contribution to gravitational-wave astronomy beyond LIGO. *Bulletin of the AAS, 51 (7), 035*, 2019.

[125] A. V. Cumming et al. Large-scale monolithic fused-silica mirror suspension for third-generation gravitational-wave detectors. *Physical Review Applied, 17, 024044*, 2022.

[126] Einstein Telescope Science Team. Einstein gravitational wave telescope conceptual design study. Technical report, Einstein Telescope, 2011.

[127] Einstein Telescope Editorial Team. Design report update 2020 for the Einstein Telescope. Technical report, Einstein Telescope, 2020.

[128] Nikhef. Artist's impression of the Einstein Telescope. Online, accessed 2nd August 2022. https://www.aei.mpg.de/einsteintelescope.

[129] S. Rowlinson et al. Feasibility study of beam-expanding telescopes in the interferometer arms for the Einstein Telescope. *Physical Review D, 103, 023004*, 2021. http://www.gwoptics.org/research/et/layout/.

[130] K. Ackley et al. Neutron Star Extreme Matter Observatory: A kilohertz-band gravitational-wave detector in the global network. *Publications of the Astronomical Society of Australia, 37, E047*, 2020.

[131] LISA Consortium. Laser interferometer space antenna. *arXiV preprint, arXiV:1702.00786*, 2017.

[132] É. É. Flanagan and S. A. Hughes. Measuring gravitational waves from binary black hole coalescences: I. signal to noise for inspiral, merger, and ringdown. *Physical Review D, 57 (8), pp 4535-4565*, 1998.

[133] P. Amaro-Seoane et al. Intermediate and extreme mass-ratio inspirals - astrophysics, science applications and detection using LISA. *Classical and Quantum Gravity, 24 (17), R113-169*, 2007.

[134] P. Binétruy et al. Cosmological backgrounds of gravitational waves and eLISA/NGO: phase transitions, cosmic strings and other sources. *Journal of Cosmology and Astroparticle Physics, 2012, 027*, 2012.

[135] LISA Consortium. LISA constellation. Online, accessed 2nd August 2022,. https://www.lisamission.org/multimedia/image/lisa-astro2020.

[136] A. V. Cumming et al. Silicon mirror suspensions for gravitational wave detectors. *Classical and Quantum Gravity, 31 (2), 025017*, 2013.

[137] H. B. Callen. Irreversibility and generalised noise. *Physical Review, 83, pp 34-40*, 1951.

[138] H. B. Callen and R. F. Greene. On a theorem of irreversible thermodynamics. *Physical Review, 86, pp 702-710*, 1950.

[139] C. Affeldt et al. Advanced techniques in GEO 600. *Classical and Quantum Gravity, 31 (22), 224002*, 2014.

[140] G. Gonzalez. Suspensions thermal noise in the LIGO gravitational wave detector. *Classical and Quantum Gravity, 17, pp 4409-4435*, 2000.

[141] F. Acernese et al. Status of Advanced Virgo. In *The European Physical Journal Conferences, 182, 02003*, 2018.

[142] S. Sakata et al. A study for reduction of radiation pressure noise in gravitational wave detectors. *Journal of Physics: Conference Series, 122, 012020*, 2008.

[143] G. M. Harry et al. Thermal noise ininterferometric gravitational wave detectors due to dielectric optical coatings. *Classical and Quantum Gravity, 19 (5), pp 897-917*, 2002.

[144] A. V. Cumming. *Aspects of mirrors and suspensions for advanced gravitational wave detectors*. PhD thesis, School of Physics and Astronomy, University of Glasgow, 2008.

[145] R. Adhikari, P. Fritschel, and S. Waldman. Enhanced LIGO, LIGO-T060156. Technical report, LIGO Scientific Collaboration, 2006.

[146] S. Rowan et al. The quality factor of natural fused quartz ribbons over a frequency range from 6 to 160 Hz. *Physics Letters A, 227, pp 153-158*, 1997.

[147] J. R. Smith et al. Mechanical quality factor measurements of monolithically-suspended fused silica test masses of the GEO600 gravitational wave detector. *Classical and Quantum Gravity, 21, pp S1091-1098*, 2004.

[148] A. V. Cumming et al. Lowest observed surface and weld losses in fused silica fibres for gravitational wave detectors. *Classical and Quantum Gravity, 37, 195019*, 2020.

[149] A. Heptonstall. *Characterisation of mechanical loss in fused silicon ribbons for used in gravitational wave detector suspensions*. PhD thesis, School of Physics and Astronomy, University of Glasgow, 2004.

[150] K. Haughian. *Aspects of materials research for advanced and future generations of gravitational wave detectors*. PhD thesis, School of Physics and Astronomy, University of Glasgow, 2012.

[151] H. J. Pain. *The physics of vibrations and waves*. John Wiley and Sons, 2005.

[152] P. Murray. *Measurement of the mechanical loss of test mass materials for advanced gravitational wave detectors*. PhD thesis, School of Physics and Astronomy, University of Glasgow, 2007.

[153] R. Nawrodt et al. Investigation of mechanical losses of thin silicon flexures at low temperatures. *Classical and Quantuam Gravity, 30, 115008*, 2013.

[154] D. F. McGuigan et al. Measurements of the mechanical Q of single-crystal silicon at low temperatures. *Journal of Low Temperature Physics, 30 (5-6), pp 621-629*, 1978.

[155] R. Nawrodt et al. High mechanical Q-factor measurements on silicon bulk samples. *Journal of Physics: Conference Series, 122, 012008*, 2008.

[156] A. Schroeter et al. On the mechanical quality factors of cryogenic test masses from fused silica and crystalline quartz. *arXiV preprint, arXiV:0709.4359*, 2007.

[157] V. B. Braginsky, V. P. Mitrofanov, and V. I. Panov. *Systems with small dissipation*. University of Chicago Press, 1985.

[158] J. Carvill. *Mechanical engineer's data handbook*. Butterworth-Heinemann, 1994.

[159] D. H. Damon. Thermal conductivity of vitreous silica at low temperatures. *Physical Review B, 8 (12), pp 5860-5865*, 1973.

[160] S. Miyoki et al. CLIO cryogenic laser interferometer observatory. In *28th International Cosmic Ray Conference*, 2003.

[161] A.-M. van Veggel. Quasi-monolithic mirror suspensions in ground-based gravitational-wave detectors: an overview and look to the future. *Philosophical Transactions of the Royal Society A, Mathematical, Physical and Engineering Sciences, 376 (2120), 20170281*, 2018.

[162] A. Dari, F. Travasso, H. Vocca, and L. Gammaitoni. Breaking strength tests on silicon and sapphire bondings for gravitational wave detectors. *Classical and Quantum Gravity, 27 (4), 045010*, 2010.

[163] The Editors of Encyclopædia. *Encyclopædia Britannica, silicon*. Britannica Books, 2021. Accessed 17 Feb 2021, https://www.britannica.com/science/silicon.

[164] Mitshibisi Materials Trading Corporation. Silicon wafers. Online, accessed 25th April 2022. https://www.mmtc.co.jp/en/products/silicon-s.html.

[165] N. Sarukura et al. *Handbook of crystal growth*. Elsevier, 2015.

[166] W. Lin and C. W. Pearce. Properties of uniform oxygen Czochralski silicon crystals. *Journal of Applied Physics, 51, pp 5540-5543*, 1980.

[167] V. Savolainen et al. Simulation of large-scale silicon melt flow in magnetic Czochralski growth. *Journal of Crystal Growth, 234, pp 243-260*, 2002.

[168] P. Rudolph and K. Kakimoto. Crystal growth from the melt under external force fields. *MRS Bulletin, 34, pp 251-258*, 2009.

[169] Y. B. Gianchandani, O. Tabata, and H. Zappe. *Comprehensive microsystems, vol. 1*. Elsevier, 2008.

[170] A. Muiznieks et al. *Handbook of crystal growth (second edition)*. Elsevier, 2015.

[171] A. Ayodele. Silicon wafers: Everything you need to know. Online, accessed 18th May 2022. https://www.wevolver.com/article/silicon-wafers-everything-you-need-to-know.

[172] M. Watanabe and S. Kramer. 450 mm silicon: An opportunity and wafer scaling. *The Electrochemical Society, Interface, pp 28-31*, 2006. https://www.eetimes.com/collaborative-advantage-design-impact-of-450mm-transition/.

[173] W. C. O'Mara, R. B. Herring, and L. P. Hunt. *Handbook of semiconductor silicon technology*. Elsevier, 1990.

[174] M. A. Hopcroft, W. D. Nix, and T. W. Kenny. What's the Young's modulus of silicon? *Journal of Microelectromechanical Systems, 19 (2), pp 229-238*, 2010.

[175] EL-CAT Inc. Properties of silicon and silicon wafers. Online, accessed 24th April 2022. https://www.el-cat.com/silicon-properties.htm.

[176] A. S. Nowick and B. S. Berry. *Anelastic relaxation in crystalline solids*. Academic Press, 1972.

[177] V. B. Braginsky, M. L. Gorodetsky, and S. P. Vyatchanin. Thermodynamical fluctuations and photo-thermal shot noise in gravitational wave antennae. *Physics Letters A, 264, pp 1-10*, 1999.

[178] Y. T. Liu and K. S. Thorne. Thermoelastic noise and homogeneous thermal noise in finite sized gravitational-wave test masses. *Physical Review D, 62, 122002*, 2000.

[179] X. Liu et al. A loss mechanism study of a very high Q silicon micromechanical oscillator. *Journal of Applied Physics, 97, 023524*, 2005.

[180] T. Metcalf, X. Liu, and M. Abernathy. Improving the mechanical quality factor of ultra-low-loss silicon resonators. *Journal of Applied Physics, 123, 235105*, 2018.

[181] B. S. Lunin. *Physical and chemical bases for development of hemispherical resonators for solid-state gyroscopes*. 2005.

[182] M. Morita et al. Growth of native oxide on a silicon surface. *Journal of Applied Physics, 68, pp 1272-1281*, 1990.

[183] B. White and R. Pohl. Internal friction of subnanometer alpha a-SiO2 films. *Physical Review Letters, 75 (24), pp 4437-4439*, 1995.

[184] A. M. Gretarsson and G. M. Harry. Dissipation of mechanical energy in fused silica fibers. *Review of Scientific Instruments, 70 (10), pp 4081-4087*, 1999.

[185] G. Harry. Mechanical loss of epoxies. In *Department of Physics, American University*, 2014.

[186] M. Phelps et al. Strength of hydroxide catalysis bonds between sapphire, silicon, and fused silica as a function of time. *Physical Review D, 98, 122003*, 2018.

[187] S. Rowan et al. Mechanical losses associated with the technique of hydroxide-catalysis bonding of fused silica. *Physics Letters A, 246 (6), pp 417-478*, 1998.

[188] Virgo Collaboration. Status of Advanced Virgo and upgrades before next observing runs. *Physica Scripta, 96 (12), 124054*, nov 2021.

[189] S. M. Aston et al. Update on quadruple suspension design for Advanced LIGO. *Classical and Quantum Gravity, 29 (23), 235004*, 2012.

[190] N. Beveridge. *Characterisation of silicon-silicon hydroxide catalysis bonds for future gravitational wave detectors*. PhD thesis, School of Physics and Astronomy, University of Glasgow, 2012.

[191] N. L. Beveridge et al. Low-temperature strength tests and SEM imaging of hydroxide catalysis bonds in silicon. *Classical and Quantuam Gravity, 28, 085014*, 2011.

[192] R. Nawrodt et al. Cryogenic setup for Q-factor measurements on bulk materials for future gravitational wave detectors. In *Proceedings of ICEC22-ICMC2008*, 2009.

[193] S. Reid et al. Mechanical dissipation in silicon flexures. *Physics Letters A, 351 (4-5), pp 205-211*, 2006.

[194] M. Alshourbagy et al. First characterization of silicon crystalline fibers produced with the μ-pulling technique for future gravitational wave detectors. In *Review of Scientific Instruments, 77, 044502*, 2006.

[195] J. Rodriguez et al. Direct detection of anchor damping in MEMS tuning fork resonators. *Journal of Microelectromechanical Systems, 99, pp 1-10*, 2018.

[196] R. Nawrodt et al. Mirror thermal noise calculation for ET. Technical report, Einstein Telescope, 2009.

[197] S. Rowan et al. Test mass materials for a new generation of gravitational wave detectors. In P. Saulson, editor, *Gravitational-Wave Detection*, volume 4856. Proceedings of SPIE, 2003.

[198] B. Shapiro et al. Cryogenically cooled ultra low vibration silicon mirrors for gravitation-alwave observatories. *Cryogenics, 81, pp 83-92*, 2017.

[199] R. Adhikari. Gravitational radiation detection with laser interferometry. *Reviews of Modern Physics, 86, 121*, 2014.

[200] T. Tomaru et al. Maximum heat transfer along a sapphire suspension fiber for a cryogenic interferometric gravitational wave detector. *Physics Letters A, 304 (3-4), pp 215-219*, 2002.

[201] M. Asheghi et al. Temperature-dependent thermal conductivity of single-crystal silicon layers in SOI substrates. *Journal of Heat Transfer, 120 (1), pp 30-36*, 1998.

[202] M. M. Herrera. *Properties of bonded silicon for future gravitational wave observatories*. PhD thesis, School of Physics and Astronomy, University of Glasgow, 2019.

[203] R. A. Richardson and F. Nori. Transport and boundary scattering in confined geometries: Analytical results. *Physical Review B, 48 (20), pp 15209-15217*, 1993.

[204] R. G. Scurlock. *Low temperature behaviour of solids*. Routledge and Kegan Paul, 1966.

[205] Y. S. Touloukian, R. W. Powell, C. Y. Ho., and P. G. Klemens. *Thermophysical properties of matter - the TPRC data series, volume 4, specific heat - metallic elements and alloys*. New York, IFI/Plenum, 1970.

[206] D. Bridges. (QUAD) quadruple suspension controls arrangement poster, LIGO-E1000617. Technical report, LIGO, 2012.

[207] C. Brif. Notes on anelastic effects and thermal noise in suspensions of test masses in-interferometric gravitational-wave detectors, LIGO-T990041. Technical report, LIGO, 1999.

[208] A. Gillespie and F. Raab. Thermal noise in the test mass suspensions of a laser interferometer gravitational-wave detector prototype. *Physics Letters A, 178 (5-6), pp 357-363*, 1993.

[209] J. E. Logan, J. Hough, and N. A. Robertson. Aspects of the thermal motion of a mass suspended as a pendulum by wires. *Physics Letters A, 183 (2-3), pp 142-152*, 1993.

[210] P. Amico et al. Monocrystalline fibres for low thermal noise suspension in advanced gravitational wave detectors. *Classical and Quantum Gravity, 21, S1009-1013*, 2004.

[211] N. Healy et al. CO_2 laser-induced directional recrystallization to produce single crystal silicon-core optical fibers with low loss. *Advanced Optical Materials, 4, pp 1004-1008*, 2016.

[212] S. Morris et al. The influence of core geometry on the crystallography of silicon optical fiber. *Journal of Crystal Growth, 352, pp 53-58*, 2012.

[213] M. Nicolai et al. Lumped parameter model for silicon crystal growth from granulate crucible. *Crystal Research and Technology, 55, 2000044*, 2020.

[214] K. Dadzis et al. Characterization of silicon crystals grown from melt in a granulate crucible. *Journal of Electronic Materials, 49, pp 5120-5132*, 2020.

[215] R. Menzel et al. Leibniz Institute for Crystal Growth. In *Einstein Telescope Meeting*, 2020.

[216] F. M. Kiessling et al. Quasi-monocrystalline silicon for low-noise end mirrors in cryogenic gravitational-wave detectors. *Physical Review Research*, 2022.

[217] J. Zhang and K. Hoshino. *Molecular sensors and nanodevices (second edition)*. Academic Press, 2019.

[218] G. Lacaille. Characterisation of bonded silicon towards assemblies for 3rd generation detectors. In *GWADW 2022, LIGO G2200872-v1*, 2022.

[219] M. Phelps. *Hydroxide catalysis and indium bonding research for the design of ground-based gravitational wave detectors*. PhD thesis, School of Physics and Astronomy, University of Glasgow, 2018.

[220] R. Jones. CO_2 laser pulling machine, LIGO-D070560. Technical report, Institute for Gravitational Research, University of Glasgow, September 2019. Version 2.

[221] K. Mühlberger, C. M. Harvey, and M. Fokine. Temperature dynamics in silicon core fibers during CO_2 laser processing. *Optics Express, 30 (1), pp 92-100*, 2022.

[222] D. Chen et al. Vibration measurement in the KAGRA cryostat. *Classical and Quantum Gravity, 31 (22), 224001*, 2014.

[223] T. Ushiba. Cryogenic mirror suspension in KAGRA. In *Proceedings of the 1st Gravitational Wave Science and Technology Symposium*, 2019.

[224] R. Birney et al. Coatings and surface treatments for enhanced performance suspensions for future gravitational wave detectors. *Classical and Quantum Gravity, 34, 235012*, 2017.

[225] G. Cagnoli et al. Silica suspension and coating developments for Advanced LIGO. *Journal of Physics: Conference Series, 32, pp 386-392*, 2006.

[226] E. D. Hall et al. Gravitational-wave physics with Cosmic Explorer: limits to low-frequency sensitivity. *Physical Review D, 103, 122004*, 2021.

[227] K. Toland et al. Development of a pulling machine to produce micron diameter fused silica fibres for use in prototype advanced gravitational wave detectors. *Classical and Quantum Gravity, 35, 165004*, 2018.

[228] R. Beardmore. Basic notes on factor of safety. Online, accessed 15th March 2022. https://roymech.org/Useful_Tables/ARM/Safety_Factors.html.

[229] G. Hammond. Suspending the test masses of gravitational wave detectors. In *SUSSP 73, St Andrews*, 2017.

[230] N. Rau et al. Limits of very small manufacturing processes. Online, accessed 23rd May 2022. https://inst.eecs.berkeley.edu/~ee290g/fa97/small.html.

[231] A. S. Bell et al. Quad monolithic fiber pulling/welding procedure, LIGO-E1000366. Technical report, LIGO, 2017.

[232] K. E. Petersen. Silicon as a mechanical material. In *Proceedings of the IEEE, 70 (5), pp 420-457*. IEEE, 1982.

[233] W. D. Sylwestrowicz. Mechanical properties of single crystals of silicon. *Philosophical Magazine, 7 (83). pp 1825-1845*, 1962.

[234] Y. Umeno et al. Ab initio study of the surface properties and ideal strength of (100) silicon thin films. *Physical Review B, 72 (16)*, 2005.

[235] D. Roundy and M. L. Cohen. Ideal strength of diamond, Si, and Ge. *Physical Review B, 64 (21), 212103*, 2001.

[236] H. Zhang et al. Approaching the ideal elastic strain limit in silicon nanowires. *Science advances, 2 (8)*, 2016.

[237] M. N. Esfahani and B. E. Alaca. A review on size-dependent mechanical properties of nanowires. *Advanced Engineering Materials, 21, 1900192*, 2019.

[238] J. C. McLaughlin and A. F. W. Willoughby. Fracture of silicon wafers. *Journal of Crystal Growth, 85, pp 83-90*, 1987.

[239] U. Bismayer et al. Measurement of subsurface damage in silicon wafers. *Precision Engineering, 16 (2), pp 139-144*, 1994.

[240] H.-F. Hadamovsky. *Werkstoffe der halbleitertechnick*. Dt. Verlag für Grundstoffindustrie, 1990.

[241] L. Gert. Structure of the surface layers of molybdenum single crystals aftermachining. *Physics, Chemistry and mechanics of Surfaces, 2, pp 583-598*, 1984.

[242] I. Zarudi and L. Zhang. Subsurface damage in single-crystal silicon due to grinding and polishing. *Journal of Materials Science, 15, pp 586-587*, 1996.

[243] J. Yin, Q. Bai, Y. Lia, and B. Zhang. Formation of subsurface cracks in silicon wafers by grinding. *Nanotechnology and Precision Engineering 1, 172*, 2018.

[244] B. Schwartz and H. Robbins. Chemical etching of silicon (IV): Etching technology. *Journal of The Electrochemical Society, 123 (12), pp 1903-1909*, 1976.

[245] Chemistry Libretexts. Intermolecular forces. Online, accessed 2nd February 2022, August 2020. https://chem.libretexts.org/Bookshelves/Physical_and_Theoretical_Chemistry_Textbook_Maps/Supplemental_Modules_(Physical_and_Theoretical_Chemistry)/Physical_Properties_of_Matter/Atomic_and_Molecular_Properties/Intermolecular_Forces.

[246] A. C. Fischer-Cripps. *Introduction to contact mechanics*. Springer, 2007.

[247] Y. V. Milman, I. V. Gridneva, and A. A. Golubenko. Construction of stress-strain curves for brittle materials by indentation in a wide temperature range. *Science of Sintering, 39, pp 67-75, https://doi.org/10.2298/SOS0701067M*, 2007.

[248] P. Tabeling. *Introduction to microfluidics*. Oxford University Press, 2011. Translated by Chen, S.

[249] A. Schifferle, T. Bandi, A. Neels, and A. Dommann. Where is the limit? yield strength improvement in silicon micro-structures by surface treatments. *Physica Status Solidi (A) Applications and Materials, 213 (1), pp 102-107*, 2015.

[250] A. A. Griffith. The phenomena of rupture and flow in solids. *Philosophical Transactions of the Royal Society of London. Series A, 221, pp 163–198*, 1920.

[251] C. E. Inglis. Stresses in a plate due to the presence of cracks and sharp corners. *SPIE Milestone Series, 137, pp 3-17*, 1913.

[252] E. Orowan. Fracture and strength of solids. *Reports on Progress in Physiscs, 12, pp 185-232*, 1949.

[253] G. R. Irwin. Analysis of stresses and strains near the end of a crack traversing a plate. *Applied Mechanics, 24, pp 361-364*, 1957.

[254] G. Cagnoli et al. Developments toward monolithic suspensions for advanced gravitational wave detectors. *The Eleventh Marcel Grossmann Meeting, pp 2379-2381*, 2008.

[255] K. V. Tokmakov et al. A study of the fracture mechanisms in pristine silica fibres utilising high speed imaging techniques. *Journal of Non-Crystalline Solids*, 2012.

[256] K.-H. Lee et al. Improved fused silica fibres for the advanced LIGO monolithic suspensions. *Classical and Quantum Gravity, 36, 185018*, 2019.

[257] Laser Micromachining Ltd. Laser Micromachining Ltd. Online, accessed 24th January 2022. https://www.lasermicromachining.com/.

[258] A. Bell. Private communication. Email, circa 2018/19.

[259] N. Rizvi. Private communication. Email, 24th January 2022.

[260] N. Rizvi. Private communication. Email, 30th May 2018. Technical report but confidential.

[261] Contec. Anticon gold wipes. Online, accessed 24th January 2022. https://www.contecinc.com/eu/products/anticon-gold-heavyweight/.

[262] Wyko. Wyko NT1100 optical surface profiler. Online, accessed 25th January 2022. http://www.physics.gla.ac.uk/igr/index.php?L1=capabilities&L2=equipment&L3=optical_profiler.

[263] Hitachi. Hitachi TM-1000 scanning electron microscope. Online, accessed 25th January 2022, translated from Japanese. https://www.hitachi-hightech.com/jp/about/news/2005/nr20050201.html.

[264] P. Romero et al. Experimental study of diode laser cutting of silicon by means of water assisted thermally driven separation mechanism. *Physics Procedia 41, pp 617-626*, 2013.

[265] RS. RS PRO 150mm digital caliper. Online, accessed 25th January 2022. https://uk.rs-online.com/web/p/calipers/8412518.

[266] A. Cumming et al. Apparatus for dimensional characterization of fused silica fibers for the suspensions of advanced gravitational wave detectors. *Review of Scientific Instruments, 82, 044502*, 2011.

[267] A. V. Cumming. Dimensional characterisation equipment for fused silica suspension fibres. Online, accessed 21 January 2022, January 2010.

[268] K. Toland. *Aspects of fused silica fibres for use in gravitational waves research*. PhD thesis, School of Physics and Astronomy, University of Glasgow, 2020.

[269] Wikipedia user Wizard191. Tensile specimens made from an aluminum alloy. Online, accessed 25th January 2022, 2010. https://en.wikipedia.org/wiki/Tensile_testing#/media/File:Tensile_specimen-round_and_flat.jpg.

[270] R. Jones. CO_2 laser pulling machine ribbon pull fork holster, LIGO-D070555. Technical report, Institute for Gravitational Research, University of Glasgow, December 2009.

[271] Huntsman. Araldite 2012. Online, accessed 20th June 2022. https://www.huntsman.com/products/araldite2000/araldite-2012.

[272] R. Jones et al. The design of the IGR MkII strength testing machine, LIGO-T1000345. Technical report, Institute for Gravitational Research, University of Glasgow, September 2010. Version 2.

[273] Omega Engineering, Inc. Omega DP25B-S-DC10 manual. Online, accessed 1st April 2022. https://assets.omega.com/manuals/M3598.pdf.

[274] Phantom. Phantom VEO product family. Online, accessed 3rd March 2022. https://www.adept.net.au/cameras/visionresearch/pdf/PhantomVEO.pdf.

[275] R. J. Jaccodine. The surface energy of Si, GaAs, and GaP. *Journal of Applied Physics, 52, pp 4623-4629*, 1981.

[276] R. Pérez and P. Gumbsch. Directional anisotropy in the cleavage fracture of silicon. *Physical Review Letters, 84 (23), pp 5347-5350*, 2000.

[277] T. Dobrescu, P. Nicoleta-Elisabeta, M. Ghinea, and A. Popescu. Plane lapping process of silicon wafers. In *Proceedings of the 1st International Conference on Mechanical and Robotics Engineering*, 2013.

[278] Y. Wang, Y. Zhao, and X. Chen. Chemical mechanical planarization from macro-scale to molecular-scale. *Materials and Manufacturing Processes, 27 (6), pp 641-649*, 2012.

[279] W. Heybore, G. Spierings, and J. van den Meerakker. Chemomechanical silicon polishing. *Journal of Electrochemical Society, 138 (3), pp 774-777*, 1992.

[280] T. Palmer, B. Lynch, C. Roberts, and Y. Lu. Attenuated total reflectance spectra of colloidal and particulate silica samples in aqueous media. *Applied Spectroscopty, 45 (6), pp 1022-1025*, 1991.

[281] J. Hasima, G. Spierings, T. Michielsen, and C. Adema. Surface preparation and phenomenological aspects of direct bonding. *Philips Journal of Reseasrch 49, pp 23-46*, 1995.

[282] R. B. McIntosh and R. A. Paquin. Chemical-mechanical polishing of low-scatter optical surfaces. *Optical Society of America, 19 (14), pp 2329-2331*, 1980.

[283] S. Hah and T. Fischer. Tribochemical polishing of silicon nitride. *Journal of Electrochemical Society, 145 (5), pp 1708-1714*, 1998.

[284] L. Chen, Y. Qi, B. Yu, and L. Qian. Sliding speed-dependent tribochemical wear of oxide-free silicon. *Nanoscale Research Letters, 12 (1), 404*, 2017.

[285] Logitech Ltd. Logitech PM5 lapping and polishing system. Online, accessed 31st January 2022. https://logitech.uk.com/wp-content/uploads/2016/07/generalbrochure_website-1.pdf.

[286] Logitech Ltd. Fused aluminium oxide powder. Online, accessed 31st January 2022. https://us.store.logitech.uk.com/index.php?route=product/product&product_id=58.

[287] T. McGregor. Private communication. Meeting, 9th October 2018.

[288] Logitech Ltd. Syton SF1 polishing fluid. Online, accessed 31st January 2022. https://us.store.logitech.uk.com/index.php?route=product/product&product_id=97.

[289] Olympus. Olympus BX51 optical microscope. Online, accessed 9th Feburary 2022. https://www.olympus-ims.com/en/microscope/bx51p/.

[290] Kruss Scientific. Drop shape analyzer DSA30B. Online, accessed 8th February 2022. https://www.kruss-scientific.com/en/products-services/products/dsa30b.

[291] RS. RS PRO digital micrometer, UKAS calibrated. Online, accessed 14th February 2022. https://uk.rs-online.com/web/p/micrometers/7048771.

[292] L. Pinard et al. Mirrors used in the LIGO interferometers for first detection of gravitational waves. *Applied Optics, 56 (4), C11-15*, 2016.

[293] C. Bundesmann and H. Neumann. Tutorial: The systematics of ion beam sputtering for deposition of thin films with tailored properties. *Applied Phyiscs, 123, 231102*, 2018.

[294] C. Patrick. Ion beam sputter deposition allows tailoring of thin film properties. Online, accessed 22nd February 2022, December 2018. https://aip.scitation.org/doi/10.1063/1.5086197.

[295] B. Wolf. *Handbook of ion sources*. CRC Press, 1995.

[296] S. Reid. Private communication. Email, 14th October 2019.

[297] A. and D. Weighing. Gemini GR-202 microbalance. Online, accessed 23rd February 2022. https://weighing.andonline.com/product/gemini-gr-series-balances/gr-202?commerce_product=23.

[298] DuPont. DuPont kapton polyimide film. Online, accessed 23rd February 2022. https://www.dupont.com/content/dam/dupont/amer/us/en/products/ei-transformation/documents/DEC-Kapton-HN-datasheet.pdf.

[299] R. P. Middlemiss et al. Measurement of the Earth tides with a MEMS gravimeter. *Nature, 531, pp 614–617*, 2016.

[300] P. Pal et al. High speed silicon wet anisotropic etching for applications in bulk micromachining: a review. *Micro and Nano Systems Letters, 9, 4*, 2021.

[301] D. Kendall and R. Shoultz. *Handbook of microlithography, micromachining, and microfabrication, volume 2: micromachining and microfabrication*. SPIE Press, 1997.

[302] Z. Bai et al. Improved HNA isotropic etching for large-scale highly symmetric toroidal silicon molds with <10-nm roughness. *Journal of Micro/Nanolithography, 18 (4), 044501*, 2019.

[303] B. Schwartz and H. Robbins. Chemical etching of silicon (III): A temperature study in the acid system. *Journal of The Electrochemical Society, 108 (4), pp 365-372*, 1961.

[304] B. Schwartz and H. Robbins. Chemical etching of silicon (II): The system of HF, HNO_3, H_2O, and $HC_2H_3O_2$. *Journal of The Electrochemical Society, 107 (2), pp 108-111*, 1960.

[305] P. Pal and S. Singh. A new model for the etching characteristics of corners formed by Si(111) planes on Si(110) wafer surface. *Engineering, 4 (11A)*, 2013.

[306] A. V. Narasimha Rao, V. Swarnalatha, and P. Pal. Etching characteristics of Si(110) in 20 percent KOH with addition of hydroxylamine for the fabrication of bulk micromachined MEMS. *Micro and Nano Systems Letters, 5 (23)*, 2017.

[307] D. Mirabella, G. Suarez, M. Suarez, and C. Aldao. Silicon wet etching: Hillock formation mechanisms and dynamic scaling properties. *Physica A, 395, pp 105-111*, 2014.

[308] S. Tan, M. Reed, H. Han, and R. Boudreau. Mechanisms of etch hillock formation. *Journal of Microelectromechanical Systems, 5 (1), pp 66-72*, 1996.

[309] H. Seidel, L. Cspregi, A. Heuberger, and G. Baumgartel. Anisotropic etching of crystalline silicon in alkaline solutions (I): I. orientation dependence and behavior of passivation layers. *Journal of The Electrochemical Society, 137 (11), pp 3612-3626*, 1990.

[310] J. D. Plummer, M. Deal, and P. D. Griffin. *Silicon VLSI technology: Fundamentals, practice, and modeling*. Pearson, 2001.

[311] A. S. Louro and J. R. Senna. Real-time, in-situ microscopic observation of bubbles and roughening in KOH etching of silicon. In *Micromachining and Microfabrication Process Technology VII*, volume 4557. Proceedings of SPIE - The International Society for Optical Engineering, 2001.

[312] I. Zubel and M. Kramkowska. Effect of isopropyl alcohol (IPA) on etching rate and surface roughness of silicon etched in KOH solution. *Sensors and Actuators A: Physical, 93 (2), pp 138-147*, 2001.

[313] P. Dobson. Private communication. Meeting, 30th April 2019.

[314] University of Glasgow. James Watt Nanofabrication Centre. Online, accessed 7th March 2022. https://www.gla.ac.uk/research/az/jwnc/.

[315] ITM Mechanical Services Limited. ITM Mechanical Services Limited. Online, accessed 7th March 2022. https://find-and-update.company-information.service.gov.uk/company/SC368381.

[316] M. Smith. Private communication. Email, 22nd October 2018.

[317] M. Smith. Risk assessment of JWNC 3:2:1 HNA etch. Technical report, James Watt Nanofabrication Centre, 2018. Accessed 9th March 2022, https://webapps.eng.gla.ac.uk/tools/risk/view.php?action=view&id=6e4853776e11aff6280a0dfccb4b8c88.

[318] Olympus. Olympus Nomarski DIC microscope. Online, accessed 8th February 2022. https://www.olympus-lifescience.com/en/microscope-resource/primer/techniques/dic/dicconfiguration/.

[319] R. Hull. *Properties of crystalline silicon*. INSPEC, the Institution of Electrical Engineers, 1999.

[320] Virginia Semiconductor, Inc. Wet-chemical etching and cleaning of silicon. Technical report, Virginia Semiconductor, Inc., 2003.

[321] D. Li, editor. *Encyclopedia of microfluidics and nanofluidics, oxidation (of silicon)*. Springer, 2014.

[322] V. Svetovoy, J. W. Berenschot, and M. C. Elwenspoek. Precise test of the diffusion-controlled wet isotropic etching of silicon via circular mask openings. *Journal of The Electrochemical Society, 153 (9), pp C641-C647*, 2006.

[323] Galvetek. Industrial etching plants. Online, accessed 7th April 202. https://www.galvatek.eu/process-plants/industrial-etching/.

[324] E. Amaldi et al. Initial operation at liquid-helium temperature of the M=2270 kg Al 5056 gravitational-wave antenna of the Rome group. *Il Nuovo Cimento C, 7, pp 338-354*, 1984.

[325] E. Bonilla et al. Noise requirements of the cryogenic shielding for next generation cryocooled gravitational wave observatories: Newtonian noise. *Physical Review D, 104, 122005*, 2021.

[326] Y. Sakakibara et al. A study of cooling time reduction of interferometric cryogenic gravitational wave detectors using a high-emissivity coating. In *AIP Conference Proceedings, 1573, pp 1176-1183*, 2014.

[327] Y. Ikushima et al. Ultra-low-vibration pulse-tube cryocooler system – cooling capacity and vibration. *Cryogenics, 48, pp 406-412*, 2008.

[328] T. Yamada et al. High performance thermal link with small spring constant for cryogenic applications. *Cryogenics, 116*, 2021.

[329] T. Yamada. *Low-vibration conductive cooling of KAGRA cryogenic mirror suspension*. PhD thesis, The University of Tokyo, 2021.

[330] Leiden Cryogenics. Leiden Cryogenics. Online, accessed 17th May 2022. http://www.leiden-cryogenics.com/.

[331] Leiden Cryogenics. Cryogenic-free dilution refrigerator user manual. Manual.

[332] Pfeiffer Vacuum. TPR 280, TPR 281 safety technical data installation. Online, accessed 22nd June 2022. https://www.idealvac.com/files/manuals/Pfeiffer_ActiveLine_TPR_280_TPR_281_Pirani_Gauge_Operating_Instructions.pdf.

[333] Edwards. Edwards wide range gauge WRG-S-NW25. Online, accessed 22nd June 2022. https://shop.edwardsvacuum.com/products/d14701000/view.aspx.

[334] Lake Shore. Sensors with lead extensions. Online, accessed 30th May 2022. https://www.lakeshore.com/Lead-Extensions.

[335] Lake Shore. Cryogenic wire. Online, accessed 30th May 2022. https://www.lakeshore.com/products/categories/overview/temperature-products/cryogenic-accessories/cryogenic-wire.

[336] Lake Shore. Appendix E: Temperature measurement system. Online, accessed 31st May 2022. https://www.lakeshore.com/docs/default-source/temperature-catalog/lstc_appendixe_l.pdf?sfvrsn=b5177ae1_8.

[337] RS. Belden control cable 4 cores. Online, accessed 31st May 2022. https://uk.rs-online.com/web/p/control-cable/5102555.

[338] Lake Shore. Lake Shore model 340 temperature controller user's manual. Online, accessed 31st May 2022. https://www.lakeshore.com/docs/default-source/product-downloads/340_manual.pdf?sfvrsn=a7420fdf_1.

[339] Lake Shore. Lake Shore model 336 temperature controller user's manual. Online, accessed 31st May 2022. https://www.lakeshore.com/docs/default-source/product-downloads/336_manual0ebc9b06cbbb456491c65cf1337983e4.pdf?sfvrsn=2e8633a3_1.

[340] Artisan Technology Group. Lake Shore Cryotronics 340 temperature controller. Online, accessed 31st May 2022. https://www.artisantg.com/Scientific/81851-1/Lake-Shore-Cryotronics-340-Temperature-Controller.

[341] Lake Shore. DT-670 silicon diodes. Online, accessed 14th May 2022. https://www.lakeshore.com/products/categories/overview/temperature-products/cryogenic-temperature-sensors/dt-670-silicon-diodes.

[342] Quantum Design. Cryogenic temperature sensors. Online, accessed 15th June 2022. https://qd-uki.co.uk/lake-shore/temperature-sensors/gaalas-diodes/.

[343] Lake Shore. DT-670 silicon diode specifications. Online, accessed 15th June 2022. https://www.lakeshore.com/products/categories/specification/temperature-products/cryogenic-temperature-sensors/dt-670-silicon-diodes.

[344] Lake Shore. Cernox specifications. Online, accessed 15th June 2022. https://www.lakeshore.com/products/categories/specification/temperature-products/cryogenic-temperature-sensors/cernox.

[345] Lake Shore. Appendix F: PID temperature control. Online, accessed 15th June 2022. https://www.lakeshore.com/docs/default-source/product-downloads/literature/lstc_appendixf_l.pdf?sfvrsn=b044ec06_4.

[346] Lake Shore. Model 340 typical sensor performance. Online, accessed 15th June 2022. https://www.lakeshore.com/products/product-detail/model-340/Typical-Sensor-Performance.

[347] Lake Shore. Appendix G: Sensor temperature response data tables. Online, accessed 15th June 2022. https://www.lakeshore.com/docs/default-source/software/germanium-rtd-files/lstc_appendixg_h.pdf?sfvrsn=232b0713_2.

[348] Lake Shore. Model 340 specifications. Online, accessed 15th June 2022. https://www.lakeshore.com/products/categories/specification/discontinued-products/discontinued-products/model-340-cryogenic-temperature-controller.

[349] Thorlabs. 1" pillar posts. Online, accessed 16th August 2022. https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=241.

[350] ICE Oxford. Low temperature varnish. Online, accessed 17th June 2022. https://www.iceoxford.com/Low-Temperature-Varnish-ICEles03-.htm.

[351] Apiezon. Apiezon N grease. Online, accssed 17th June 2022. https://apiezon.com/products/vacuum-greases/apiezon-n-grease/.

[352] ICE Oxford. Eccobond 286. Online, accessed 20th June 2022. https://www.iceoxford.com/Epoxy-Resin-Adhesive-Suitable-for-Cold-Windows-ICEles47.htm.

[353] AimTTi. 1908 series. Online, accessed 16th August 2022. https://www.aimtti.com/product-category/multimeters/aim-1908.

[354] F. Travasso on behalf of the Virgo Collaboration. Status of the monolithic suspensions for Advanced Virgo. *Journal of Physics: Conference Series, 957, 012012*, 2018.

[355] J. McKenna. Private communication. Email, 12th October 2020.

[356] Henkel. Ablestik 286 twin pack. Online, accessed 24th June 2022. https://www.hitek-ltd.co.uk/_files/ugd/404d82_3b9eb05f187b4c648ae057363abacffa.pdf.

[357] A. Perin, R. M. Jareño, and L. Metral. Study of materials and adhesives for superconducting cable feedthroughs, LHC project report 504. Technical report, Large Hadron Collider Project, European Laboratory for Particle Physics, 2001.

[358] R. S. Kern et al. Cryogenic synopsis from the testing of the fully equipped ESS high β cavity ESS086-P01. Technical report, Department of Physics and Astronomy, Uppsala University, 2018.

[359] J. Polinski. Materials in cryogenics. In *European Course in Cryogenics 2010, CERN*, 2010.

[360] R. I. Amils et al. Thermal conductivity of silver loaded conductive epoxy from cryogenic to ambient temperature and its application for precision cryogenic noise measurements. *Cryogenics, 76, pp 23-28*, 2016.

[361] Sil'tronix. Sil'tronix. Online, accessed 27th September 2022. https://www.sil-tronix-st.com/en/home.

[362] Zygo. Zygo GPI XP-D interferometer. Online, accessed 25th July 2022. http://www.physics.gla.ac.uk/igr/index.php?L1=capabilities&L2=equipment&L3=zygo.

[363] K. Haughian et al. Mechanical loss of a hydroxide catalysis bond between sapphire substrates and its effect on the sensitivity of future gravitational wave detectors. *Physical Review D, 94, 082003*, 2016.

[364] L. Cunningham et al. Re-evaluation of the mechanical loss factor of hydroxide-catalysis bonds and its significance for the next generation of gravitational wave detectors. *Physics Letters A, 374, pp 3993-3998*, 2010.

[365] Sentech. SE850 spectroscopic ellipsometer. Online, accessed 25th July 2022. http://www.physics.gla.ac.uk/igr/index.php?L1=capabilities&L2=equipment&L3=ellipsometer.

[366] H. Armandula and M. van Veggel. Silicate bonding procedure (hydroxide-catalysis bonding), LIGO-E050228. Technical report, LIGO Scientific Collaboration, September 2010. Version 2.

[367] P. Duthil. Material properties at low temperature. *arXiV preprint, arXiV:150107100*, 2015.

[368] Thorlabs. Total transmission of 5 mm thick, uncoated sapphire windows. Online, accessed 27th July 2022. https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=3982&pn=WG30530.

[369] H. D. Young and R. A. Freedman. *University physics (13th edition) with modern physics*. Addison-Wesley, 2012.

[370] AimTTi. AIM 1908P multimeter data sheet. Online, accessed 16th August 2022. https://resources.aimtti.com/datasheets/AIM-1908_multimeter_data_sheet-Iss4.pdf.

[371] Lake Shore. Lake Shore Wire. Online, accessed 17th August 2022. https://www.lakeshore.com/docs/default-source/product-downloads/catalog/lstc_wire_l.pdf?sfvrsn=df94ba65_1.

[372] Eletronics Notes. Wire wound resistor: wirewound power resistor. Online, accessed 18th August 2022. https://www.electronics-notes.com/articles/electronic_components/resistors/wirewound-resistor.php.

[373] R. Patterson, A. Hammoud, and S. Gerber. Performance of various types of resistors at low temperatures. Technical report, NASA Glenn Research Center, 2001.

[374] J. E. Naefe, R. W. Johnson, and R. R. Grzybowski. High-temperature storage and thermal cycling studies of thick film and wirewound resistors. *IEEE Transactions on Components and Packaging Technologies, 25, pp 45-52*, 2002.

[375] V. Musilova, P. Hanzelka, T. Kralik, and A. Srnka. Low temperature radiative properties of materials used in cryogenics. *Cryogenics, 45, pp 529-536*, 2005.

[376] Y. S. Touloukian, R. W. Powell, C. Y. Ho., and P. G. Klemens. *Thermophysical properties of matter - the TPRC data series, volume 1, thermal conductivity - metallic elements and alloys*. New York, IFI/Plenum, 1970.

[377] L. M. Phinney, J. D. Kuppers, and R. C. Clemens. Thermal conductivity measurements of SUMMiT V polycrystalline silicon. Technical report, Sandia National Laboratories, 2006.

[378] S. Uma et al. Temperature-dependent thermal conductivity of undoped polycrystalline silicon layers. *International Journal of Thermophysics, 22 (2), pp 605-616*, 2001.

[379] S. Ju, X. Liang, and X. Xu. Out-of-plane thermal conductivity of polycrystalline silicon nanofilm by molecular dynamics simulation. *Journal of Applied Physics, 110 (5), 054318*, 2011.

[380] J. Yang. *Theory of thermal conductivity*. Springer, Boston, 2004.

[381] J. W. Sharp, S. J. Poon, and H. J. Goldsmis. Boundary scattering and the thermoelectric figure of merit. *Physica Status Solidi (a), 187 (2), pp 507-516*, 2001.

[382] A. V. Inyushkin et al. Ultrahigh thermal conductivity of isotopically enriched silicon. *Journal of Applied Physics, 123, 095112*, 2018.

[383] A. K. McCurdy, H. J. Maris, and C. Elbaum. Anisotropic heat conduction in cubic crystals in the boundary scattering regime. *Physical Review B, 2 (10), pp 4077-4083*, 1970.

[384] H. A. Vodden and J. L. Williams. Pure silicon. In *Proceedings of the Institution of Mechanical Engineers*, 1965.

[385] T. K. Liu. Lecture 4: Properties of carriers in semiconductors; carrier drift; conductivity and resistivity. Online, accessed 31st August 2022. https://inst.eecs.berkeley.edu/~ee130/sp13/lectures/Lecture4.ppt.

[386] O. C. Allkofer et al. The absolute cosmic ray flux at sea level. *Journal of Physics G: Nuclear and Particle Physics, 1 (6), L51-52*, 1975.

[387] K. Yamamoto et al. Effect of energy deposited by cosmic-ray particles on interferometric gravitational wave detectors. *Physical Review D, 78, 022004*, 2008.

[388] Aerospace Specifications Metals Inc. Aluminum 6061-T6. Online, accessed 8th September 2022. https://asm.matweb.com/search/SpecificMaterial.asp?bassnum=ma6061t6.

[389] K. Q. Chang. Calculation of eddy current loss of a disc. Master's thesis, Faculty of the Graduate School, Universiity of Missouri-Rolla, 1971.

[390] R. Kumar. *Aspects of suspension design for the development of advanced gravitational wave detectors*. PhD thesis, School of Physics and Astronomy, University of Glasgow, 2013.

[391] L. Cunningham, C. Torrie, and A. Cumming. Development of an FEA model to investigate the thermal noise in the bond layer of the ETM TM for AdLIGO thermal noise simulation, LIGO-T080163. Technical report, Institute for Gravitational Research, University of Glasgow, 2010.

[392] ANSYS, Inc. ANSYS mechanical APDL command reference. Technical report, ANSYS, Inc., 2010.

[393] R. P. Middlemiss et al. A MEMS gravimeter with multi-axis gravitational sensitivity. *arXiV preprint, arXiv:2102.05528*, February 2021.

[394] N. W. McLachlan. *Theory of vibrations*. Dover Publications, 1951.

[395] D. J. McLachlan and L. K. Chamberlain. Atomic vibrations and melting point in metals. *Acta Metallurgica, 12 (5), pp 571-576*, 1964.

[396] G. Cagnoli et al. Damping dilution factor for a pendulum in an interferometric gravitational waves detector. *Physics Letters A, 272, pp 39-45*, 2000.

[397] E. Elliffe. *Aspects of thermal noise modeling in ground-based gravitational wave detectors and developments of hydroxide catalysis bonding for space-based gravitational wave detectors and other optical applications*. PhD thesis, School of Physics and Astronomy, University of Glasgow, 2005.

[398] ANSYS, Inc. ANSYS modeling and meshing guide. Technical report, ANSYS, Inc., 2004.

[399] Intel. 3rd gen Intel Xeon scalable processors. Online, accessed 12th August 2022. https://www.intel.co.uk/content/www/uk/en/products/docs/processors/xeon/3rd-gen-xeon-scalable-processors-brief.html.

[400] Intel. 6th gen Intel Core desktop processor. Online, accessed 12th August 2022. https://www.intel.co.uk/content/www/uk/en/products/platforms/details/skylake-s.html.

[401] M. Phelps et al. Measurements of the Young's modulus of hydroxide catalysis bonds, and the effect on thermal noise in ground-based gravitational wave detectors. *Physical Review D, 97, 102004*, 2018.

[402] ANSYS, Inc. ANSYS contact technology guide. Technical report, 2004.

[403] J. Cheah. FEA stress singularities. Online, accessed 12th August 2022. https://www.ansystips.com/2017/05/fea-stress-singularities.html.

[404] K. Haughian. Private communication. Email, 25th March 2022.

[405] P. Fritschel and D. Coyne. Advanced LIGO systems design, LIGO-T010075. Technical report, LIGO Scientific Collaboration, 2015.

[406] C. A. Cantley et al. Monolithic stage conceptual design for Advanced LIGO ETM/ITM, LIGO-T050215. Technical report, LIGO Scientific Collaboration, 2006.

[407] M. van Veggel et al. Final design document ETM/ITM ears, LIGO-T0900447. Technical report, LIGO Scientific Collaboration, 2014.

[408] K. A. Strain. Note on revision of monolithic quad pendulum "d" parameters to enable revised ears and welding, LIGO-T0900556. Technical report, LIGO Scientific Collaboration, 2008.

[409] B. Sorazu, K. A. Strain, I. S. Heng, and R. Kumar. Violin mode amplitude glitch monitor for the presence of excess noise on the monolithic silica suspensions of GEO 600. *Classical and Quantum Gravity, 27 (15), 155017*, 2010.

[410] ANSYS, Inc. Private communication. Meetings and emails.

[411] March May. March May AC-3P-MD magnetically driven, single-stage centrifugal chemical fluid pump. Online, accessed 6th April 2022. https://marchmay.co.uk/pumps/ac-range/ac-3p-md/.