

Monolithically Integrated Membrane-In-The-Middle Cavity Optomechanical Systems

by

Graham Justin Hornig

A thesis submitted in partial fulfillment of the requirements for the degree of

Master of Science

in

Photonics and Plasmas

Department of Electrical and Computer Engineering

University of Alberta

© Graham Justin Hornig, 2020

Abstract

This thesis describes the theory, fabrication, and characterization of on-chip, integrated Fabry-Pérot microcavities embedded with thin silicon nitride membranes. The approach combines a previously developed process for fabricating self-assembled Si / SiO₂ / Ta₂O₅ curved mirror cavities, with a XeF₂-based sacrificial etch process to allow for free standing PECVD silicon nitride membranes to be released internally. Membranes with a fundamental mechanical resonant frequency ~ 10 MHz were coupled to wavelength-scale cavities operating in the ~ 1550 nm wavelength range, and the optical and mechanical properties of the combined system were studied. The estimated optomechanical coupling coefficient $g_0/2\pi \sim 0.1$ MHz is significantly higher than that previously reported in membrane-in-the-middle systems. In further contrast to previous approaches, the work here presents a scalable and monolithic surface-micromachining process for systems that are otherwise typically large and complex.

These devices are also described within the context of existing cavity optomechanics literature. The small size and high areal chip density of these devices may allow them to serve as potential candidates in a variety of applications such as sensing, frequency transduction, and quantum information.

Preface

This thesis contains the description of devices that are the product of a collaborative effort for process development among the research group of Professor Ray DeCorby. The theory of operation and design of the optomechanical system was a joint effort between Dr. DeCorby and myself. The process development, particularly the etch hole implementation described in Chapter 3, was a joint effort between James Maldaner, Sanaa Al-Sumaidae, and myself. Appendix D, in particular the description of the sputtering SOP, is a result of past group research and nanoFab staff. The material dispersion models in Appendix F are adapted from measurements taken by Timothy Harrison, Lintong Bu, and Kyle Scheuer.

Chapter 3 of this thesis has in portion been submitted for publication to the journal Optics Express as of August 6th, 2020 under G. J. Hornig, S. Al-Sumaidae, J. Maldaner, L. Bu and R. G. DeCorby, “Monolithically integrated membrane-in-the-middle cavity optomechanical systems.” The manuscript has been accepted and is set to be published in 2020.

Acknowledgements

I would like to first and foremost thank my supervisor Dr. Ray DeCorby for your continued mentorship throughout this project. Your patience and support has allowed me to grow immensely as both an engineer and researcher. Thank you for making my time here an enjoyable and interesting experience.

The work here would not be possible without the collaborative effort of my colleagues. A big thank you to: Tim Harrison for all of your help with basically everything in the lab and the fab, and for being one of the most reliable people I know; Sanaa Al-Sumaidae for always lending a hand without hesitation; James Maldaner for your affinity in data analysis; Danny Pulikkaseril for the great coffee and conversation; Seyed Azmayesh-Fard for your creative ideas and life advice; Lintong Bu for your curiosity and attention to detail; and Kyle Scheuer for your electric skills on both circuits and guitar. I also want to extend additional thanks to Mohammad Bitarafan for his assistance early on in my masters, and to the entire nanoFab staff whose process advice and diligent care of the facility make the nanoFab an invaluable part of research at the university.

Next, I want to thank my parents Harold and Linda Hornig, and my brother Alex Hornig for all of their endless love and encouragement. None of what I have accomplished in my life would have been possible without you. To all my friends over the years, I want to also thank you for all you've done to enrich my life and make every day worth living.

Finally, I want to acknowledge the financial support of the Natural Sciences and Engineering Research Council of Canada (NSERC,) Alberta Innovates, and the University of Alberta.

Contents

1	Introduction	1
1.1	Cavity Optomechanics	2
1.2	Membrane-In-The-Middle Systems	4
1.3	Buckled Dome Microcavities	8
1.4	System Integration	9
1.5	Outline of Thesis	11
2	Theory	12
2.1	Optics	12
2.1.1	The Bragg mirror	12
2.1.2	The Fabry-Pérot Cavity	15
2.2	Mechanics	24
2.2.1	Buckled Dome Mechanics	24
2.2.2	Mechanics of a Membrane	27
2.3	Optomechanics	30
2.3.1	Cavity Optomechanical Systems	35
2.3.2	Optomechanical Cooling	37
3	Monolithically integrated membrane-in-the-middle cavity optomechanical systems	40
3.1	Fabrication	41
3.1.1	Sputtering	42
3.1.2	PECVD Silicon Nitride	43
3.1.3	Low Adhesion Layer	45
3.1.4	Hotplate Buckling	47
3.1.5	Cobra Etching	48
3.1.6	XeF ₂ Sacrificial Layer Etching	50
3.2	Process Development	55
3.2.1	Free-Standing Membranes	55
3.2.2	First Generation: Si/SiO ₂ Based Top Mirror	56
3.2.3	Second Generation: Ta ₂ O ₅ /SiO ₂ Based Top Mirror	62
3.3	Characterization	65
3.3.1	Simulations	67
3.3.2	Experimental Setup	69
3.3.3	Experimental Measurements	72
3.3.4	Discussion	79
4	Conclusions and Future Work	82
4.1	Summary	82
4.2	Future Work	83
	References	85

Appendix A Wafer Survey	94
Appendix B Extended Theory	97
B.1 Optical Loss Expressions	97
B.2 Optical Transfer Matrices	97
Appendix C Norcada Inc. Membrane Development	103
C.1 Process Design	103
C.2 Fabrication	108
Appendix D Further Process Details And Operating Procedures	116
D.1 Process Details	116
D.1.1 Photolithography	116
D.1.2 Sputtering	120
D.1.3 Piranha Cleaning	122
D.1.4 Wet Deck Processing: Lift off and Mask Removal . . .	122
D.1.5 Zygo Optical Profilometer	123
D.1.6 Confocal Laser Scanning Microscope (CLSM)	123
D.1.7 Alphastep Step Profilometer	124
D.1.8 Flexus	125
D.2 Standard Operating Procedures (SOP)	127
D.2.1 Doug	127
D.2.2 Coupling Station	131
Appendix E Mask Details	133
Appendix F Material Dispersion Models	135
Appendix G Thin Film Simulator GUI	139

List of Tables

2.1	Optical constants for select thin film materials at 1550 nm. . .	14
2.2	Mechanical Properties of all relevant materials for use in the mechanical analysis.	26
3.1	Recipes used in Sputtering System 2, ‘Doug’.	45
3.2	Recipes used in the Oxford Plasmapro 100 Cobra Metal Etch system.	49
3.3	Approximate etch rates for various materials, as tested in the Cobra system.	50
3.4	A survey of the various possible methods for removing a sacrificial layer within the domes.	51
3.5	XeF ₂ etch process default recipe parameters.	54
3.6	Measured throughputs of the experimental setup, with no sample in the path.	71
3.7	Material properties of silicon nitride used in the membrane simulation.	75
C.1	Physical properties of the various membranes provided by Norcada.	105
D.1	Photoresist used in the fabrication process, with their corresponding recipes.	118
E.1	Geometry configurations of the dies used in this mask.	134
E.2	Hole configurations present on each die.	134
F.1	Fitting parameters for the dispersion models of relevant thin film materials.	136

List of Figures

1.1	Schematics of popular cavity optomechanical configurations. . .	2
1.2	Prominent optomechanical systems from literature exhibiting quantum behaviour.	5
1.3	Various Membrane-based Optomechanical systems from the literature.	6
1.4	Developments in non-uniform membrane geometries from the literature.	7
1.5	A cartoon of the delamination buckling process.	8
1.6	An overview of the representative devices produced by the buckling process.	10
1.7	An illustration of the proposed integrated device.	11
2.1	An illustration of a thin-film Bragg mirror on a substrate. . .	13
2.2	Transfer matrix simulations of two typical mirrors.	15
2.3	An illustration of a simple Fabry-Pérot cavity.	16
2.4	A plot of the transmission of the simple Fabry-Pérot.	17
2.5	An illustration of a half-symmetric cavity.	20
2.6	Plots of the first 25 Hermite-Gaussian modes.	21
2.7	Plots of the first 25 natural Laguerre-Gaussian modes.	23
2.8	Plots of the first 25 linearly polarized Laguerre-Gaussian modes.	23
2.9	Predicted vibrational mode frequencies of the top mirror using both the shell and plate model.	27
2.10	Plots of the eigenmodes of a circular membrane.	29
2.11	A plot illustrating the two-dimensional overlap between the optical and mechanical modes.	31
2.12	A plot of a Lorentzian function, and the location of the maximum slope.	32
2.13	The maximum derivative of the cavity transmission as linewidth is changed.	33
2.14	An illustration of a simple optomechanical cavity model to describe optical feedback cooling.	38
2.15	Diagrams of the work done by the membranes motion if the light in the cavity responds both instantaneously and retarded with respect to the cavities resonance.	39
3.1	A schematic overview of the fabrication process.	41
3.2	A cartoon illustration of the sputtering process.	43
3.3	A schematic of the Doug sputtering system.	44
3.4	A schematic of the nanoFab XeF ₂ etching system.	54
3.5	Microscope images of various stand-alone membranes after XeF ₂ etching trials.	56
3.6	Microscope images of three side-by-side membranes observed with different membrane integrity.	57

3.7	A schematic demonstrating a possible explanation for the concentric stepping visible on the membranes.	57
3.8	Microscope images of select individual membranes with the mask applied, prior to Cobra etching.	58
3.9	An overview of results from the XeF ₂ etching trials with and without an additional pre-XeF ₂ Bosch pulse.	60
3.10	Four domes from the XeF ₂ 20 pulsed no-Bosch sample.	61
3.11	Images from SEM and HiM scans on the devices, showing flakes near the etch holes.	62
3.12	A microscope and corresponding schematic providing an explanation for the multilayer XeF ₂ etching observed on the devices.	63
3.13	Microscope images of various membranes, tracked at different stages of the etching process.	64
3.14	Microscope images of various membranes, with a calibrated etch time chosen.	65
3.15	Microscope images of curious line patterns visible on the devices between the etch holes.	66
3.16	A schematic cross section and microscope image of a finished device.	66
3.17	Confocal microscope images and scans of the final device, in addition to a plot and fit of the buckled dome profile to theory.	67
3.18	Transfer-matrix predictions for the planar-equivalent model of the cavity.	69
3.19	A photograph of the experimental setup used to characterize the devices.	70
3.20	A schematic block diagram of the experimental setup.	71
3.21	Plots of the measured spectral transmission of the final device, along with a similar sample but with no etch holes or released membrane.	73
3.22	Measured FFT spectra for the finished devices.	75
3.23	A plot of the measured PSD spectrum over the simulation predicted mechanical mode frequencies, and a thermomechanical calibration for the fundamental mode.	77
A.1	An overview of various different types of devices visible on the wafer.	96
B.1	A diagram of the electric fields for a general transfer matrix problem.	98
C.1	A schematic of the proposed process for integrating the mirror dome cavity technology with Norcada's high quality mechanical membranes.	104
C.2	Microscope images of examples of typical poor liftoff of PTFE found during processing.	106
C.3	A photograph of the host wafer with taped on devices, after the photoresist was spun on.	109
C.4	Micrope images of successfully buckled membrane devices.	110
C.5	Zygo data from a successfully buckled membrane device.	110
C.6	Zygo data from another buckled membrane device, with a large degree of over buckling.	111
C.7	Microscope images of various membrane devices, with patterned PTFE after liftoff.	112
C.8	Microscope images showing the membrane devices after a 4 period 1550 nm Si/SiO ₂ mirror was deposited on top of them.	113

C.9	Microscope images showing the membrane devices after buckling.	113
C.10	Zygo data from imaging the reverse side of two buckled membrane samples.	114
D.1	A schematic demonstrating how a mask is used for the patterning of a material through liftoff.	117
D.2	A photograph of a Si sputtering target that has been used too long, and broken.	121
E.1	The layout of a dome with etch holes on the mask. Labelled dimensions correspond to Table E.1.	134
F.1	Dispersion model of sputtered thin-film a-Si.	137
F.2	Dispersion model of sputtered thin-film SiO_2	137
F.3	Dispersion model of sputtered thin-film Ta_2O_5	138
F.4	Dispersion model of PECVD thin-film Si_3N_4	138
G.1	A screenshot of the MATLAB Thin Film Simulator GUI. . . .	140

Chapter 1

Introduction

When light is captured and confined in the space between two mirrors, it can resonate and have a profound impact on enclosed matter such as confined atoms or mechanical objects. In the case of a mechanical object that is free to move, the momentum intrinsic to the resonant light can be engineered to apply appreciable forces with high precision. The study of this interaction between light and solid mechanical objects, a field known as ‘optomechanics’, has enabled experiments of extreme sensitivity. Optomechanical systems have demonstrated notable feats including the recently heralded detection of gravitational waves [1], and the placement of mechanical objects into a quantum ground state [2]. Light can be strongly coupled to a variety of different mechanical systems, without the level of noise that typically plagues alternative electrical measurement schemes. Advancements in this technology in recent years has opened the door to allow for optomechanical systems to have a significant technological impact on numerous fields of research.

In this work, a unique class of devices is presented that adapt a highly versatile optomechanical scheme of much recent technological interest [3] and place it on-chip through a process that is integrated, straightforward to fabricate, and scalable. The developments described here might pave the way for moving these experiments out of the laboratory, and into future systems that could require millions of these devices on a single wafer. This is made possible through past developments of integrated, high quality optical cavities by the DeCorby group, and this thesis aims to functionalize these cavities by coupling




Figure 1.1: Schematics of popular cavity optomechanical configurations, adapted from Ref. [4]. (a) A canonical Fabry-Pérot cavity keeping one mirror fixed and the other free to move, as a simple harmonic oscillator (modelled as a mass on a spring.) (b) A membrane-in-the-middle configuration, where both mirrors are fixed and instead the membrane position dispersively modifies the cavity resonance. (c) A microtoroidal resonator, where the mechanical element is physically part of the cavity itself. (d) A photonic crystal beam resonator, where the beam forms both the cavity, mirrors, and mechanical element.

a further integrated, high-quality mechanical element.

1.1 Cavity Optomechanics

As already briefly mentioned, cavity optomechanical systems promise to provide significant advances in a variety of fields. Specifically, the coupling to mechanical objects allows for these systems to be of immense value in precise measurements of force, mass, and acceleration. Accelerometers with high bandwidths and sensitivities that approach quantum limits have been created [5] which, for example, could enable the engineering of high precision navigation systems [6]. Indirectly, force measurements are sometimes used as an intermediary to measure other phenomena of interest, including magnetic fields, the precise measurement of which is critical for medical and other applications [6, 7]. Mechanical elements have also been demonstrated to effectively

serve as platforms for processing light. This includes switches and delay lines (through optomechanically induced-transparency) [8], pulse storage [9], and frequency conversion [10]. Frequency (wavelength) conversion in particular is of much interest, in that optomechanical systems offer significantly more flexibility than traditional conversion with non-linear materials [6]. Conversion between different wavelengths (*e.g.* from visible to NIR) is a requirement of many telecommunication systems. Mechanical elements can facilitate conversion between vastly differing frequencies, such as transferring information between optical and microwave photons [11], and in some cases even radio waves [12]. This is of great interest in upcoming quantum information technologies [11, 13, 14].

Optomechanical systems come in many forms. Perhaps the most basic, canonical arrangement is simply allowing one of the end mirrors in an optical cavity to be free to move as a simple harmonic oscillator (Fig. 1.1(a)). Changes in the mirror position directly impact the cavity resonance frequency by modulating the cavity length. Difficulties arise however in simultaneously creating high reflectivity mirrors that are also good mechanical resonators. This can be overcome by decoupling the motion from the end mirror, as in the membrane-in-the-middle (MIM) configuration shown in Fig. 1.1(b). In a MIM system, a thin membrane is suspended inside of the optical cavity, and both end mirrors are held fixed. Membranes of very high mechanical quality have been fabricated and have demonstrated a large dispersive effect on the cavity (similar to allowing one of the end mirrors to move,) while also having very low optical absorption and scattering losses [3]. There are a variety of other approaches in the literature that have achieved much success in particular when placed within integrated, on-chip systems. These include microtoroids (Fig. 1.1(c)) in which light is coupled evanescently through a nearby fibre into a ring shaped cavity. Mechanical modes are within the ‘breathing’ of the central mechanical element. Photonic crystal beam resonators take a similar approach of making the cavity one with the mechanical element (Fig. 1.1(d),) and have been used successively to demonstrate impressive quantum properties [2, 14, 15].

One reason optomechanical systems are so coveted is their potential for

experimental observations of the quantum ground state. Through using a mechanism known as ‘sideband cooling’, the light sent into the cavity can be strategically detuned causing a feedback effect between the optical and mechanical components of the system. The force applied by the light can damp the intrinsic thermal motion of the mechanical resonator, effectively cooling it to temperatures below that which could be achieved with a cryogenic cooling system. The mechanical resonator can be nearly ‘stopped’, meaning the time-averaged vibrational energy contained within it is sufficiently low that it is forced into its ground state - the lowest possible energy as predicted by quantum mechanics. One of the first experimental demonstrations of ground-state cooling (Fig. 1.2(a)) employed an aluminum drum resonator in a microwave cavity [16]. In recent years, this has also been extended to optical systems such as in Ref. [17] (Fig. 1.2(b)) with a microtoroidal resonator, and in Ref. [2] (Fig. 1.2(c)) through a phononic-crystal-isolated photonic-crystal beam resonator. These systems are interesting on their own - the implications of the placement of relatively large objects in quantum states is not fully understood. From a practical standpoint however, it is their potential as platforms for manipulating and studying single quanta of information that is particularly compelling. Placement of these objects in a controlled quantum state also allows for fundamental studies of the very nature of quantum mechanics [18, 19]. Careful consideration of the quantum properties of these systems can also enable the generation of so-called ‘squeezed states’ of light [20]. By managing uncertainty trade-off’s, these squeezed states can subvert conventional limitations and thus be used to make measurements of high precision that are otherwise classically impossible [21].

1.2 Membrane-In-The-Middle Systems

One optomechanical platform of particular interest is the so-called ‘membrane-in-the-middle’ (MIM) system (Fig. 1.1(b)), which uses a thin membrane placed between the cavity mirrors as a mechanical element. The first demonstrations of a MIM system with dispersive optomechanical coupling was reported in




Figure 1.2: Prominent optomechanical systems from literature exhibiting quantum behaviour. (a) [16] Aluminum membrane placed into a quantum ground state through sideband cooling with a microwave system, (b) [17] Quantum coherent coupling between a microtoroidal mechanical mode and an optical mode, (c) [2] Patterned beam used as a platform for quantum behaviour, facilitated by surrounding phononic crystal structures to isolate the mechanical modes from the environment. (Images are reproduced with permission from their respective sources.)

2008 by the Harris group at Yale university [3]. A 50 nm thick, 1 mm wide square membrane was placed within a Fabry-Pérot optical cavity formed between two fixed, macroscopic mirrors. As the membrane moves within the cavity it dispersively detunes the cavity resonance. Judicious placement of the membrane within the cavity can result in either a linear or quadratic coupling of the cavity optical resonance to the membrane position; quadratic coupling is of particular interest for quantum non-demolition measurements [22, 23, 24]. In 2014 MIM devices were demonstrated as an effective, integrated system with the potential for transduction of coherent quantum states between microwave and optical frequencies by the Lehnert group at the University of Colorado [13]. By patterning the periphery of the membrane with a metal (Fig. 1.3(c)) the membrane forms one end of a capacitor within a microwave-frequency LC resonator circuit. Displacement of the membrane cause changes in capacitance, and thus modulates the resonance of the microwave circuit. Some more recent developments in MIM systems have included work on so-called ‘MATE’ (membrane-at-the-edge) systems, which through bending of a membrane frame allows for it to be placed closer than otherwise possible to one of the end mirrors [25]. These systems, while still utilizing a relatively large cavity length (~ 10 cm), allow for a higher degree of optomechanical




Figure 1.3: Various Membrane-based Optomechanical systems from the literature. (a) [26] Fibre based MIM system, with a notably small cavity length ($\sim 80 \mu\text{m}$) (b) [27] A ‘compact’ MIM experimental system, showing their relatively large size and the difficulty of true integration. (c) [13] Experimental realization of transduction between optical and microwave domains with the use of a membrane as an interface. (Images are reproduced with permission from their respective sources.)

coupling than is possible when the membrane is placed closer to the centre. However, the highest degree of optomechanical coupling is ultimately achieved by a smaller, wavelength-scale cavity, which remains until now an outstanding goal in the field of MIM style systems [25].

There is interest in making these systems smaller and more integrated, in order to compete with other quantum-enabled optomechanical systems. Reduced size means a reduced mass, thus potentially making them easier to cool to reduce thermal noise. In addition, a smaller cavity length allows for cavity operation at a lower optical mode order and affords an increase in optomechanical coupling and thus system efficiency. Systems on the order of centimetres are typical (Fig. 1.3(b)) for most ‘macroscopic’ experiments, such as in Ref. [27]. Demonstrations of small cavity sizes, approximately 80 microns, have been reported with fibre based systems ([26], Fig. 1.3(a)). However, further size reduction with similar hybrid-integrated systems would require the membrane to be replaced with an even smaller object, such as a nanowire [28]. The




Figure 1.4: Developments in non-uniform membrane geometries. (a) [29, 30] A ‘trampoline’ style resonator consisting of a small suspended central pad. Room temperature quantum operation may be possible with such a geometry owing to careful design of the tether structure. (b) [31] A phononic crystal patterned membrane, also utilizing a similar tether design as in a, which can enable drastic improvements in mechanical quality factor. (c) [32] A perforated integrated membrane, patterned with holes (via electron-beam lithography) to allow for sacrificial layer release. (Images are reproduced with permission from their respective sources.)

size and complexity of these systems also limits their future scalability (*i.e.* to large numbers).

Some of the practical issues associated with MIM systems may be overcome through optimization of the membrane itself. Recently, ‘trampoline’ style membranes have been developed [29, 30] (Fig. 1.4(a)) in which a central pad is suspended by tethers to the surrounding frame. The low mass of the membrane, combined with stress-engineering of the tether structure has enabled extremely high mechanical quality factors at room-temperature, making coveted room temperature quantum operation a possibility. Similarly, the mechanical quality can also be increased through the patterning of phononic crystal structures within the membrane surface [31] (Fig. 1.4(b)). Phononic crystals allow for vibrations in the membrane to be tightly controlled, and can drastically increase the mechanical quality factor over that which is possible with a simple uniform membrane. Even with these improvements, most membranes are fabricated separately from their accompanying optical cavities and inserted later during the experiment. Developments in on-chip membranes with an in-place release process could lead to advances in the integration and




Figure 1.5: A cartoon of the buckling process. Compressive stress in the top mirror causes it to buckle in areas where adhesion is poor. This can be controlled by patterning a thin low adhesion layer (~ 20 nm thick) prior to sputtering of the top mirror.

scaling of this technology. As an example from the MEMS literature, Figure 1.4(c) shows a membrane released through holes patterned on its top surface by removing an enclosed sacrificial layer below. Despite the holes, and even slight etching of the membrane itself, a high mechanical quality was still achieved [32].

1.3 Buckled Dome Microcavities

The work in this thesis is built upon technology developed over the past decade by the DeCorby group. Through controlled thin-film delamination, a process has been developed to allow for the unique ability to create on-chip, monolithic, high quality Fabry-Pérot cavities [33]. A highly reflective Bragg mirror, or quarter wave stack (QWS), is deposited onto a substrate as shown in Figure 1.5. A patterned layer of a low adhesion material similar to PTFE follows. Another QWS is then sputtered on top of this layer, such that it is under high compressive stress. The low adhesion layer causes the top mirror to have a weaker bond to select areas of the substrate, where the low adhesion layer is located. Due its intrinsic compressive stress, upon heating the top mirror will selectively ‘buckle’ along the patterns of low adhesion. The result is a highly symmetric, and smooth curved mirror structure with an enclosed internal cavity.

Depending on the patterning of the low adhesion layer, the buckling mir-

rors can be formed into different types of three dimensional optical structures. With a circular pattern, delamination causes a circularly symmetrical dome-like structure to form (Fig. 1.6(a)). The height of the final cavity, and thus it's optical resonance, is closely correlated to the diameter of the original circle patterned into the low adhesion layer. This process has been shown to produce high quality half-symmetric Fabry-Pérot cavities [34]. Alternatively, elongated rectangular shapes can be patterned that allow for the formation of waveguide structures (Fig. 1.6(b)), that can be coupled to from the side of the substrate [35]. Combining these dome and waveguide designs allows for the creation of so-called channel-connected domes (Fig. 1.6(c)) that allow for light to be transferred from one dome to the next within the chip itself [36]. This process can in principle be used to create intricate high density networks of functional structures. QWS mirrors may be designed and constructed to operate at an arbitrary wavelength, limited only by the availability of appropriate materials. This process is also completely monolithic, meaning that all processing and construction is done on a single substrate using standard clean-room processes such as lithography and thin-film deposition. Many other methods exist for the creation of similarly integrated on-chip waveguides and cavities, but typically involve some form of bonding and alignment to a second substrate [37]. Wafer bonding makes processing more difficult, and is undesirable from a mass production perspective.

The dome microcavities present themselves as excellent platforms for functionalized Fabry-Pérot devices. Work is being done on the introduction of dipole emitters into the cavities through potentially fluidic suspensions (*e.g.* PbSe quantum dots,) gas-phase emitters (*e.g.* optically trapped Rb atoms) and/or solid state embedded dopants (*e.g.* Er doped SiO₂ films) [38, 39].

1.4 System Integration

The goal of the work in this thesis is to implement the first monolithically integrated membrane-in-the-middle cavity optomechanical system. The advancement of any optomechanical system towards practicality will require increased




Figure 1.6: An overview of the representative devices produced by the buckling process. (a) [34] Domes with a high degree of cylindrical symmetry. (b) [35] Tapered Waveguides with a wavelength dependent cut-off position, that can be used for spectral interrogation. (c) [36] Channel-connected domes, that allow for devices to be interconnected.

robustness and scalability, which has not been a hallmark of most MIM experiments to date. While, as shown, some efforts have been made to reduce the size of these systems, an on-chip fully integrated system has not yet been demonstrated.

The size reduction in the membrane also provides further advantages for optomechanical applications. Smaller membranes have lower effective masses and higher resonance frequencies, giving them a larger zero-point motion. This can potentially result in a higher optomechanical coupling rate to the surrounding optical cavity. As demonstrated in recent advances in membrane technology [29, 30], this may have the potential to allow for the observation of quantum phenomena at room temperature. Small-area, high-frequency resonators promise to be the next step in the development of MIM technology, and there is no better place for these delicate devices than an integrated platform.

The demonstration of a MIM system with this potential for integration will provide an important milestone for MIM technology and could be of interest for the future of quantum transduction systems. Figure 1.7 is an artist's depiction of the completed integrated MIM system presented here. In this illustration, the membrane is enclosed within the Bragg mirror cavity. Holes have been etched, top down, through the mirror and membrane to allow for membrane release.




Figure 1.7: An illustration of the proposed integrated device, shown in a slice profile (not to scale.) A schematic beam (red) shows where the optical mode would sit inside of the cavity.

1.5 Outline of Thesis

This thesis begins with the introduction of background theory in Chapter 2. Principles in optics and mechanics are presented separately, and then brought together with a discussion of optomechanics. A presentation of the devices and fabrication processes follow in Chapter 3, which includes the characterization and measurements of a few devices indicative of what is possible with the process. A conclusion and summary follow in Chapter 4, along with a set of appendices that offer more detail on theory and fabrication.

Chapter 2

Theory

2.1 Optics

2.1.1 The Bragg mirror

One of the key technologies in controlling light is the mirror. There are many kinds of mirrors, but of particular relevance and use here is the Bragg mirror. A Bragg mirror is composed of alternating layers of two dielectric materials with different refractive indices. These dielectric materials, by themselves and in isolation, do not provide a significant reflectivity. However, through their careful design in a multilayer configuration, interference effects can be exploited to create extremely reflective mirrors at the designers wavelength of choice, and in some cases over a large range of angles. A Bragg mirror may also be referred to as a dielectric Bragg reflector (DBR,) or (assuming certain restrictions on layer thicknesses are met) a quarter wave stack (QWS.)

Designing such a mirror is fairly simple. First, two materials must be selected - one of a higher index and the other of a lower index. The larger the difference in refractive indices between the two materials, the more reflective a mirror for a set number of periods (in addition to a wider operating band around the design wavelength.) Next, a wavelength of design should be selected. In all work here, we have chosen to operate at the near-infrared wavelength of 1550 nm, which is a standard telecommunications wavelength. From the choice of the centre operating wavelength, the thickness of each layer is typically chosen to be that of a ‘quarter-wave.’ At this particular thickness,




Figure 2.1: An illustration of a thin film Bragg mirror on a substrate (typically silicon.)

the reflectance of the mirror is maximized for the chosen pair of materials. The thickness is found according to [40]:

$$t_q = \frac{\lambda_{\text{Bragg}}}{4n_q}. \quad (2.1)$$

Here, t_q is the thickness of the layer of material q , λ_{Bragg} is the design wavelength in free space, and n_q is the real part of the refractive index of material q . The two materials, with these calculated thicknesses, may now be layered and alternated with one another to form a periodic thin-film stack. Figure 2.1 shows an illustration of a typical Bragg mirror, with N periods stacked on top of a substrate. Each set of the two materials is referred to as a period. The peak reflectance for an N -period QWS mirror at the design wavelength is given by the expression [40]:

$$R_B = \left(\frac{1 - \left(\frac{n_{\text{out}}}{n_{\text{in}}}\right)\left(\frac{n_{\text{H}}}{n_{\text{L}}}\right)^{2N}}{1 + \left(\frac{n_{\text{out}}}{n_{\text{in}}}\right)\left(\frac{n_{\text{H}}}{n_{\text{L}}}\right)^{2N}} \right)^2. \quad (2.2)$$

Here, R_B is the peak reflectance of the Bragg mirror (at normal incidence and λ_{Bragg}); n_{out} , n_{in} , n_{H} , and n_{L} are the (assumed real) refractive indices of the output medium, input medium, high index material, and low index

	a-Si	SiO ₂	Ta ₂ O ₅	Si ₃ N ₄
n	3.47	1.46	2.08	2.0

Table 2.1: Optical constants for select thin film materials at 1550 nm. Full dispersion models are provided in Appendix F.

material respectively, and N is the number of periods. When considering ideal, lossless materials, mirror reflectance can be made arbitrarily closer to 1 through simply adding more periods to the QWS. For a large number of periods such that $R_B \rightarrow 1$, the range of wavelengths around λ_{Bragg} that is subject to high reflection (*i.e.* the stop band width, see Fig. 2.2) may also be predicted analytically through the following equation [40]:

$$\Delta\lambda_{\text{Bragg}} = \lambda_{\text{Bragg}} \left(\frac{4}{\pi} \right) \arcsin \left(\frac{|n_H - n_L|}{n_H + n_L} \right). \quad (2.3)$$

Here it should be noted that for a particular design wavelength, the width of the stop band depends on the difference between the refractive indices of the two materials. In practice, typical material combinations that are used in the 1550 nm range will include Si/SiO₂, and Ta₂O₅/SiO₂. They are easily sputtered as thin-films, have a low optical absorption, and have an index contrast that makes them highly reflective for a reasonable number of periods. Typical optical properties of these sputtered thin films at 1550 nm are given in Table 2.1, and full dispersion models (which are used when relevant in simulations) are provided in Appendix F. Constants for PECVD silicon nitride, discussed later, are also included.

The spectral response of two typical Bragg mirrors, used in the work described here, are simulated at normal incidence via transfer matrices (Fig. 2.2.) A higher index contrast, by using a-Si instead of Ta₂O₅ as the high index material, gives the advantage of a higher reflectance with less periods, and a larger width of the stop band. Nominal thicknesses here for a QWS layers are $d_{\text{Ta2O5}} = 186$ nm, $d_{\text{SiO2}} = 265$ nm, and $d_{\text{Si}} = 112$ nm. This gives total thicknesses for the 3 period Si/SiO₂ mirror and 7.5 period Ta₂O₅/SiO₂ mirror of ~ 1100 nm and ~ 3300 nm respectively.




Figure 2.2: Transfer matrix simulations of two typical mirrors. The lower contrast of the Ta_2O_5 based mirror (dashed) means more periods are required to obtain a similar reflectance as Si based mirror (solid) the at the design wavelength (1550 nm). The curves shown are for 7.5 and 3 period mirrors, respectively.

As light reflects off of the Bragg mirror, it penetrates a finite extent into the surface, typically further in comparison to the reflection from a metal mirror for example, which would be more of a hard boundary. Since the reflective properties of the Bragg mirror arise due to constructive interference of sub-reflections at the periodic interfaces of the layers inside, the true reflective boundary is at a ‘penetration depth’ into the mirror. This penetration depth may be approximated by [41, 42]:

$$L_p = \frac{\lambda_{\text{Bragg}}}{4\bar{n}(n_H - n_L)}, \quad (2.4)$$

where L_p is the penetration depth into the mirror, and $\bar{n} = \frac{n_H t_H + n_L t_L}{t_H + t_L}$ is the effective medium index of the mirror. For the two mirrors shown in Fig. 2.2, Equation 2.4 predicts penetration depths of 364 nm and 94 nm for the Ta_2O_5 and a-Si based mirrors respectively.

2.1.2 The Fabry-Pérot Cavity

When two mirrors with finite transmittance are placed close together (Fig. 2.3), they can be used to capture and confine light between them. As the trapped light bounces back and forth between the two mirrors, a pattern of interference occurs, causing resonance and potentially high local intensities of light. The wavelength of resonance is highly dependent on the mirror spacing.




Figure 2.3: A simple Fabry-Pérot cavity, where light is confined by repeated reflections between two mirrors and at resonant wavelengths can constructively interfere.

Suppose two mirrors are placed a distance L apart. A ray of light inside of the cavity incurs a phase after travelling a round trip back and forth, given by the expression [43]:

$$2\phi = 2k_z L, \quad (2.5)$$

where $k_z = 2\pi/\lambda_0$ is the component of the wavevector along the optical axis. In order for resonance to occur, the light must constructively interfere with itself after it reflects and travels around the cavity. This means that the round trip phase shift 2ϕ must obey $2\phi + 2\phi_R = m2\pi$ ($m = 1, 2, 3\dots$), where ϕ_R is an additional phase shift due to reflection from each mirror (*i.e.* where the mirrors are assumed to be identical here for simplicity). Temporarily neglecting ϕ_R , this condition can be rewritten in terms of wavelength:

$$\lambda_o = \frac{2Ln}{m} \cos \theta. \quad (2.6)$$

At normal incidence ($\theta = 0^\circ$) and for an air cavity ($n = 1$), this expression reduces to $\lambda_o = \frac{2L}{m}$. This simple expression shows that in order for a cavity to be resonant to light of a particular wavelength, its length must be roughly (after considering the non-zero mirror dependent ϕ_R *i.e.* mirror penetration) an integer multiple of a half wavelength, starting at the shortest possible cavity




Figure 2.4: The transmission of the simple Fabry-Pérot (Equation 2.8) with respect to the L/λ . Resonance occurs at wavelengths that are half-integer multiples of the cavity length. The plot shown is for a representative case with $R=0.9$ and assuming lossless, identical mirrors.

which is a half-wavelength long. The transmission and reflection of the cavity may be predicted through [43]:

$$R_E = \frac{4R \sin^2(\phi)}{(1 - R)^2 + 4R \sin^2(\phi)}; \quad (2.7)$$

$$T_E = \frac{T^2}{(1 - R)^2 + 4R \sin^2(\phi)}. \quad (2.8)$$

Equation 2.8 is shown in Fig. 2.4 plotted with respect to the length-to-wavelength ratio. When the cavity length is an integer multiple of the wavelength of the light, resonance occurs and the cavity will transmit light. Off resonance, the Fabry-Pérot will tend to reflect most of the light. The spectral sensitivity of the Fabry-Pérot, for mirrors of reflectance R , is quantified by the cavity reflection finesse \mathcal{F} [43]:

$$\mathcal{F} = \frac{\pi\sqrt{R}}{1 - R}. \quad (2.9)$$

Finesse is the average number of round trips that a resonant photon makes before it leaves the cavity through either transmission or absorption. It is directly related to the full-width half-maximum (FWHM) of the phase $\Delta\delta$ of the (fundamental) optical mode (in the limit of a large finesse) by [43] :

$$\Delta\phi \approx \frac{\mathcal{F}}{2\pi}. \quad (2.10)$$

It is typically desirable to have a very high finesse in a cavity, and doing so comes from engineering a balance between cavity reflectance and loss. The quality Q of the mode, or the ratio of the wavelength to FWHM, is closely related to the cavity finesse, but scaled by the mode order m [43]:

$$Q = m\mathcal{F}. \quad (2.11)$$

This implies that a significantly higher optical Q can be achieved if operating at a higher mode order. It follows however that larger cavity lengths will be required for a particular wavelength. Shorter cavities will have a larger mode spacing or ‘free spectral range’. The free spectral range λ_{fsr} of a Fabry-Pérot around the wavelength λ_0 is given by [43]:

$$\lambda_{\text{fsr}} = \frac{\lambda_0^2}{L}. \quad (2.12)$$

The free spectral range is the spacing in wavelength between adjacent longitudinal modes in the cavity. Ignoring material dispersion this spacing is uniform across all frequencies, but in wavelength it is not due to the inverse relationship between wavelength and frequency.

Reflectance, transmittance, and loss of a mirror are related through a conservation of energy relation, which states [40]:

$$T + R + A + S = 1, \quad (2.13)$$

where T is the transmittance of the mirror, R is the reflectance of the mirror, A is the absorptance of the mirror (inclusive of bulk scattering due to voids or other material inhomogeneities,) and S is the scattering of the mirror at all interfaces due to surface roughness (as per Equation B.3.) When considering maximizing the efficiency of a Fabry-Pérot, it may seem always best to use as many periods as possible, as to maximize mirror reflectance. However, real materials are optically lossy. Practical use of the Fabry-Pérot requires some light to transmit through at least one of the mirrors, and a mirror with many periods will be very thick and absorb a significant portion of the transmitted

light, even with low-loss materials. This balance between reflectance and the inevitable mirror absorption is quantified by the potential transmittance (PT):

$$PT = \frac{T}{1 - R} = \frac{T}{T + A + S}. \quad (2.14)$$

PT is an indicator of how much ‘potentially’ transmitted light is lost in a design due to absorption and scattering. It is desirable to have simultaneously highly reflective mirrors while also being able to have an appreciable amount of light to measure that leaves the mirrored cavity. In this context, the PT becomes a useful parameter for designing devices with these trade-off’s in mind.

Consider forming a Fabry-Pérot with multilayer Bragg mirrors at each end, instead of ideal perfect reflectors. With DBRs, the electric field penetrates a finite distance into each mirror effectively changing the cavity length and resonant wavelength from the hard reflection assumption of Equation 2.5. If mirrors of two different compositions are used, the stop band of the Fabry-Pérot will be roughly equivalent to the overlap of the two, with a transmitted resonant mode given by a modified Equation 2.5, letting $L \rightarrow L_{\text{eff}} = L + L_{\text{p},1} + L_{\text{p},2}$ to compensate for the penetration into mirrors 1 and 2. A corrected single-pass phase for the DBR-based cavity is given by [44] :

$$\phi = \frac{\pi}{\lambda}L + \pi \left(1 - \frac{\lambda_{\text{Bragg}}}{\lambda} \right) (L_{\text{p},\text{Si}} + L_{\text{p},\text{Ta}}), \quad (2.15)$$

which through rearrangement predicts the resonant wavelength for $\phi = 2\pi$:

$$\lambda = \frac{\pi L - \lambda_{\text{Bragg}}(L_{\text{p},\text{Si}} + L_{\text{p},\text{Ta}})}{2\pi - \pi(L_{\text{p},\text{Si}} + L_{\text{p},\text{Ta}})}. \quad (2.16)$$

Within a Fabry-Pérot, consideration needs to be made towards keeping light confined perpendicular to the optical axis. A cavity where both mirrors are flat is highly sensitive to misalignment, and extremely difficult to obtain a ‘stable’ resonance from. A common approach to overcoming this problem is to make one or both of the mirrors curved, so as to ‘redirect’ the light towards the central optical axis. Consider a so-called ‘half symmetric’ cavity (Fig. 2.5,) where one mirror is flat and the other is approximately spherical with a radius of curvature R_2 . The curved mirror allows for optical confinement, and




Figure 2.5: An illustration of a half-symmetric cavity, where one mirror is flat ($R = \infty$,) and the other is curved with a radius R_1 .

in doing so introduces a set of transverse modes, in addition to the expected longitudinal modes (as was predicted in the case of the 1-D Fabry-Pérot.) In the limit of $R_2 \gg L$, spherical mirror cavities will always support stable Laguerre-Gaussian (LG) and Hermite-Gaussian (HG) modes [43]. For both LG and HG modes, their lateral extent in the cavity is directly influenced by the mirror curvature, and may be described by the beam width $w(z)$ (Fig. 2.5):

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_0}\right)^2}, \quad (2.17)$$

where $w_0 = \sqrt{z_0 \lambda_0 / \pi}$ is the minimum beam waist size, and z_0 the Rayleigh range, given for the half-symmetric cavity as:

$$z_0 = (LR_2)^{1/2} \left(1 - \frac{L}{R_2}\right)^{1/2}. \quad (2.18)$$

The Rayleigh range is the distance along the optical axis from the position of the minimum beam waist size z_0 at which the cross-sectional area of the beam doubles [43]. Knowing the Rayleigh range also allows for a precise definition

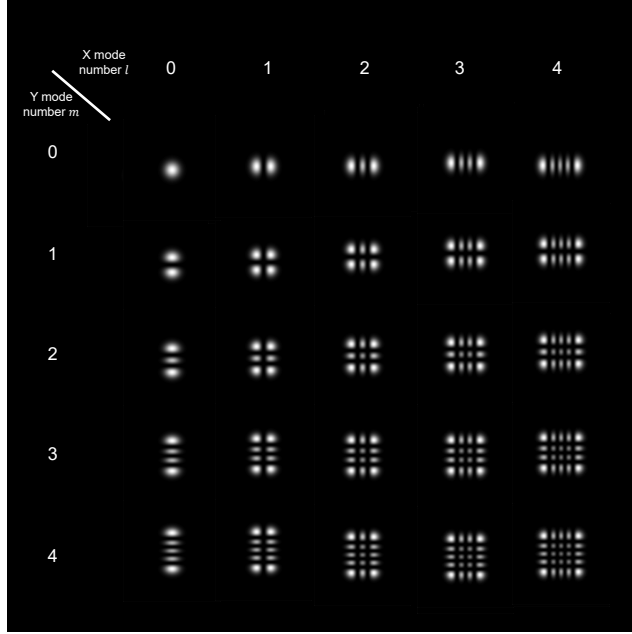


Figure 2.6: The first 25 Hermite-Gaussian modes, computed using Equation 2.21. Labels indicate the two indices as (l,m) .

of the radius of curvature of the beam wavefront along the optical axis $R(z)$:

$$R(z) = \frac{1}{z}(z^2 + z_0^2). \quad (2.19)$$

Both the LG and HG modes are independent and orthogonal sets of basis solutions to the paraxial wave equation [43]:

$$\frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial \Psi}{\partial r} \right) + 2ik \frac{\partial \Psi}{\partial z} = 0. \quad (2.20)$$

Here the equation is given in cylindrical coordinates $(\Psi(r, \theta, z))$. This equation provides an accurate description for both the electric ($\Psi = E$) and magnetic ($\Psi = H$) fields under the paraxial approximation, which assumes that most rays of interest within the mode field are nearly parallel to the optical axis and well described by a small-angle approximation. By solving Equation 2.20 in Cartesian coordinates for the electric field, an expression for the HG modes can be obtained:

$$E_{l,m}(x, y, z) = \left(E_0 \frac{w_0}{w(z)} \right) H_l \left(\frac{\sqrt{2} x}{w(z)} \right) H_m \left(\frac{\sqrt{2} y}{w(z)} \right) \exp \left(-\frac{x^2 + y^2}{w^2(z)} \right) \exp \left(-i \frac{k(x^2 + y^2)}{2R(z)} \right) \exp(-i\phi_z), \quad (2.21)$$

where E_0 is the electric field amplitude, and H_n represents the n^{th} -order Hermite polynomial. These modes are plotted in Fig. 2.6, for a variety of the mode indices l and m . Very similarly, the LG modes are obtained by instead solving Equation 2.20 in cylindrical coordinates and can be written:

$$E_{l,s}(r, \phi, z) = \left(E_0 \frac{w_0}{w(z)} \right) \left(\frac{\sqrt{2} r}{w(z)} \right)^{|l|} L_{l,s} \left(\frac{2r^2}{w^2(z)} \right) \exp \left(-\frac{r^2}{w^2(z)} \right) \exp \left(i \left(-\frac{kr^2}{2R(z)} - kz - \phi_z \pm s\phi \right) \right), \quad (2.22)$$

where l is the radial index and s is the azimuthal index. This equation represents two distinct orthogonal solutions of opposite-direction circular polarization, indicated by the \pm in the phase term. The ‘+’ mode is plotted in Fig. 2.7; however note that the ‘+’ and ‘-’ modes are indistinguishable in an intensity plot. In our experimental work, light used to excite these modes is always generated by a linearly polarized laser. In this case, the observed field can be represented as a linear combination of these two solutions:

$$E_{\text{LG,linear},l,s}(r, \phi, z) = E_{+,s}(r, \phi, z) \pm E_{-,s}(r, \phi, z), \quad (2.23)$$

which represents the two possible orthogonal linear polarizations. The intensity map of the ‘+’ solution to Equation 2.23 is shown in Fig. 2.8 for many modes.

While both LG and HG modes are valid solutions for the ideal cavity; in practice one will often observe some superposition of the two. The presence of LG modes is typically an indicator of a high degree of cylindrical symmetry. In the case of the TE_{00} mode, a simpler expression for the Gaussian beam within the cavity is sometimes also useful [45]:




Figure 2.7: The first 25 natural Laguerre-Gaussian modes, computed using Equation 2.22. Labels indicate the two indices as (l,s) .

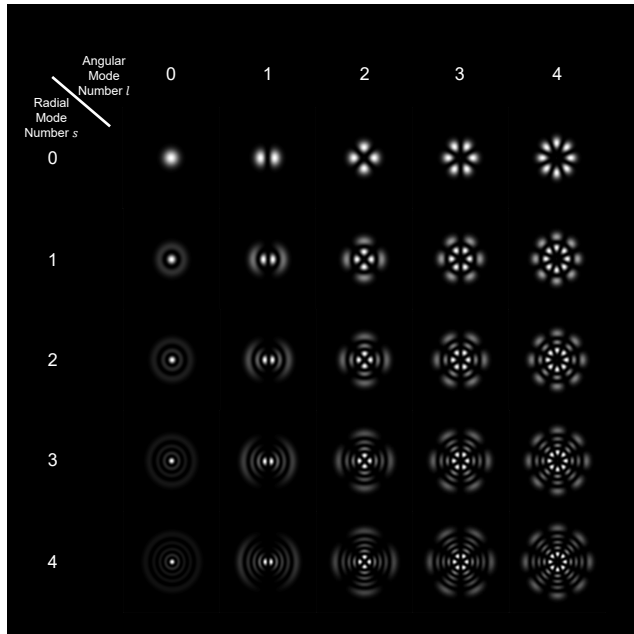


Figure 2.8: The first 25 linearly polarized Laguerre-Gaussian modes, formed from linear combination of the ‘natural’ modes, computed using Equation 2.22. Labels indicate the two indices as (l,s) .

$$\psi(x, y) = \frac{2}{\pi w_o^2} \exp\left(\frac{-2(x^2 + y^2)}{w_o^2}\right). \quad (2.24)$$

The round trip phase is now modified and unique for each mode:

$$\phi = kz - (l + m + p) \arctan\left(\frac{z}{z_0}\right). \quad (2.25)$$

In addition to the regular expected phase incurred through propagation, Equation 2.25 accounts for an additional ‘Gouy’ phase - a result of the beam’s transverse confinement. With their own phases and thus resonant wavelengths, families of transverse modes now begin to form for each longitudinal mode. For HG modes, these resonant frequencies can be predicted according to [46]:

$$f_{mpq} = \frac{c}{2L} \left(q + \frac{m + p + 1}{\pi}\right) \arccos\left(\sqrt{1 - \frac{L}{R_2}}\right). \quad (2.26)$$

In the limit of the mirror curvature $R_2 \gg L$, the spacing between any two HG modes is given by [46]:

$$\Delta f_{\text{HG}} = \frac{c}{2\pi z_0} \Delta(m + p). \quad (2.27)$$

2.2 Mechanics

2.2.1 Buckled Dome Mechanics

Buckling Process

Buckling is the deformation of a thin-film due to high internal stresses and poor adhesion to the surface below. While typically undesirable, this is used to our advantage here by engineering some areas of the wafer to have better adhesion than others, through the deposition and patterning of a low adhesion layer. A ‘Teflon-like’ fluoropolymer produced from the passivation step of a DRIE process (further described in Chapter 3) acts as a surface passivation layer. This material has the benefit of being highly transparent, so that its impact on the optical properties of the cavities is negligible. The resultant buckles are well described by existing solutions for the problem of the elastic buckling of a clamped circular plate [47]. In order for buckling to occur, the

combined stress of the multilayer thin-film stack must overcome the critical stress σ_c , predicted by elastic buckling theory:

$$\sigma_c = 1.2235 \left(\frac{E}{1 - \nu^2} \left(\frac{h}{a} \right)^2 \right), \quad (2.28)$$

where E is the effective Young's modulus of the layer (which is a composite Bragg mirror in this work,) ν is the Poisson ratio of this layer, h is the thickness of the top mirror, and a is the radius of the circle of poor adhesion. Once buckled the cross-sectional profile may be predicted by [47, 48] :

$$\Delta r = \delta \cdot [0.2871 + 0.7129 \cdot J_0(\mu r)]. \quad (2.29)$$

Here, $\Delta(r)$ is the height, J_0 is the Bessel function of the first kind and order zero, $\mu=3.8317$, and δ is the peak height of the dome. The peak height may either be measured and fit, or else predicted for a known biaxial thin-film compressive stress σ by [48]:

$$\delta = h \sqrt{1.96 \left(\frac{\sigma}{\sigma_c} - 1 \right)}. \quad (2.30)$$

This approximate expression holds true for $\sigma \gg \sigma_c$. However, it fails to predict the full non-linear nature of the buckling process, in which for large a 's the buckling may take on a higher order (and non symmetric) buckled mode profile. Non-linear contributions to the buckling from plastic deformation will lead to deviations in the true buckling height from the predicted profile. From the buckled profile, a radius of curvature within the central region may be predicted [33]:

$$R_{C0} = \frac{(2a)^2}{8\delta} + \frac{\delta}{2}. \quad (2.31)$$

The radius of curvature (R_{C0}) of the upper mirror is a critical parameter for modelling the optical modal behaviour of the Fabry-Pérot. The onset of buckling is typically induced by controlled heating of the sample. By heating up the substrate, this induces both thermal expansion among the various materials as well as decomposition of the low adhesion layer (~ 400 °C for the material used

Material	E (GPa)	ν	ρ (kg/m ³)
a-Si	60-100	0.22	2300
SiO ₂	72	0.17	2200
Ta ₂ O ₅	100-140	0.23	6800
4p Si/SiO ₂ QWS	74	0.18	3900
7.5p Ta ₂ O ₅ /SiO ₂ QWS	93	0.2	4200

Table 2.2: Mechanical Properties of all relevant materials for use in the mechanical analysis [49, 50, 51]. The composite mirrors have their effective properties calculated as a thickness-weighted mean of their constituent materials. All mirrors are designed at 1550 nm.

here), which reduces adhesion and allows for the release of stress as described above. In some rare cases, spontaneous buckling has been observed in which the devices form immediately after deposition of the high stress top mirror.

Dynamical Properties

The suspended buckled top mirror is expected (and has been shown) to exhibit mechanical vibrational resonance frequencies close to those of the membrane described in the next section. As discussed by Bitarafan *et al.* [48], the resonance frequencies of the dome may be predicted by modelling the top mirror as either a plate or a shell. It was shown that while the predictions of the models converge at dome diameters over 100 μm , below this the real behaviour is somewhere in between. The predicted frequencies of these two models for the plate and shell respectively are given by [48]:

$$\omega_{\text{P},n} = \frac{\beta_n}{a^2} \sqrt{\frac{D}{\rho h}}, \quad (2.32)$$

$$\omega_{\text{S},n} = \sqrt{\omega_{\text{P},n}^2 + \frac{E}{\rho R_{\text{S}}^2}}. \quad (2.33)$$

Here $\omega_{\text{P},n}$ and $\omega_{\text{S},n}$ are the resonant angular frequency for the n^{th} mode of the shell plate and shell models respectively, β_n is a mode-specific tabulated value ($\beta_1=10.216$, $\beta_2=21.261$, $\beta_3=34.877$), a is the lateral radius of the dome, and

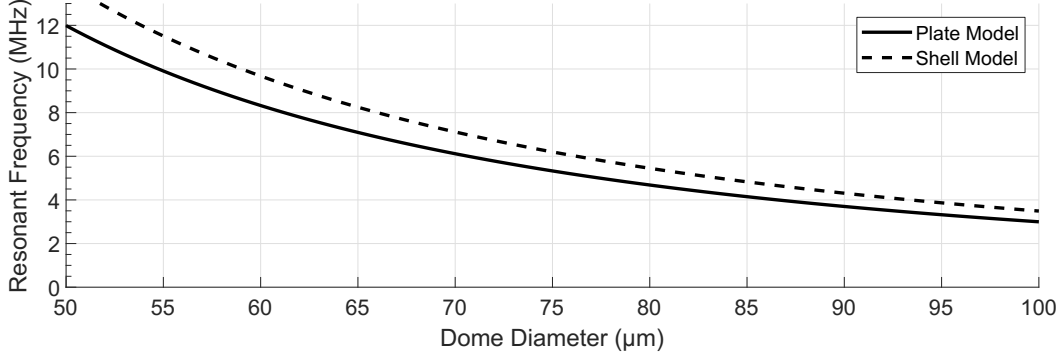


Figure 2.9: Predicted vibrational mode frequencies of the top mirror using both the shell and plate model, for a 7.5 period $\text{Ta}_2\text{O}_5/\text{SiO}_2$ based mirror.

$D = Eh^3/(12(1 - \nu^2))$ is the flexural rigidity, E is the Young’s modulus, and ρ is the mass density of the top mirror, treated as an effective medium. Finally, R_S is the effective radius of curvature of the buckled mirror (Eq. 2.31).

The resonances of a 7.5-period $\text{Ta}_2\text{O}_5/\text{SiO}_2$ 1550 nm QWS, representative of the top mirrors used in this work, is modelled as both a plate and shell and plotted in Fig. 2.9 as a function of base diameter. In the case of the shell model, Equation 2.31 is used to approximate the radius of curvature for a dome of height $\sim 3 \mu\text{m}$.

2.2.2 Mechanics of a Membrane

Consider a thin and wide piece of material, held clamped and tight at its boundaries so as to keep it suspended under tension. Rigidity does not play a part in keeping the objects shape, and much like a trampoline it returns to its original position if perturbed - solely due to the tension at which it is held. A ‘membrane’ model can be used to accurately describe the mechanical motion of this object, under the following assumptions [52]: (1) the thickness is small compared to its extent in the in-plane x and y directions; (2) internal stress is entirely tensile; (3) material resistance to bending is negligible ; and (4) the vibrations are in the out-of-plane z direction. Assuming time harmonic solutions of the form $u = w(x, y)e^{j\omega t}$, with angular frequency ω , uniform biaxial tensile stress $\sigma = T'_{xx} = T'_{yy}$, uniform mass density ρ , and radius a , a membranes motion may be described by [52]:

$$\frac{\partial^2 w(x, y)}{\partial x^2} + \frac{\partial^2 w(x, y)}{\partial y^2} + \frac{\rho a^2}{\sigma} \omega^2 w(x, y) = 0, \quad (2.34)$$

where $w(x, y)$ is the out-of-plane displacement for each point on the membrane surface. For certain geometries (*e.g.* rectangular or circular,) closed-form analytical solutions to Equation 2.34 exist [52]. For the case of a circular membrane, the set of eigenmodes that solve Equation 2.34 in cylindrical coordinates are:

$$w_{m,n}(r, \theta) = a_{m,n} \cos(m\theta) J_n(\omega_{m,n} r), \quad (2.35)$$

with accompanying eigenfrequencies

$$\omega_{m,n}/2\pi = \frac{1}{2\pi} \sqrt{\frac{\sigma}{\rho}} \frac{\alpha_{mn}}{a}. \quad (2.36)$$

Here $a_{m,n}$ is a normalization constant such that the maximum displacement for each mode is unity, J_n is the Bessel function of the first kind and order n , σ is the membrane stress, ρ is the membrane mass density, a is the membrane radius, and α_{mn} is a constant related to the zeros of the Bessel function. Figure 2.10 shows the shapes of the eigenmode solutions given by Equation 2.35. These modes are numbered as (m,n) , where m and n are the mode numbers used. Each eigenmode has a distinct shape (Fig. 2.10,) and varying patterns of displacement (anti-nodes shown as red and blue sections.) Based on how much of the membrane ‘participates’ in the displacement of a particular mode, an effective mass m_{eff} can be determined and is a fraction of the total mass. Effective mass is defined as [53]:

$$m_{\text{eff},n}(x_0, y_0) = \rho t_m \frac{\iint dxdy |w_n(x, y)|^2}{|w_n(x_0, y_0)|^2}, \quad (2.37)$$

where t_m is the thickness of the membrane, and $m_{\text{eff},n}(x_0, y_0)$ is the effective mass for mode n at the point (x_0, y_0) . The point of evaluation will typically be at a maximum of vibration for the particular mode (*e.g.* in the centre for the (0,1) mode as shown in Fig. 2.10). The integral in Equation 2.37 is taken over the entire area of the membrane.




Figure 2.10: Eigenmodes (m,n) from Equation 2.35 of the circular membrane. Red and blue regions indicate anti-nodes, of opposite phase, and white indicates nodal positions of minimal motion.

Vibrations from these modes lose their energy from various interactions with the external environment. This ‘damping’ can come in many forms, with the most notable being (1) viscous or acoustic damping [54] - the surrounding air providing resistance to any movement; (2) clamping losses [32] - the manner in which the membrane is attached at its boundaries that will change how vibrations are coupled to the substrate; and (3) material losses [55] - stresses and strains within the membrane itself that can cause unintended heating. These damping mechanisms will combine to broaden the linewidth of the mechanical mode, decreasing the mechanical quality Q_M . This can be measured in the frequency domain, by the width of the mechanical mode, or in the time domain through a ‘ring down’ measurement of how fast the membrane vibration decays when subject to an external impulse [3]. Silicon nitride in particular is known for its exceptional mechanical quality factors, and a high Q_M at high stresses. This is the topic of much research, and predictions of Q_M require a detailed treatment [56].

The work here however is conducted at atmospheric pressure, under which acoustic damping is typically dominant. By performing experiments in a vac-

uum chamber with pressure approximately < 1 Torr, other forms of loss may become more apparent [54]. The Q factor of a membrane resonator dominated by acoustic losses may be estimated by [54]:

$$Q_{M,n} = \frac{\rho_m t_m f_{M,n}}{\rho_{\text{air}} c_{\text{air}}}. \quad (2.38)$$

Here, t_m is the thickness of the membrane, ρ is the mass density of the membrane or air, c_{air} is the speed of sound in air, and $f_{M,n}$ is the vibrational frequency of the particular membrane mode n .

2.3 Optomechanics

Let this membrane now be placed within the centre of a Fabry-Pérot cavity, in a MIM configuration. The membrane is placed at a fixed position, and then allowed to oscillate small distances about this position. As the membrane moves, it modifies the resonant wavelength of the optical cavity. The amount of influence it exerts over the cavity as it oscillates G (sometimes referred to as the ‘frequency-pull parameter’) depends on its resting position relative to the cavity mirrors. When the membrane is positioned between a node and anti-node of the optical standing wave within the cavity, its influence is at a maximum, and an analytical expression is given by [3]:

$$G_{\text{max}} = 2|r_m|\omega/L, \quad (2.39)$$

where ω is the optical angular frequency of the cavity, L is the cavity length, and r_m is the electric field reflectivity of the membrane (modelled as a simple dielectric slab) given by [57]:

$$r_m = \frac{(n_m^2 - 1) \sin k_z n_m t_m}{(n_m^2 + 1) \sin k_z n_m t_m + i \cdot 2n_m \cos k_z n_m t_m}, \quad (2.40)$$

where n_m and t_m are the refractive index and thickness of the membrane, respectively. When placed between a node and an anti-node, $G = d\lambda/dz$ is conveniently linear for small variations of the membrane around this position.




Figure 2.11: A slice of the two-dimensional overlap functions for the optical Gaussian of unity area (ψ , solid,) and the normalized displacement function for the (0,1) membrane mode (w_{01} , dashed.) These two functions are integrated over the membrane surface to obtain an overlap efficiency η .

Alternatively, placement of the membrane at an anti-node point has the potential to provide quadratic (z^2) coupling which is of interest in some quantum information applications [22, 23, 24].

Both the optical and vibrational modes have distinct two-dimensional shapes and profiles (Fig. 2.11). The extent of their overlap is critical in determining how much influence the vibrational mode can have on the optical mode and vice versa. This can be quantified by the unitless overlap parameter η , given by the following integral [45]:

$$\eta = \left| \iint dxdy \psi(x,y)w_{mn}(x,y)/a_{m,n} \right|. \quad (2.41)$$

Here, $\psi(x,y)$ represents the optical beam profile for a Gaussian (in units of m^{-2} , Equation 2.24,) and $w_{mn}(x,y)$ represents the 2D displacement profile of the mechanical membrane mode (m,n) (Equation 2.35,) and $a_{m,n}$ is the corresponding maximum amplitude of the displacement profile (also from Equation 2.35.) The area underneath this Gaussian function is unity, and the normalized membrane mode profile provides an ‘efficiency map’ of the membrane mode’s influence on the optical mode.

In measurement of such a system, the resonant wavelength (and shifts in the wavelength caused by motion of the membrane) is usually not tracked directly. Instead, a laser is tuned to the cavity optical resonance and the transmitted light is measured by a photodetector or camera. One approach to detect shifts




Figure 2.12: The Lorentzian function (solid,) and its derivative (dashed) for a centre wavelength of 1550 nm and a $\Delta\lambda$ of 10 nm. The derivative is plotted in orange on the right y axis. The maxima's of the derivative (vertical lines,) are given by Equation 2.44.

in the cavity resonance is the so-called ‘tuned-to-slope’ technique, in which the laser is tuned slightly off resonance and onto the slope of the Lorentzian lineshape function of the cavity (Fig. 2.12). The transmitted light from an optical mode can be represented by a Lorentzian of the form:

$$T(\lambda) = \left(1 + 4 \left(\frac{\lambda - \lambda_0}{\Delta\lambda} \right)^2 \right)^{-1}, \quad (2.42)$$

where λ_0 is the centre wavelength, and $\Delta\lambda$ is the FWHM. If the laser is tuned slightly to the right or left of the centre of the mode, small shifts in cavity resonance will change the transmission according to the derivative of Equation 2.42:

$$\frac{\partial T(\lambda)}{\partial \lambda} = -8 \frac{\lambda - \lambda_0}{\Delta\lambda^2} \left(1 + 4 \left(\frac{\lambda - \lambda_0}{\Delta\lambda} \right)^2 \right)^{-2}. \quad (2.43)$$

As the wavelength is detuned relative to λ_0 , Equation 2.43 will give the expected changes in cavity transmission, and thus measurement sensitivity of the system. The change in transmission with wavelength is referred to as the slope parameter $S = dT/d\lambda$ (Fig. 2.12, dashed.) The detuning which provides the maximum possible S can be determined through solving for the zero of the second derivative (indicated as vertical lines on Fig. 2.12):

$$\lambda_{\max} = \lambda_0 \pm \frac{\sqrt{3}}{6} \Delta\lambda, \quad (2.44)$$




Figure 2.13: The maximum derivative of the transmission (S_{\max}) as linewidth is changed (shown from 10 pm to 10 nm.)

which inserted back into Equation 2.43 results in a maximum slope of:

$$S_{\max} = \frac{\partial T(\lambda_{\max})}{\partial \lambda} = \mp \frac{3\sqrt{3}}{4} \frac{1}{\Delta \lambda} \approx \mp 1.3 \frac{1}{\Delta \lambda}. \quad (2.45)$$

The final signal measured by the photodetector will be highly dependent on the particular experimental configuration. This includes factors such as coupling efficiency of the laser into the dome (E_{in}), the peak transmission of the Fabry-Pérot (T_0), and the collection efficiency of the output optics (E_{out}). If the laser produces a power P_{laser} incident on the device, then the power that is transmitted during operation P_0 can be found through:

$$P_0 = P_{\text{laser}} \cdot E_{\text{in}} \cdot T_0 \cdot E_{\text{out}}. \quad (2.46)$$

If we assume that we can measure the average output power incident on the photodetector, P_0 can be determined directly without the need for detailed analysis of the measurement system through Equation 2.46.

All of these described parameters may be combined in order to determine a measurement sensitivity factor α , which directly relates the motion of the resonator to the signal received by the photodetector [53]. Assuming the photodetector is pre-calibrated (*i.e.* the measured signal is in terms of optical

power W), α is given as:

$$\alpha = (P_0 \cdot G_\lambda \cdot \eta \cdot S)^2. \quad (2.47)$$

While $\sqrt{\alpha}$ allows for the direct correlation of membrane movement to optical power, α itself is in units of W^2/m^2 and instead correlates the power spectral density (PSD) of the membrane's motion to the optical power on the photodetector.

In experiment, the movements of the membrane will typically be random and uncorrelated due to thermal fluctuations. A PSD of these vibrations can be accurately modelled through use of the equipartition theorem. When a body is at thermal equilibrium with its environment, the equipartition theorem states that each of its degrees of freedom should have an equal average kinetic energy [53]. Further, if the energies of these degrees of freedom contribute quadratically to the total energy ($E \propto z^2$), the average energy of these modes will be $k_B T/2$ where k_B is the Boltzmann constant, and T is the temperature. Consider the oscillations of a one-dimensional simple harmonic oscillator. The time averaged potential energy membrane can be expressed as [53]:

$$\langle U \rangle = \frac{1}{2} k_{\text{eff}} \langle z \rangle^2 = \frac{1}{2} m_{\text{eff},n} \Omega_n^2 \langle z_n \rangle^2, \quad (2.48)$$

where $\langle U \rangle$ is the time averaged position of the membrane mode, $k_{\text{eff},n}$ is the effective spring constant for mode n , $m_{\text{eff},n}$ is the effective mass for mode n , and Ω_n^2 is the angular frequency of mode n . Observing the quadratic dependence in z , the equipartition theorem can be applied as follows:

$$\langle U \rangle = \frac{1}{2} m_{\text{eff},n} \Omega_n^2 \langle z_n \rangle^2 = \frac{k_B T}{2}. \quad (2.49)$$

Upon rearrangement, we find

$$\langle z_n \rangle = \pm \sqrt{\frac{k_B T}{m_{\text{eff},n} \Omega_n^2}}. \quad (2.50)$$

This expression provides an estimate of the average displacement $\langle z_n \rangle$ of a membrane mode n , which is subject to thermal fluctuations from the environment. These oscillations are expected to be random and incoherent. This

result can be combined with the mechanical susceptibility of a one-dimensional simple harmonic oscillator to obtain an expression for the PSD of a one dimensional micromechanical resonator [53]:

$$S_{zz}(f) = \frac{k_B T f_{M,n}}{2\pi^3 m_{\text{eff},n} Q_{M,n} \left[(f^2 - f_{M,n}^2)^2 + \left(f \frac{f_{M,n}}{Q_{M,n}} \right)^2 \right]}, \quad (2.51)$$

where the expression is for mode n , T is the temperature in Kelvin, $f_{M,n}$ is the natural resonance frequency of the mechanical mode, and $Q_{M,n}$ is the quality factor of the mechanical mode.

The total optical signal PSD measured $S_{\text{WW}}(f)$ then can be predicted using this displacement spectrum [53]:

$$S_{\text{WW}}(f) = S_{\text{WW}}^{\text{W}} + \alpha S_{zz}(f), \quad (2.52)$$

where S_{WW}^{W} is a constant white noise offset associated with the experiment, from sources such as detector dark current and shot noise, and α is a gain factor in units of W^2/m^2 that correlates physical displacement in units of m^2/Hz to the photodetector optical power signal measurements in units of W^2/Hz (assuming the actual electrical signal measured by the detector has been properly calibrated in terms of optical power) [53].

2.3.1 Cavity Optomechanical Systems

To frame this work in the greater context of cavity optomechanics, it is useful to briefly define a few important parameters commonly used to describe these systems. As was shown, the optomechanical interaction may be quantified by the frequency pull parameter G , relating movement of the mechanical element to a change in the cavity resonant frequency. A more fundamental parameter, of interest in quantum applications, is the single photon or vacuum coupling strength g_0 [4]:

$$g_0 = G \cdot x_{\text{ZPF}}. \quad (2.53)$$

The vacuum coupling strength describes the efficiency of the interaction between a single photon and a single vibrational phonon. It is dependent on the zero-point-fluctuation amplitude x_{ZPF} [4]:

$$x_{\text{ZPF}} = \sqrt{\frac{\hbar}{2m_{\text{eff}}\Omega_{\text{m}}}}. \quad (2.54)$$

where m_{eff} is the effective mass of the mechanical mode, and Ω_{M} is the vibrational resonance angular frequency. The so-called optomechanical coupling strength g , which takes g_0 and scales it by the optical power (number of photons), can now be found with [4]:

$$g = g_0\sqrt{\bar{n}_{\text{cav}}}, \quad (2.55)$$

where \bar{n}_{cav} is the average number of photons in the cavity:

$$\bar{n}_{\text{cav}} = \frac{1}{\Delta^2 + \left(\frac{\kappa}{2}\right)^2} \frac{P_{\text{in}}}{\hbar\omega_{\text{L}}}. \quad (2.56)$$

where $\Delta \equiv \omega_{\text{L}} - \omega_{\text{cav}}$ is the laser detuning relative to the optical resonance, ω_{L} is the angular frequency of the laser, and P_{in} is the input laser power. The maximum coupling is achieved when the detuning is set to $\Delta \approx -\kappa/2$, which results in:

$$\bar{n}_{\text{cav,max}} = \frac{P_{\text{in}}}{\kappa\hbar\omega_{\text{L}}}. \quad (2.57)$$

Optomechanical interactions may be placed into one of two regimes based on the optomechanical coupling strength g and the cavity optical loss κ . These regimes are the ‘strong coupling’ regime, in which $g > \kappa$, and the ‘weak coupling’ regime, in which $g < \kappa$. The strong coupling regime is of particular interest for quantum systems, in which the mechanical and optical modes will begin to ‘hybridize’ with one another, and the previously single mechanical mode will split into two new modes [4]. These coupled modes form the basis for information transduction between the optical and mechanical regimes, and even coherent quantum state transfer with the right conditions met [4]. Another important relationship is the relative values of the mechanical

oscillator frequency and the cavity linewidth. In the limit of $\Omega_M > \kappa$ one reaches the so-called ‘resolved sideband’ regime, which is a typical criterion for the possibility of the mechanics to be cooled into their quantum ground state through optomechanical sideband cooling [6].

2.3.2 Optomechanical Cooling

When the conditions within the optomechanical system are right, light in a cavity can be tuned such that it can apply a force to the mechanical element. This so-called ‘dynamical back action’ can act to effectively heat or cool the mechanical element - *i.e.* accelerating or damping its thermal motion. This effect can be implemented in all optomechanical configurations, including MIM systems. Consider the canonical Fabry-Pérot optomechanical cavity (Fig. 2.14), in which one of the mirrors is fixed and the other is free to move as a simple harmonic oscillator. As light resonates in the cavity, a radiation pressure is applied to all surfaces from which the photons are reflected. This means that the movable mirror on the right receives a radiation force towards the right, for example.

As the free mirror moves left or right, it changes the cavity length and thus cavity optical resonant wavelength. If the input light is slightly detuned from the cavity resonance, movement of the optical mode Lorentzian (resonant wavelength) will cause the fixed laser light to sit at different points along the Lorentzian, modulating the amount of optical power stored within the cavity (*i.e.* at the detuned wavelength) with respect to the mirror position. Since radiation pressure is proportional to optical power (or number of photons,) the mirror will now feel a force that is a function of its position, creating a sort of feedback loop. When an object moves through a force, the energy it exchanges with its environment is quantified by work. Work can be expressed through a line integral of a force through space, or simply just the area underneath a force vs. distance curve $W = \int F(x)dx$. As the mirror moves through the changing force (Fig. 2.15(a)) the work done by the mirror is either positive or negative, depending on the direction of its motion. However, provided the radiation force responds instantaneously to changes in the mirror position,




Figure 2.14: A simple optomechanical cavity model, with one mirror fixed and the other free to oscillate. Note that the radiation pressure (or force) is always in the same direction when applied to the right mirror, regardless of its position.

the net area over a cycle is zero ($W_{net} = \oint F(x)dx=0$), equally cancelled by the motion left and right (indicated by green and violet lines respectively in Fig. 2.15(a).) In this case, there is no net exchange of energy between the mirror's mechanical motion and the cavity. Figure 2.15(a) shows this for a red-detuning, but the effect is the same for a blue-detuning in the case of an instantaneous response.

If the cavity's ability to respond to changes is slower than the mirror vibrations ($\kappa < \Omega_M$), the force felt by the mirror no longer perfectly traces the linear ideal curve of the Lorentzian (Fig. 2.15(b)). The force seen upon a change in position is 'retarded', resulting in a force representative of the old position of the mirror some finite time before the cavity can catch up. This introduces a hysteresis in force felt by the mirror, so a slightly different force is felt on the mirror depending on the direction it travels. The net work integral is no longer zero, since the area traced under the curve is different whether the mirror is moving left or right. A non-zero work implies that energy is being exchanged between the cavity and the mirror. In the case of red-detuning,




Figure 2.15: A diagram of the work done by the membranes motion if the light in the cavity responds (a) instantaneously, and (b) retarded to the changes in resonance. In the retarded case, there is a hysteresis that is introduced depending on the direction travelled, and a non-zero work is now done by the mirror on the optical cavity. Note that as the cavity length changes, it is the Lorentzian curve that moves left to right, while the laser stays fixed with respect to wavelength.

the work done by the mirror is negative, meaning energy leaves the mirror and results in ‘cooling’. The description here is applicable to blue-detuning as well, in which case the slope of force is now positive instead, and the net work becomes positive. The result is the mirror ends up having work done on it by the cavity, so that ‘heating’ occurs instead.

Chapter 3

Monolithically integrated membrane-in-the-middle cavity optomechanical systems

Considerable effort has been put into the development of dome microcavity systems by the DeCorby group over recent years. The process has been refined to a state where large arrays of devices can be created through a highly controlled microfabrication process. While interesting in their own right, these devices have until recently lacked some sort of functionalization. Integrated Fabry-Pérot cavities of this size can have some intriguing properties when coupled to some sort of matter - be it fluids, atoms, solid-state emitters, or mechanical elements. In order to introduce something into the cavities, the fabrication process must be modified to somehow allow for, at some point, access to the internal cavity region for the placement of the object of interest.

The work in this section was done in tandem with efforts to create open access domes for use with emitters and fluids, described in Ref. [39]. The timeline of the project may be divided into three parts. First, the stand alone XeF_2 sacrificial etch based release of silicon nitride membranes was studied. In subsequent phases, we explored the production of two fully integrated devices: one wafer with a Si/SiO_2 top mirror, and a second wafer with a $\text{Ta}_2\text{O}_5/\text{SiO}_2$ top mirror. The final trial yielded promising results, and is the focus of the characterization efforts that follow. A submitted publication [58], under review at the time of writing, mirrors much of what is discussed in this chapter.




Figure 3.1: A schematic overview of the fabrication process. (a) Bottom mirror, sacrificial layer, and etch stop are sputtered onto the substrate. (b) PECVD silicon nitride is deposited on the wafer, and annealed to obtain tensile stress. (c) A low adhesion layer is patterned as circles to define where the domes will form after buckling. (d) The top mirror is sputtered on top of the patterned low adhesion layer. (e) Through controlled heating, the top mirror undergoes controlled buckling and forms an air cavity. (f) Access holes are etched with a mask through the use of RIE, penetrating the surface of the membrane. (g) The membrane is released through a XeF_2 based etch of the sacrificial layer through the access holes. If required, additional periods may be sputtered after this step to further increase mirror reflectivity.

3.1 Fabrication

An overview of the fabrication process is shown in Figure 3.1, with the steps described in detail in the following sections. To start the process, a double-side-polished silicon wafer (University Wafer, $\langle 100 \rangle$ orientation, $500\ \mu\text{m}$ thick, P-doped, resistivity $\sim 10\ \Omega\ \text{cm}$,) was cleaned in a Piranha solution for 15 minutes. Sputtering was then performed in a DC magnetron sputtering system, and a 3-period Si/SiO₂-based 1550 nm QWS mirror was deposited. The completed mirror was immediately followed by a 30 nm thick ‘etch stop’ layer of SiO₂, and a 200 nm thick ‘sacrificial layer’ of Si.

3.1.1 Sputtering

Sputtered films provide desirably low optical loss (*i.e.* films are highly uniform, homogeneous, smooth, and free from contaminants,) and tunable compressive stress. Figure 3.3 shows a schematic of the sputtering system used for this work. The substrate to receive the films is placed in a vacuum chamber opposite to a ‘target’ containing the material to be sputtered. The chamber is pumped to a vacuum pressure of $\sim 10^{-6}$ Torr, and then a large voltage bias (~ 400 to 500 V) is applied between the target and substrate. Argon gas is introduced into the chamber, which becomes ionized due to the large bias. The ionized argon causes a plasma to form, which encourages ion bombardment of the target surface. These impacts cause plumes of the target material to be released, and travel across the chamber to cover the substrate. An illustration of this process is shown in Figure 3.2. This particular system also utilizes what is known as a ‘magnetron’ configuration on the surface of the target. Through introducing a permanent magnetic field within the guns, secondary electrons produced during target bombardment are kept close to the target surface and are used to increase the sputtering rate and efficiency. Target health and integrity is also further preserved through the use of a ‘pulsed DC’ (PWM) supply, as opposed to a constant DC bias. By rapidly switching the bias from positive to negative in a square wave configuration and with a variable duty cycle, the build-up of foreign materials on the target surface can be mitigated and the efficiency of the sputtering is increased. Build-up of foreign material on the target surface (*e.g.* redeposited from the target or an oxide,) can lead to poor film quality and electrical arcing, which damages the target.

Recipes and rates for various materials are given in Table 3.1. The substrate will typically be heated to 150 °C,¹ critical in achieving high compressive stress. Some materials, such as the oxides, are sputtered ‘reactively’. Oxygen (O₂) gas is introduced into the chamber alongside the argon, and reacts with

¹The substrate holder is heated through a set of lamps shining on the rear side. A PID controller monitors the temperature of the holder through an attached thermistor. It is suspected that this temperature set point is inaccurate, for a number of reasons, so how close or repeatable the depositions are to this temperature is unknown.




Figure 3.2: A cartoon illustration of the sputtering process. Ionized argon atoms form a plasma, and are propelled by the external bias into the target. The impacts cause plumes of the target material to be released, which coat the substrate.

the target material to form an oxide on the substrate surface ². Further details on the practical use of Doug can be found in Appendix D.

3.1.2 PECVD Silicon Nitride

On top of the sputtered films, a plasma-enhanced chemical vapour deposition (PECVD) system is used to grow a thin layer of silicon nitride to act as the future membrane. Silicon nitride is a popular material of choice for mechanical membranes due to its ability to vibrate with exceptionally high mechanical Q's [45].

Low pressure CVD (LPCVD) is the most popular process for creating silicon nitride layers suitable for membranes, and is notably the process of choice in the production of commercial membranes by Norcada Inc., an Edmonton-based company which was consulted for some advice on the current project. One of the primary reasons for performing CVD at low pressures is that it significantly reduces the amount of gas-phase nucleation during the reaction [59]. When thin films are formed on the substrate surface, as opposed to elsewhere in the chamber, a more consistent and higher quality thin-film can be

²As of February 20, 2020 the pressures for the SiO₂ and Ta₂O₅ recipes were adjusted from 4.0 mTorr to 3.7 mTorr. This was due to a drop in the measured stress, believed to be associated with a reconfiguration of the gas flow inlet location during December 2019.

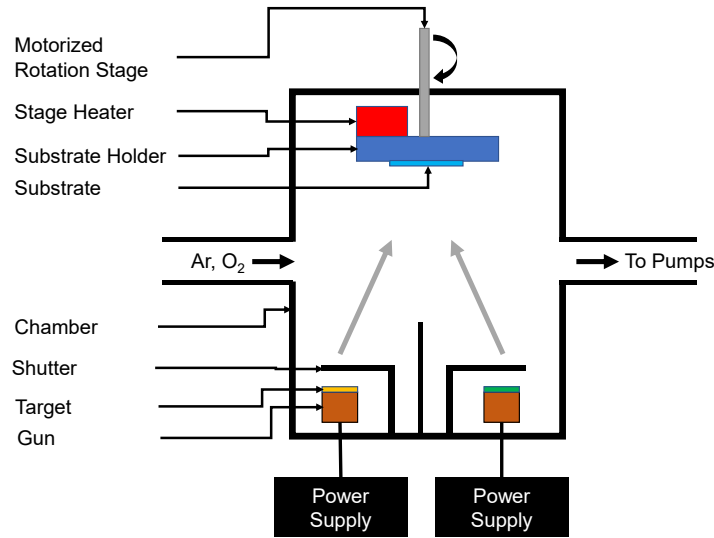


Figure 3.3: A schematic of the Doug sputtering system, described in detail in the main text.

formed. LPCVD will typically require extremely high substrate temperatures, ranging from 700 °C to 900 °C. On the other hand, PECVD can achieve similar results to LPCVD at much lower temperatures (300 °C to 400 °C) by utilizing a plasma instead to supply the necessary activation energy for the chemical processes. LPCVD would have been preferred for this process, however due to equipment access restrictions this was not possible. LPCVD and PECVD are interchangeable in the process here, and thus membrane quality can be potentially improved in future work [60]³. Nevertheless, PECVD has also been used to successfully create MEMS silicon nitride membranes [61, 62, 63]. Specific to this project, a PECVD based process is further appealing since within the next year it is expected that the entire thin film deposition process will be moved to a new PECVD system within the nanoFab, which will allow for a wide variety of advanced materials to be integrated into the process, such as low loss hydrogenated amorphous silicon [64]. With all materials being deposited in the same system, both the deposition of our mirrors and the silicon nitride could be done together without breaking vacuum, thereby reducing

³A notable caution with using LPCVD would be subjecting the substrate and its mirror to high temperatures. Annealing tests in related work by the DeCorby lab have concluded Si/SiO₂ based mirrors endure these annealing processes with no issues.

Table 3.1: Recipes used in Sputtering System 2, ‘Doug’.

	Non-Reactive		Reactive	
	Si		SiO ₂	Ta ₂ O ₅
Gun	1		2	2
Temperature (°C)	150		150	150
Power (W)	200		200	200
Off Time (µs)	0.5		0.8	5
Frequency(kHz)	150		150	20
Ar Flow (sccm)	50		50	40
O ₂ Flow (sccm)	0		3.2	20
Ramp Time (s)	600		600	60
Burn In Time (s)	60		120	120
Ramp Pressure (mTorr)	10		10	7
Deposition Pressure (mTorr)	3.0		3.7	3.7
Deposition Rate (nm/s)	0.142		0.265	0.231

complications such as contamination that may compromise the quality of the membrane.

PECVD nitride is deposited such that it starts with a compressive stress (typically around 200 MPa compressive.) However, through high temperature annealing tensile stress can be achieved [65]. Tensile stresses are required for a functioning membrane, as the membrane must be held tight on all sides for it to properly resonate and stay in place after it is released from its sacrificial layer below. In addition, particularly notable in the case of silicon nitride, high stresses can result in correspondingly higher mechanical quality factors [56]. Annealing of the wafer is done within a quartz tube furnace (Tystar General Anneal) where it is heated in a nitrogen environment for 2 hours at 600 °C, resulting in a stress of approximately 770 MPa tensile.⁴

3.1.3 Low Adhesion Layer

On the surface of the annealed silicon nitride, photolithography is next performed using a set of two masks. First, a mask containing alignment marks is

⁴Due to interference of the Flexus stress measurement system with the bottom mirror layers, it was not possible to get an accurate stress measurement. Values are instead taken from a test wafer processed side-by-side with the device wafer.

patterned via photolithography in AZ1512 resist. A standard recipe was used, resulting in an approximately 1.1 μm thick photoresist layer. The substrate was taken to the inductively coupled plasma reactive ion etch system (Cobra ICPRIE, described later on) for a brief etch (~ 5 min.) The mask is removed in acetone, and the etched alignment marks are then used for the patterning of a second mask for the low adhesion layer liftoff. This mask contains circles of either 50 μm or 100 μm in diameter. The mask is designed such that, after development, the regions where the chrome features are not present correspond to the open circle regions in the photoresist.

The low adhesion layer utilized here is a material similar to polytetrafluoroethylene (PTFE), formed on the device surfaces through a vapour-phase deposition. This material is commonly used as a passivation layer in cyclic DRIE processes, such as in ‘Bosch’ etching. Fluoropolymers in particular are known for their unusually low adhesion to other materials. Fluoropolymers of this sort are perhaps most well known as Teflon, which finds household use in non-stick cookware exactly due to this property. When referring to a fluoropolymer in this work, it is understood to mean a chain of CF_x molecules of indeterminate length and stoichiometry. During a CVD-like deposition, C_4F_8 is flowed into a chamber with a plasma present. The result is an isotropic polymer coating, with fluorine-to-carbon ratios that can be highly dependent on deposition parameters such as pressure, plasma power, and substrate temperature. The recipe used has been adapted and optimized for our buckling process. While thin-film fluoropolymers may be similar in their basic chemical structure to PTFE or commercial Teflon, they can vary widely in terms of important physical properties such as film adhesion, hydrophobicity, and vapourization temperature. The composition of the films in terms of polymer chain length and stoichiometry is determined by many factors, and can be difficult to analyse.

The Alcatel AMS110 system located in the nanoFab was used for the deposition of the fluoropolymer. Processing first involves a two-stage chamber clean using SF_6 and then O_2 with a dummy wafer loaded into the machine. This is critical to remove contaminants and oxides that may have formed as

a result of other users' processes. The chamber is then conditioned for 30 seconds using the established deposition recipe, with the dummy wafer still loaded. This coats all exposed surfaces in the chamber, and so by conditioning beforehand the process has time to stabilize. The Alcatel is an older system, and has been known to have numerous issues with reliability and uniformity in its processing. For this reason, as of early 2020 we are in the process of switching the deposition process over to a new machine, the Oxford PlasmaPro Estrelas 100. However, the Estrelas system is in high demand in the nanoFab facility, and access is often difficult to obtain.

Once the low adhesion layer has been deposited, the photoresist is lifted off through sonication in acetone for approximately 1 hour. The wafer is subsequently cleaned with IPA, water, and dried. The substrate is then returned to the sputtering system, and 4.5 periods of a Ta₂O₅/SiO₂ 1550 nm Bragg mirror are sputtered in a similar manner to the bottom mirror.

3.1.4 Hotplate Buckling

In order to induce the 'self-assembly' of the domes into their buckled state, the substrate must be heated. Buckling has typically been done on a standard lab hotplate in the fume hood inside the nanoFab characterization lab. The sample is placed on the hotplate, with an infrared temperature sensor monitoring the temperature of the wafer, which tends to deviate from the hotplate 'setpoint'. The entire setup is placed underneath a fume hood as a precaution, since the vapourization of PTFE is known to produce harmful fluoride compounds. A microscope with a digital camera attached to a computer enables monitoring of the buckling process from above, and is a good visual indicator of when buckling is complete. A typical buckling schedule is as follows. The hotplate is pre-heated, with no sample present, to 150 °C. Once the temperature has been allowed to stabilize, the sample is placed on the hotplate underneath the infrared thermometer. Every minute, the temperature setpoint on the hotplate is raised by 50 °C and the setpoint and infrared measured temperatures are recorded. Buckling will typically begin to occur at 300 °C to 400 °C. A typical sample will stay on the hotplate for a total of 7 minutes, reaching a final

temperature of ~ 400 °C. The buckling process is a balance of obtaining a mostly buckled wafer, since not all features will buckle completely at the same time, and not heating the wafer for too long or at too high a temperature, since this will result in the destruction and cracking of buckles. Development is ongoing to move the buckling process to a rapid thermal annealing (RTA) system, which is a highly controlled and uniform form of heating which has demonstrated higher yields and consistency over hotplate buckling.

3.1.5 Cobra Etching

The cavity access design that involves creating holes down from the top of the wafer requires an anisotropic etch that has a low selectivity among the various materials that may be used in the top mirror and membrane. RIE processes using fluorine-based gases such as CHF_3 and SF_6 have a high etch rate for silicon and tantalum compounds, and fulfill a dual role of continually producing fluoropolymers for sidewall passivation, much like the previously described DRIE process. The nanoFab's Oxford PlasmaPro 100 Cobra system allows for the plasma-driven RIE using CHF_3 and SF_6 , and existing recipes are accessible on this machine that have demonstrated reasonable etch rates through the required materials. Anisotropy of the etch process and a low selectivity between Si and SiO_2 may be achieved through the introduction of hydrogen ions into the etch process [59]. This is achieved in part with the choice of CHF_3 gas here, instead of a similar chemical such as CF_4 . Further, the use of CHF_3 in the process is highly comparable to the process used for the aforementioned CVD of fluoropolymers in the Alcatel system. CHF_3 is highly likely to form a polymer due to excess carbon atoms present during silicon fluorination and removal [59]. This results in the coating of post-etched surfaces with fluoropolymers. This is undesired in most processes - however, it may be beneficial here to protect the etch hole side walls from any etching during the subsequent XeF_2 etch step. Some evidence of this was observed during processing, and is discussed later in Section 3.2.2. If buckling is attempted after etch holes have been created, there tends to be very low yields with most devices not buckling at all. One possibility is that this fluoropolymer layer acts

Table 3.2: Recipes used in the Oxford Plasmapro 100 Cobra Metal Etch system, for etching of the alignment marks and the access holes.

SiO₂+Ta₂O₅+TiO₂ Etch	
Pressure Setpoint (mTorr)	90
CHF ₃ Gas Flow (sccm)	50
SF ₆ Gas Flow (sccm)	10
Ar Gas Flow (sccm)	25
HF Power (W)	300
ICP Power (W)	1000
He Chuck Backing Flow Pressure (torr)	10
Chuck Temperature Setpoint (°C)	20
Typical DC Bias (V)	667

as a glue between the top mirror and the layer below, preventing them from de-adhering. Further details of this phenomena are described by Maldaner in Ref. [66].

As mentioned, the recipe used in the Cobra system was one previously developed by nanoFab Staff, and it was not modified in any way except for changing the duration of the total process to match the thickness of the layers being etched. A standard cleaning recipe is typically done for 10 minutes before processing, and then the etch recipe is run with a dummy wafer loaded for 5 minutes in order to condition the chamber. The etching recipe is given in Table 3.2. Prior to etching, the wafer is masked via lithography using AZ 4620 resist with an approximate thickness of 12.5 μm . A modified recipe from the standard procedure has been developed by our group, which expedites otherwise lengthy rehydration times. The mask has etch holes of various sizes, numbers, and placement diameters. Both the modified photoresist recipe and details on this mask are given in Appendices D and E respectively. Care is taken to align this mask to the alignment marks previously etched onto the surface below.

Prior to testing the etch step on the actual devices, estimates of etch rates were determined by etching wafers with single sputtered layers of each material. These were masked with a Sharpie pen drawn on portions of the wafer. After etching, the Sharpie mark was removed and the depth was determined

Table 3.3: Approximate etch rates for various materials, as tested in the Cobra system with the ‘SiO₂+Ta₂O₅+TiO₂ Etch’ recipe.

Material	Bulk Etch Rate (nm/min)	Comment
a-Si	350	
SiO ₂	190	
Ta ₂ O ₅	80	
Si ₃ N ₄	550	Approx., very fast
AZ4620 (Mask)	160	Approx., likely lower

with a step profilometer in order to obtain the change in height. Table 3.3 shows the results of these trials. Based on these ‘bulk’ estimates, a 4.5-period Ta₂O₅/SiO₂ top mirror should require approximately 17 minutes of etching to get through the top mirror and membrane, without going too far through the 200 nm Si sacrificial layer. The 350 nm/minute etch rate of Si is quite fast, so the precision of the etch to stop somewhere within the middle of the 200 nm sacrificial layer needs to be within less than a minute. The masked device wafer was cleaved into smaller pieces for testing, and it was found that 18 minutes was adequate to reach the sacrificial layer without compromising the underlying SiO₂ etch stop layer (as evidenced by reasonable results from the XeF₂ sacrificial etch step that followed.) Once the etching was complete, the mask was cleaned by soaking in acetone for 5 minutes, followed by a clean with IPA and water, drying with nitrogen, and heating on a hotplate for 10 minutes at 110 °C to remove any residual moisture.

3.1.6 XeF₂ Sacrificial Layer Etching

In order for the membrane to be free and suspended, it must first be released from the sacrificial layer below. In choosing the chemistry of the sacrificial layer, it is critical to consider the selectivity of the etchant that will be used. Most etchants do not have a ‘perfect’ selectivity to other materials, meaning that some of the material one wishes to keep will inevitably be removed as well as the sacrificial layer during the etch. In a general sense, the following criteria must be satisfied in order for the chemistry of a particular sacrificial layer and etch process to be a good fit for our devices: (1) must remove the sacrificial

Table 3.4: A survey of the various possible methods for removing a sacrificial layer within the domes [67, 68]

Method	Sacrificial Layer	Typical Etchant	Advantage	Issue
Metals	Al, Cu, etc.	HCl, FeCl ₃ , H ₂ O ₂ , H ₃ PO ₄	Selectivity	Porosity in Metal
Organic	PDMS, PTFE, Parylene etc.	Acetone, Heat, O ₂ plasma	Selectivity	Temp. Stability
KOH	Si	Wet KOH	Common Process	Etches both Si and SiO ₂
HF	SiO ₂	Wet/Vapour HF	Common Process	Etches Si ₃ N ₄
XeF ₂	Si	Vapour XeF ₂	Selectivity	Si is used in mirrors

layer material, while leaving the materials used in our devices untouched. These device materials could potentially include a-Si, SiO₂, Ta₂O₅, and Si₃N₄; (2) must release membranes in a way that avoids ‘stiction’ - the common phenomena of MEMS devices being statically pulled down upon drying after a wet etch via surface tension; and (3) must function well when etching in a space with limited outside access, since the narrow and deep etch holes may restrict the introduction of fresh reactants and spent by-products.

Many methods were considered, and the details of some potential candidates are listed in Table 3.4. From a selectivity standpoint, a metal or organic sacrificial layer would be the best, since the devices are otherwise free from these types of chemistries. Pure metals (*e.g.* Al, Cu, Ag) could easily be sputtered alongside the bottom mirror materials, and thus could be easily introduced into the existing sputtering process. Etching of these elements will typically use wet chemical compounds that specifically target metals, and have virtually no effect on oxides or Si based compounds. There is a history of the use of metal sacrificial layers in MEMS, but it is uncommon [69]. Some nitride membrane release has interestingly been done with electron beam evaporated germanium as a sacrificial layer, and H₂O₂ as an etchant to achieve an excellent selectivity and appreciable etch rate [63]. However, the wet etching of metals with hydrogen-based chemistries is known to produce hydrogen bubbles,

and in confined spaces this has the potential to both impact etch rate and etch uniformity, and even cause high local pressures that can rupture sensitive membrane structures [70]. On top of this, all wet etches require a constant refreshment of etchants and removal of by-products (typically induced by agitation), which is expected to be made even more difficult by the etch holes. Past work on similar devices within the DeCorby group has indicated some difficulty with liquid intrusion into the domes, likely due to the hydrophobic nature of residual PTFE inside.

An organic sacrificial layer brings similar benefits as a metal layer, in that the only other organic polymer compound in the final devices is the fluoropolymer low adhesion layer, which at this stage has already served its purpose and could be removed without harm. The use of a polymer-based sacrificial layer also has a history in MEMS technology [71, 72]. The low vapourization temperature of many polymers allows them to in principle be removed through heating, with no chemicals or concerns around selectivity involved. This allows the polymer to decompose and exit the devices in a gaseous state. However, the PECVD process used to deposit the silicon nitride membrane layer (which is the next step after sacrificial layer deposition) must operate at substrate temperatures of 300 °C to 400 °C, which would be high enough to decompose the polymer prior to proper release. Further, if the process was moved to an LPCVD nitride at some point in the future, deposition temperature at this stage are expected to possibly double. For example, the PTFE compound used here as the low adhesion layer has a vaporization temperature on the higher end of most organic polymers at approximately ~ 400 °C.

More common etches for these materials, such as KOH or HF were also considered. KOH is commonly used to release membrane structures, due to it's ability to etch crystalline Si at sharp angles. KOH etches are done in liquid form, and thus to avoid stiction would require a subsequent step of critical point drying, in addition to the other complications associated with wet etching as discussed for a metal chemistry above. KOH also has an affinity for etching SiO₂, which would mean it would also attack the top mirror materials making it largely incompatible. Hydrofluoric acid (HF) is also commonly used

in microfabrication to etch SiO_2 , and is available in both a wet and vapour phase form. It is highly selective to SiO_2 over Si, which would avoid the selectivity complications of the KOH etch. Some initial testing was done with vapour HF on releasing nitride membranes. However, HF was found to have a poor selectivity to silicon nitride and there was considerable damage to the thin and delicate nitride membrane after the etch.

Xenon difluoride (XeF_2) is a relatively uncommon etchant with a similar etching chemistry to HF, but instead using the noble gas xenon as a fluorine carrier. XeF_2 has been demonstrated with an extremely high selectivity of silicon relative to silicon dioxide, tantalum pentoxide, as well as fluoropolymers [67]. The use of XeF_2 in MEMS design to isotropically etch deep cavities of silicon is well noted in the literature [73, 74, 75, 76]. The vapour-phase nature of XeF_2 has the added benefit of mitigating any difficulties associated with stiction. The process of etching is very similar to other fluorine-based etch chemistries, such as those used in RIE systems, but using Xe as a carrier for the gas allows reactions to be done at room temperature without a plasma and for a very high selectivity to be achieved. The nanoFab has a custom-built XeF_2 etching system, based off of that described in Ref. [77]. Xenon difluoride is located in a canister as commercially sourced solid crystals. This canister is connected to the ‘etching’ chamber via an ‘expansion’ chamber, shown in Figure 3.4. The etching process is done in pulses and consists of three steps. First, the expansion chamber is isolated from the reaction chamber and exposed to the XeF_2 crystals, which sublime and fill the expansion chamber with gas. The expansion chamber is then isolated from the crystals and opened to the reaction chamber. The sample is exposed to the gas and the etch occurs. After a set amount of time, both the expansion and etch chambers are pumped and purged to remove any remaining XeF_2 and by-products, and the process repeats itself. The standard settings for a cycle are shown in Table 3.5. These parameters are from the original recipe developed by the nanoFab.

Prolonging the etch time per cycle would likely reduce the number of cycles required. This is because compared to the samples used in the nanoFab’s characterization tests, there is a relatively low volume of silicon being removed




Figure 3.4: A schematic of the nanoFab XeF_2 etching system. Crystals are first sublimated into the expansion chamber, and then this gas is exposed to the sample over a series of cycles.

Table 3.5: XeF_2 etch process default recipe parameters. These are unchanged from the nanoFab SOP. A cycle typically takes approximately 3.5 minutes.

Parameter	Value
Base Pressure (mTorr)	150
Expansion Time (s)	60
Etch Time (s)	60
Total Etch Cycles	Variable, typically a multiple of 5
Pressure Tolerance (%)	0

in our samples. Thus, the requirement for reactant refreshment is likely lower. On the other hand, reducing the etch time may result in a more controllable etch. It is difficult for the reactants to reach the sacrificial layer through the access holes, and so after the initial reactants have removed the silicon, there is probably a reduced etch rate as the fresh reactants struggle to refresh below in the cavity. Brief pulses, followed by chamber pumping may allow for a more predictable etching.

With the XeF_2 etching complete, the wafer is finished with the sputtering of three additional periods of a $\text{Ta}_2\text{O}_5/\text{SiO}_2$ 1550 nm mirror. This results in the top mirror having a total of 7.5 periods. The additional periods were added after buckling has occurred, due to concerns that otherwise the top mirror would be potentially too thick and rigid for proper buckling. Filling or bridging of the etch holes was not observed with the addition of the extra periods. Any sputtered material overtop of these regions likely stuck to the internal side walls of the etch holes, which should not have an impact on the

optical performance of the cavity.

3.2 Process Development

3.2.1 Free-Standing Membranes

Prior to the integration of the membranes within the cavity, they were first developed on their own to verify that the release and etch process was viable. The sacrificial Si and membrane SiN layers were deposited (as in the full device process described above,) but onto a thermally oxidized silicon wafer. The etch hole mask was then used to pattern etch holes as in the full process (with just a brief ICPRIE etch,) and the wafer was exposed to the same 10 cycles of XeF₂.

The results varied widely depending on the diameter of the hole pattern, and the size of the holes themselves. Pictures of larger membranes (with patterns intended for 100 μm domes, $D \sim 40 \mu\text{m}$, Fig. 3.5) indicate a cleanly formed undercut and a membrane surface that seems somewhat uniform. There are concentric ring outlines visible that appear to reflect the geometry of the holes (Fig. 3.5 (g-i).) These appear to be centred around the holes, and are likely related to the XeF₂ etching process.

One of these membranes was intentionally dislodged with a step profilometer tip (Fig. 3.6(a).) The oxide etch stop surface below (blue in the image) is smooth, indicating the membrane was fully released and that the sacrificial Si was completely removed. This is also encouraging from an optical perspective, since an incomplete removal of the sacrificial Si layer would result in a rough lower surface and could be detrimental to cavity finesse. The dislodged membrane still has the observed concentric rings. We believe this is due to the membrane being slightly etched from underneath to varying degrees (depending on the distance from the etch hole.) While XeF₂ is highly selective to Si₃N₄, the etch rate is still higher than that of SiO₂ and is highly dependent on stoichiometry and stress [68]. The ring patterns might be attributable to the discontinuous, pulsed nature of the XeF₂ etch process (Fig. 3.7.) With each pulse, the exposed section of the underlying membrane ceiling is etched,




Figure 3.5: Microscope images of various stand-alone membranes after XeF_2 etching trials. (a) Incomplete etch with a central ‘island’ of sacrificial layer still present. (b) Completely released membrane ($D \sim 100 \mu\text{m}$). (c) Completely released membrane ($D \sim 50 \mu\text{m}$). (d) Membrane with etch holes too large and too close together, resulting in bridging between holes and a collapse of the central structure. (e) Incomplete etch of a smaller membrane. (f) Membrane in which etching proceeded significantly slower than the others, without the etch cavities joining yet into one. (g-i) Well released membranes with a variety of hole configurations. (Scale bars on each image denote $10 \mu\text{m}$.)

and this creates a stepping pattern. As the etch proceeds, a single pulse needs to remove a volume of sacrificial layer increasing at a rate proportional to the diameter of the circle squared. Additional pulses will show a decreasing spatial extent (when viewed from above) relative to earlier pulses, causing rings at the periphery of the membrane to be closer than those near the holes. The varying exposure of different surfaces to the etch is also visible on the lower oxide etch stop, to a much lesser extent, with the visibility of a faint imprint (underneath the dislodged membrane in Fig. 3.6a) of the etch holes in the SiO_2 where the oxide would have received the most exposure to XeF_2 .

3.2.2 First Generation: Si/SiO₂ Based Top Mirror

With evidence that a process had been developed to create suspended membranes, this step was next integrated into the established dome fabrication process. As a first trial, devices were deposited as described above, however with a top mirror composed of Si/SiO₂ instead of Ta₂O₅/SiO₂. This material combination offers a higher index contrast over Ta₂O₅/SiO₂, and so fewer pe-




Figure 3.6: Microscope images of three side-by-side membranes observed with different membrane integrity. Images are to scale but have been cropped to remove empty space between individual membranes. (a) A dislodged membrane intentionally torn off by the step profilometer probe next to its original location. (b) A damaged membrane, with a large crack and large areas with multicolour interference patterns, possibly an indication of it being in close contact with the surface below. (c) An intact membrane. Scale bar (bottom right) denotes 100 μm .

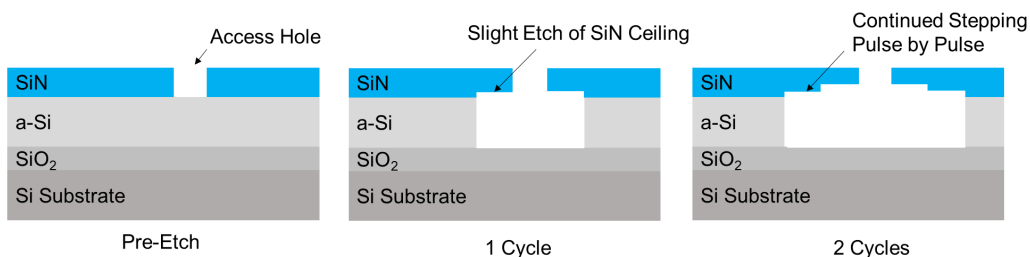


Figure 3.7: A schematic demonstrating a possible explanation for the concentric stepping visible on the membranes (not to scale.) Slight etching of the membrane occurs at every step, with a new piece becoming exposed each time.

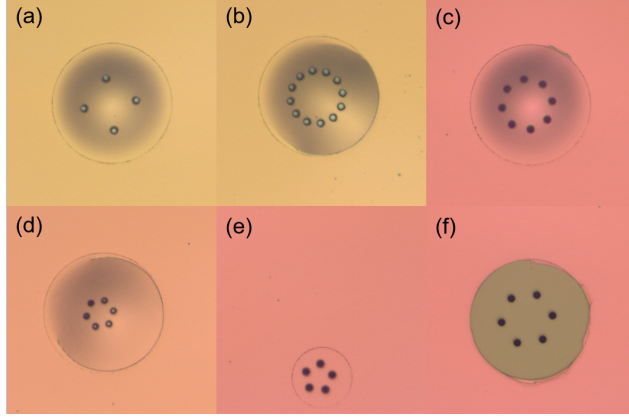


Figure 3.8: Microscope images of select individual membranes with the mask applied, prior to Cobra etching. (a-d) 100 μm diameter domes with various etch hole configurations. (e) A 50 μm diameter unbuckled dome. (f) The original site of a 100 μm dome that has been popped off due to sonication.

riods are required to achieve a desired reflectance. On this sample, a 4-period 1550 nm Si/SiO₂ top mirror was sputtered, starting with Si. Note that the mirror capping layer was chosen as SiO₂ in order to prevent the etching of the top surface by XeF₂.

The sample was buckled on the hotplate, according to the standard procedure. It was observed that while the 100 μm diameter domes buckled, the 50 μm diameter domes did not (which constituted half of the wafer.) This was likely due to uneven heating of the substrate on the hotplate. Zygo optical profilometer scans indicated a typical height for the 100 μm diameter buckled domes on this wafer was approximately 2.4 μm . The wafer was then cleaved into four similar sized pieces. One of these pieces was used for characterizing the Cobra etch process, in order to ensure that the appropriate etch time was used. With 18 minutes of the Cobra etching, it was determined that the etching did indeed impact the membrane surface below as desired.

While the use of Si as a top mirror material allows for the creation of a more reflective mirror with fewer periods, its dual use as the sacrificial layer means that there is the danger that the top mirror may also be unintentionally etched. The side walls of the etch holes expose the top mirror Si layers to attack by the XeF₂. The thickness of each Si layer in the mirror was ~ 100 nm, meaning significantly less area was exposed relative to the the size of the bottom of the

hole. Although the XeF_2 etch is isotropic, it was hoped that the sacrificial layer could be sufficiently removed before the top mirror was etched enough to interfere with the device operation.

In an attempt to mitigate this issue, on some samples immediately prior to the XeF_2 etch, the devices were subject to a single DRIE Bosch pulse in the Alcatel system. The Bosch process is a cyclic etching process, which cycles between an isotropic PTFE deposition (identical to the materials used here as a low adhesion layer,) and an RIE etch to allow for the formation of deep and straight holes. PTFE on side walls acts as an etch barrier during the RIE step. In the same respect, PTFE might be used to protect the side walls during the XeF_2 etch, since PTFE is compatible and similarly inert to XeF_2 . It was hoped that a single Bosch cycle would allow for the coating of all surfaces with PTFE, and then the brief directional RIE etch would clear the bottom of the access hole for it to be exposed to the subsequent XeF_2 etch.

A single Bosch pulse was done on one test piece, and it was loaded for XeF_2 etching alongside a similar piece that received no Bosch pulse. Ten cycles of etching were performed, and the pieces were removed and inspected. There was etching that occurred, but either in the form of small circles around the etch holes or large ‘flower’ patterns that occurred at the periphery of the domes as well. There was no noticeable difference between the piece that received the Bosch step and the one that did not. The samples were loaded for an additional 10 cycles of XeF_2 and inspected again. The small holes around the etch holes did not expand in size, but the flowering at the edges of the domes did spread. This is shown in Figure 3.9. It is possible that a single Bosch pulse did not provide a sufficiently thick and contiguous PTFE surface. The 18 minute Cobra etch used here likely reached the substrate based on an estimation using the bulk etch rates. This meant that the largest section of Si available for etching was the substrate layer at the bottom. This created the multicolour flower patterns that were visible since what was being viewed overhead was a composite interference pattern caused by all of these layers being etched out. The inconsistent ‘texture’ of the etching was likely because the substrate had no etch stop to set a floor for etching in the out-of-plane

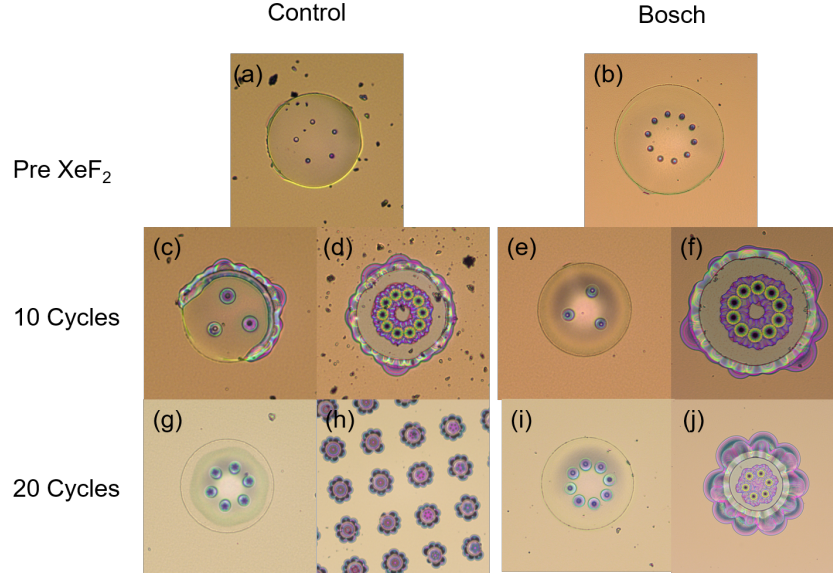


Figure 3.9: An overview of results from the XeF_2 etching trials with and without an additional pre- XeF_2 Bosch pulse. Images are not to scale relative to one another, and are microscope images of randomly chosen domes from the sample at each step.

direction, causing an uneven bottom surface (*e.g.* the pink surface in Fig. 3.9). We concluded that if the etching was in fact occurring on the still-present domes, we should see a similar set of flower patterns. However, it seems that when comparing these internal dome membrane release patterns with those of the much more transparent $\text{Ta}_2\text{O}_5/\text{SiO}_2$ top mirror in the next section, they are of comparable geometry. For instance the outlines, such as those seen in Figure 3.10, may indeed be devices with released membranes (without significant etching of the top mirror from the access hole side walls.) Owing to the difficulty of seeing inside of the integrated devices with the opaque Si/SiO_2 top mirror, it is unknown what a ‘properly’ released membrane would look like from the outside. On some samples (Fig. 3.10) there is evidence of a faint discolouring underneath the dome pattern that appears to reflect the geometry of the holes.

Let the access holes be approximated as a circle of radius r , so that the exposed area of the sacrificial layer is then πr^2 . The top mirror consists of 4 layers (alternated with SiO_2) of Si , each approximately 103 nm in thickness. This is a total of 412 nm of side wall thickness composed of Si . The circumference of




Figure 3.10: Four domes from the XeF_2 20 pulsed no-Bosch sample. Contrast has been enhanced to show the internal patterns visible within the confines of the dome, that reflect the hole geometry that indicate the presence of a feature underneath that could be the released membrane. Optical profilometer scans indicate this is distinct from the contours of the dome, which is still circular.

the access hole is $2\pi r$, so the total sidewall area of the top mirror consisting of Si is therefore $2\pi r \times 412 \text{ nm}$. Let the hole radius be approximately $2 \mu\text{m}$, which is typical for these holes. The approximate exposed area of the sacrificial layer is then $13 \mu\text{m}^2$, while the exposed area of the Si in the mirror is $5 \mu\text{m}^2$. There is over twice the area for the XeF_2 to access in the preliminary stages of the etch. However, once the etching reaches the bottom of the sacrificial layer and proceeds laterally, the sacrificial layer is only 200 nm thick and more etching should actually occur of the mirror in terms of exposed surface area. This cannot explain why it seems that the top mirror layers stayed mostly intact.

Further investigation revealed that, much as was the intention with the addition of a single cycle of the Bosch etch to protect the access hole side walls with a fluoropolymer, the fluorine based Cobra RIE etch is known to produce a surface protecting fluoropolymer as a by-product of its reaction. Fluoropolymer contamination is a known issue within RIE processes [59]. The production of this polymer helps contribute to the downwards anisotropy of the etch, which is desired in our process. Preliminary etch hole tests were followed up by SEM scans, in an attempt to view the inside of the holes created. Images

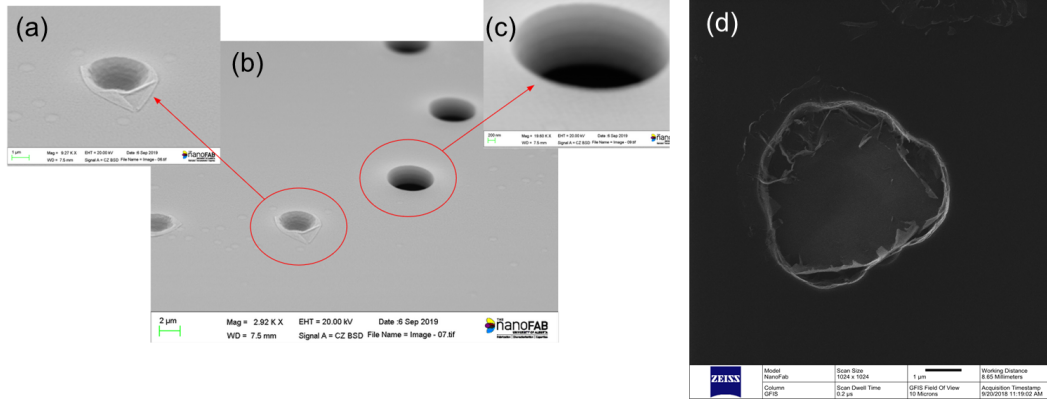


Figure 3.11: (a-c) SEM Images of holes from preliminary etching tests, possibly showing a PTFE-like coating flaked around some holes. (d) A helium ion microscope (HiM) image of an etch hole from the same sample. Scans with the HiM, unlike the SEM, did not require a gold coating and higher detail of the flake is visible.

taken from these scans (Fig. 3.11) show large amounts of flake-like deposits, around the top of the holes and even lining the insides in some cases. This flaky layer, based on past observations with the fluoropolymer low adhesion layers, is similar in appearance to thin film PTFE. Similar observations were also made in the Helium ion microscope (Fig. 3.11(d)) without the gold coating that the SEM required.

3.2.3 Second Generation: $\text{Ta}_2\text{O}_5/\text{SiO}_2$ Based Top Mirror

The complications of the top mirror containing silicon encouraged the development of a Si-free top mirror, using Ta_2O_5 as a high index material instead. Processing was identical to that described for the first generation devices, except with the top mirror as a 4.5-period 1550 nm QWS $\text{Ta}_2\text{O}_5/\text{SiO}_2$ mirror (starting and ending with Ta_2O_5 .) The wafer was buckled on the hotplate. This sample buckled unprecedentedly early - after approximately 3 minutes and with the hotplate setpoint at 300 °C, the sample read between 200 °C to 300 °C and was removed due to large amount of the domes popping off in some sections. Inconsistent heating resulted in large portions of this wafer not buckling, while others popped off and were destroyed.




Figure 3.12: A closeup of a ‘multilayer’ XeF_2 etch, with this sample clearly being over-etched at 24 minutes in the Cobra system, with a plausible explanation that various layers of Si in the bottom mirror are being etched, and are causing the visible multiple layers and rings.

Since the top mirror thickness was different relative to the previous Si/SiO₂ devices, etch tests had to be redone in order to determine the optimal RIE etch time for the access holes. The wafer was split into smaller dies, and these were subject to 10, 15, 18, 19, and 24 minutes of etching with the same recipe. The samples that were etched for 10 and 15 minutes showed no obvious visual changes after XeF_2 etching, indicating the Cobra etch did not reach the sacrificial Si layer. The 19 and 24 minute samples exhibited ringed features, as shown in Fig. 3.13 for three different domes after different numbers of XeF_2 cycles. These rings are believed to be due to the Si layers in the bottom mirror being removed, as shown in the schematic in Figure 3.12. If the XeF_2 is able to reach these layers, it will etch them just as it does the sacrificial layer. Considering the etch rate of Si (Table 3.3), only a minute would be required to go through the 200 nm sacrificial layer and into the bottom mirror.

The 18 minute sample was found to show uniform etching and creation of membrane features. On these devices, a large portion of the domes intended to be 50 μm in diameter tended to ‘over-buckle’ (visible on all samples shown in Fig. 3.14). The top mirror can be seen extending in one direction outside of the PTFE patterning, and changing the shape of the buckle. The buckle itself is still roughly circular. Curiously, throughout the wafer, all of the over-buckling tends to occur in the same direction (towards the ‘top-right’, with reference to the wafer rotated with the major flat at the bottom.) This over-buckling




Figure 3.13: Various membranes, tracked at different stages of the etching process, for a section of the wafer that received a 24 minute etch in the Cobra as a test. The device in the middle row has had its dome unintentionally been damaged and removed prior to the RIE etch, with the top surface being the membrane.

was seen on the sample immediately after buckling, so it is not related to the patterning of the etch holes. It is likely this issue was an anomaly with this particular wafer due to factors such as cleanliness. The intended patterned diameter of the domes may also affect the degree of overbuckling that occurs. Larger diameter domes ($\sim 100 \mu\text{m}$) did not noticeable overbuckle.

Further microscope study of these devices revealed some interesting lines visible when the focus is on the membrane plane. Two photographs of the same dome, taken at slightly different foci are shown in Figure 3.15, show highly symmetric line patterns that appear to reflect the geometry of the holes. It is unclear whether this is physically the shape of the membrane, or some sort of optical illusion. One possible explanation is these lines are due to stress in the membrane. An effect called photoelasticity is known to create visual contours according to stress, which are caused by birefringence in the material. Appreciable photoelasticity has been observed in similarly patterned silicon nitride membranes [78]. A detailed analysis of this is left to further work, however the symmetry of these lines is somewhat intuitive for a circular membrane and is likely indicates the presence of an unbroken,




Figure 3.14: Various membranes devices from the 18 minute sample. While imperfect in many ways, they distinctly do not show the multilayer contours from the 24 minute sample (Fig. 3.13).

stressed membrane. It may also be interesting to see what effect, if any, using a polarized light for microscope imaging could have on these lines.

3.3 Characterization

A large number of well-aligned devices were identified, and one representative ‘case-study’ device (shown in Fig. 3.17(a),) is studied in detail in this section. This device combines a dome of base diameter $74\ \mu\text{m}$ and peak (buckled mirror) height $\sim 3.2\ \mu\text{m}$ with an etch array pattern comprising 5 equally spaced holes (each $4\ \mu\text{m}$ in diameter on the mask, see Appendix E,) placed on a ‘ring’ of diameter $\sim 30\ \mu\text{m}$. Due to the isotropic nature of the XeF_2 etch, this hole pattern resulted in a ‘flower-shaped’ suspended membrane with 5-fold symmetry. A confocal microscope (imaging at $408\ \text{nm}$) was used to characterize many devices on the wafer, and this study revealed consistently intact membranes at the expected depth (*e.g.* as shown in Figs. 3.17 (b) and (c)). Note that the membranes lie $\sim 3.3\ \mu\text{m}$ below the planar surface (outside the buckled regions) of the wafer, corresponding to the thickness of the 7.5-period $\text{Ta}_2\text{O}_5/\text{SiO}_2$ mirror. The confocal microscope was further used to measure the profile of the




Figure 3.15: Microscope images of the same dome, before any XeF_2 etching but with the mask removed, at two different focuses. Focus on the membrane plane (b) reveals line features reflecting the geometry of the holes, enlarged in (c). This was present on virtually all hole configurations. (d,e) Microscope images of a closely spaced hole configuration at two different close focuses, with a curious line pattern visible on the membrane plane that connects all holes to one another. (f) A particularly bold and symmetrical 5 point star pattern.




Figure 3.16: (a) Schematic cross-sectional view (not to scale) of a buckled dome microcavity with an embedded free-standing Si_3N_4 membrane. For the devices discussed here, the sacrificial layer is ~ 200 nm thick, and sets the spacing between the membrane and the bottom mirror. (b) Microscope image of a completed membrane-in-cavity device. The green ‘flower’ shape is the suspended membrane, and the circular interference fringes arise from the buckled profile of the upper mirror. As is evident, many of these first-generation devices suffered from imperfect alignment between the etch hole pattern (thus the suspended membrane) and the buckled dome microcavity.




Figure 3.17: Images of the ‘case-study’ MIM cavity described in the main text. (a) Standard microscope image; the concentric interference fringes arise from the profile of the buckled upper mirror and the ‘flower-shaped’ region in the centre is the sacrificial etched cavity and suspended membrane. (b) Confocal microscope image with the focus set at the top surface of the upper mirror (outside the buckled regions). (c) Confocal image with the focus set $\sim 3.3 \mu\text{m}$ below the top surface, where the membrane layer is located. (d) Profile of the buckled dome (upper mirror of the cavity) as determined using the confocal microscope (blue,) and fit to profile predicted for a circular delamination buckle (red.)

upper (buckled) mirror for many domes, and these profiles were confirmed to be consistent with elastic buckling theory (Equation 2.29.) Fig. 3.17(d) shows the measured profile for the case study dome alongside to the predicted profile (with $\delta = 3.2 \mu\text{m}$ and $a = 37 \mu\text{m}$). The excellent agreement indicates that neither the membrane nor the etch holes had a significant impact on the mirror shape, consistent with previous results [39].

3.3.1 Simulations

Figure 3.18 shows transfer-matrix predictions for a planar-equivalent model of the ‘case-study’ dome from Fig. 3.17. Here, a 150 nm thick membrane is separated from the bottom and top mirrors by 200 nm and 3.2 μm thick air gaps, respectively. The SiN membrane was assigned a real refractive index of 2, and the other materials were modelled using dispersion relations (see Appendix F.) As shown in Fig. 3(a), a single resonant mode with FWHM linewidth $\sim 0.29 \text{ nm}$ ($Q \sim 5000$) and centred at $\sim 1550 \text{ nm}$ is predicted within the stop-bands of the mirrors.

In these devices, the membrane is not near the middle of the cavity, unlike what is typical in other MIM models. As shown in Fig. 3.18(b), the reso-

nant mode can be viewed as having longitudinal order of 4 or 5, depending on whether the SiN membrane (and underlying etched air gap) is viewed as part of the cavity or as part of the bottom mirror. For example, if one views it as part of the cavity, then the effective cavity length is $L \sim 5 \cdot \lambda \sim 3.9 \mu\text{m}$ which is the sum of the air gaps, the optical thickness of the SiN membrane ($\sim 0.3 \mu\text{m}$), and the penetration depths into the upper and lower Bragg mirrors. For conventional MIM systems, with a membrane assumed to be much thinner than the overall cavity length and positioned near the middle of the cavity, a sinusoidal variation of the cavity resonance frequency with membrane position can be predicted [3]. From these expressions, a maximum frequency pull parameter (for optimal membrane placement mid-way between field nodal and anti-nodal positions) may be predicted using Equation 2.39 [57]. In the case-study device above, $|r_m| \sim 0.57$ and so $G_{\text{max}}/2\pi \sim 55 \text{ GHz/nm}$. The structure here deviates significantly from these assumptions, in that the membrane is not centrally placed and its thickness is not vanishingly small compared to the cavity length. In fact, it more closely resembles the ‘membrane-at-the-end’ (MATE) system [25], for which G_{max} can actually exceed the previous estimate. In any case, a more exact treatment is afforded by direct transfer-matrix solutions obtained for varying membrane position. As shown in Fig. 3.18(b), the ‘resting’ position of the membrane is near a field anti-nodal position. Thus, we can expect the pull parameter to be lower than the optimal value above, as verified by the plots shown in Figs. 3.18(c) and (d), which predict $G/2\pi \sim 25 \text{ GHz/nm}$ or $G_\lambda \sim d\lambda/dz \sim 0.185 \text{ nm/nm}$ for the case-study device.

The single-photon optomechanical coupling strength g_0 can also be predicted (Equation 2.53.) Using the values $m_{\text{eff}} \sim 90 \text{ pg}$ and $\Omega_M \sim 10 \text{ MHz}$ discussed below gives $x_{\text{ZPF}} \sim 2.5 \text{ fm}$ and $g_0/2\pi \sim 0.1 \text{ MHz}$ for our case study device. This predicted g_0 value is already comparable to the highest values reported in the literature [4, 28], and could likely be increased in future process iterations as discussed further on.

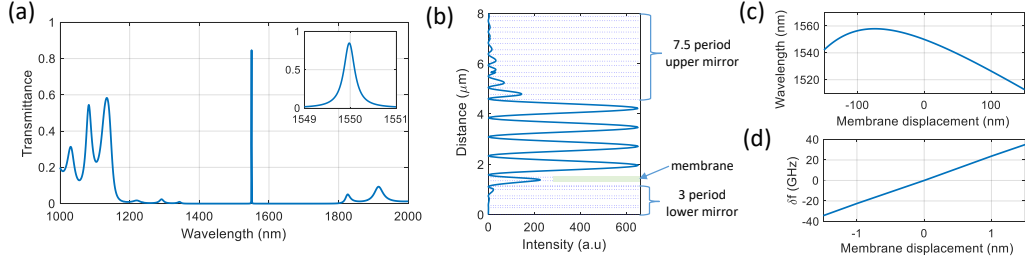


Figure 3.18: Transfer-matrix predictions for the planar-equivalent model of the cavity shown in Fig. 3.17. (a) Transmittance spectrum showing a resonant mode at ~ 1550 nm with line width ~ 0.29 nm ($Q \sim 5000$). (b) Field intensity ($E \cdot E^*$) profile for the resonant mode from part a. The SiN membrane essentially acts as an additional layer in the bottom mirror, and is roughly centred on a field anti-node. (c) Variation in resonant wavelength with change in membrane position. (d) Change in resonant frequency versus membrane displacement. The predicted linear optomechanical coupling strength is $G/2\pi \sim 25$ GHz/nm.

3.3.2 Experimental Setup

A schematic of the experiment used for measurement is shown in Fig. 3.20. Light is sent from a fibre coupled laser (Santec TSL-710) through a variable attenuator (Oz Optics BB-500-11-1300/1550-9/125-S-50-3A3A-3-1,) and then into a broadband reflective fibre collimator (Thorlabs RC02FC-P01.) The free space collimated beam propagates ~ 3 cm through a 20x microscope objective (NA=0.40,) which is focused onto the sample with a spot diameter of $D \sim 3.5$ μm . This bottom ‘collimator assembly’ is fixed together and mounted on an XYZ stage, in addition to a two axis tilt-tip stage. The sample itself sits on an independent XYZ and tilt-tip stage (Thorlabs XYF1.) This platform can accommodate samples of varying size. A thin aluminum plate (75 mm x 26 mm) with a hole in its centre has been machined to support the sample while allowing light to pass from underneath unobstructed. Collecting the transmitted light above the sample is a set of interchangeable 10x and 50x infinity corrected microscope objectives (Zeiss Epiplan-Neofluar.) These are mounted into a switchable turret to allow for quick swapping from a higher to a lower magnification. These objectives focus on the sample and send light upwards towards a 90:10 1550 nm beam splitter (Thorlabs BS039 mounted in Thor-




Figure 3.19: A photograph of the experimental setup used to characterize the devices. The experiment is shown in detail in the schematic of Figure 3.20.

labs CCM1-4ER.) 10% of the light is sent vertically upwards towards a digital infrared camera (Raptor Photonics Ninox 640 NX1.7-VS-CL-640) where it is focused as an image through a 200 mm tube lens (Thorlabs TTL200-S8.) This allows for imaging and active alignment of the sample. The remaining 90% of the light is sent horizontally towards multimode fibre through a large beam air spaced doublet collimator (Thorlabs F810FC-1550.) This entire ‘collection assembly’ is mounted vertically onto a rail, and the rail is attached to an XY stage. Light collected by the collimator and multimode fibre is incident on to a high speed photodetector (Resolved Instruments DPD80,) interfaced to the computer. The photodetector includes all necessary electronics, including an analog to digital converter, that allows for it to be operated over a USB 3.0 interface. Measurements may be taken up to 80 Msamples/s, and these are further post-processed into the frequency domain to allow for spectral analysis.

Measured coupling losses between the various stages of the experimental setup are given in Table 3.6. These losses can include coupling losses between fibres and equipment, inefficiencies associated with reflections and aberrations




Figure 3.20: A schematic block diagram of the experimental setup. Details of each component are discussed in the main text.

Location	Fraction of Laser Power	Measurement Device
Before Collimator	81%	IS
Sample Stage	79%	FS
After 90% Beam-splitter	24%	FS
Out of MMF, before PD	1%	IS

Table 3.6: Measured throughputs of the experimental setup, with no sample in the path. Measurement Device is indicated as either a fibre-coupled integrating sphere (IS) or free-space photodetector (FS).

in the microscope objectives (due to them being optimized for visible operation,) and loss due to misalignment of the multimode fibre with the input beam, due to limitations on the sensitivity of the XY stage for the collection assembly and the small core diameter of the multimode fibre of $\sim 50 \mu\text{m}$. Throughputs were characterized by inserting either a fibre-coupled integrating sphere photodetector or a free-space detector where appropriate at particular points in the experimental setup.

3.3.3 Experimental Measurements

Spectral scans were obtained by focusing the optical mode onto the camera, and summing the intensity of all pixels within the region of the modes while taking care to avoid saturation. Comparative measurements were made between the case study device and an un-etched ‘control’ dome of similar size (peak height $\sim 3.0\ \mu\text{m}$), with results shown in Fig. 3.21. Note that the low-amplitude, periodic ripple in both scans is due to interference effects between the surfaces of the silicon substrate, and could be eliminated in future work by the addition of an anti-reflection coating on the bottom surface of the substrate. Aside from this, the scans are consistent with theoretical predictions for a half-symmetric Fabry-Pérot cavity. For example, the upper mirror has effective radius of curvature $R_{C0} \sim 100\ \mu\text{m}$ here (see Fig. 3.17(d)), suggesting a fundamental mode spot size radius $w_0 \sim (\lambda/\pi)^{1/2} \cdot (L \cdot R_{C0})^{1/4} \sim 3.1\ \mu\text{m}$ and transverse mode spacing $\Delta\lambda_T \sim \lambda^3/(2 \cdot \pi^2 \cdot w_0^2) \sim 20\ \text{nm}$ both of which are in good agreement with the measured results. Higher-order transverse modes exhibit non-degeneracy (multiple sub-peaks), especially for the etched dome, which can be attributed to slight deviations from spherical symmetry for the buckled mirror. Consistent with this, while the control dome exhibited clear Laguerre-Gaussian (LG) modes (Fig. 3.21 insets), the ‘case study’ dome tended to exhibit Hermite-Gaussian (HG) modes.

For the case study dome, the measured linewidth of the fundamental TE_{00} resonance is $\sim 0.3\ \text{nm}$ ($Q \sim 5000$), in excellent agreement with the transfer-matrix predictions (Fig. 3.18). Similar results were obtained on a large number of other devices. Note that this corresponds to a finesse $\mathcal{F} \sim 10^3$, similar to values reported in earlier work [36]. Since membrane-induced scattering/absorption losses were neglected in the transfer matrix treatment, we can conclude that this finesse is limited mainly by the low period count of the mirrors (*i.e.* reflectance-limited finesse). This lends further evidence to the good optical quality of the surfaces left behind by the XeF_2 sacrificial etch, consistent with the microscope images in Fig. 3.17. Significantly increased finesse should be possible with improved Bragg mirrors.

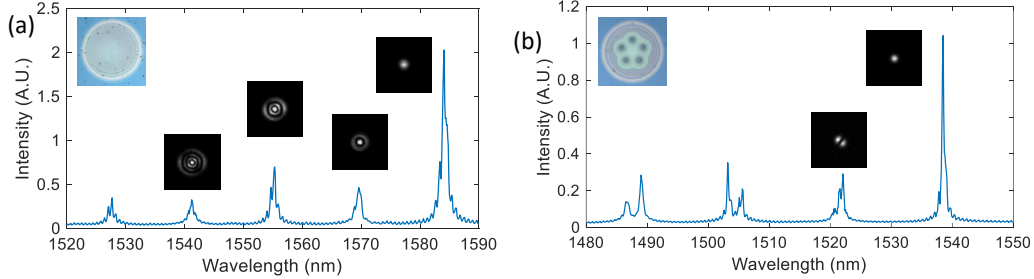


Figure 3.21: (a) Spectral transmission scan for a dome without etch holes, thus with no released membrane. The insets show mode-field intensity profiles imaged at the resonant wavelengths indicated. (b) Spectral transmission scan for the case study dome from Fig. 3.17. The insets show selected mode-field intensity profiles, as in part a.

As was explored in the Chapter 2, vibrations of an internal membrane are expected to cause changes in the cavity resonant wavelength proportional to its position. By detuning the laser slightly from the optical mode, a so-called ‘tune-to-slope’ technique can be used to map the time domain cavity resonance to changes in cavity transmission. This technique has been used successfully in the past by Bitafaran [48] to characterize the thermal resonances of the buckled top mirror, which induces changes in cavity resonant wavelength by directly modulating the cavity length. In all measurements shown here, the laser is always red-detuned from the cavity resonance. PSD’s measured from several different domes are given in Figure 3.22, including that of a dome without any holes as a ‘control’ for the study. Top mirror resonances are observed at roughly 7.2 MHz and 20.6 MHz. The fundamental resonance at 7.2 MHz is well predicted at this diameter by the plate and shell models (see Fig. 2.9.) Very similar spectra were observed for a large number of membrane-free domes. As shown in Fig. 3.22(b), dramatically different vibrational spectra were observed for cavities containing suspended SiN membranes. For the case study dome, and many other similar devices, the vibrational features associated with the buckled mirror are still present as expected, superimposed on several other resonant lines that we attribute to the thermal vibrational motion of the SiN membrane. For example, if the flower-shaped membrane is treated very approximately as a circular membrane of diameter $\sim 40\ \mu\text{m}$ (and the etch holes

are ignored), analytical expressions (see Equation 2.36) predict a fundamental vibrational frequency ~ 10 MHz given the tensile stress (~ 770 MPa) mentioned above. The large feature at ~ 10.5 MHz in Fig. 3.22(b) is then possibly attributable to a fundamental vibrational resonance (analogous to the (0,1) mode of a circular membrane, Fig. 2.10), and that the higher-frequency resonances could similarly be attributed to higher-order vibrational modes (again analogous to (m,n) modes of a circular membrane). Moreover, the stronger resonance peaks at ~ 22 MHz and ~ 32 MHz lie at approximately the expected frequencies for the (0,2) and (0,3) modes of a circular membrane having the aforementioned properties. Since only the $(0,n)$ vibrational modes lack a nodal point at the centre of the membrane (Fig. 2.10), the higher amplitudes of these peaks can be explained by the increased overlap between these mechanical modes and the fundamental optical mode field [45], which is nearly centred on the membrane for this cavity. Finally, it is also worth noting that the quality factor ($Q_M \sim 50$) of these membrane-attributed vibrational modes is consistent with values reported for similar SiN membrane resonators at atmospheric pressure [54], with Q_M limited mainly by acoustic radiation into the surrounding air medium.

To gain further insight, mechanical properties were modelled using a finite element analysis software package (COMSOL.) The suspended SiN feature was modelled as a membrane, with an assumed shape (including holes) based on the confocal microscope image shown in Fig. 3.17(c). Clamped boundary conditions were applied to membrane edges, while the hole peripheries were treated as free boundaries. For the membrane, material properties used are given in Table 3.7.

Squeeze film damping effects [54] must be considered in order to obtain the observed resonances at the experimentally measured tension. Trapped air below the membrane becomes compressed due to the vibrations of the membrane. As described by Southworth [54], the high frequency of our membranes places them in the so-called ‘elastic damping’ regime. Within this approximation at high frequencies, the effect of the trapped gas on the membrane may be approximated as a ‘gas spring’:




Figure 3.22: FFT spectra extracted from ‘tuned-to-slope’ measurements for (a) a ‘regular’ cavity with no etch holes, and thus no released membrane, and (b) the case-study cavity shown in Fig. 3.17. Both were captured with time-averaged power at the detector $P_0 \sim 23 \mu\text{W}$. The resonances in part a are attributable to vibrational modes of the buckled upper mirror of the cavity, and are also present in part b. The additional resonances in part b are attributed to the vibrational modes of the released SiN membrane inside the cavity.

Parameter		Value
Outer Diameter (μm)	D	50
Mass Density (kg/m^3)	ρ	3100
In-Plane Stress (MPa)	σ	770
Membrane Thickness (nm)	t_m	150
Young’s Modulus (GPa)	E	270
Poisson’s Ratio	ν	0.27

Table 3.7: Material properties of silicon nitride used in the membrane simulation.

$$k_g = \frac{P_a A}{d_{\text{gap}}}, \quad (3.1)$$

where P_a is the ambient pressure (atmospheric,) A is the area of the membrane, and d_{gap} is the distance from the membrane to the layer below (approximately 200 nm.) The unusual shape of our membrane means the spring is better represented as an ‘areal’ spring. This is simply

$$k_A = \frac{k_g}{A} = \frac{P_a}{d_{\text{gap}}}, \quad (3.2)$$

in units of N/m³. Approximating the area of the non-circular membrane, an estimate of $k_A \sim 5 \times 10^{11}$ N/m³ is obtained, or over the area of a 40 μm diameter simple circular membrane $k_g \sim 630$ N/m. This spring is on the same order as the effective spring of the membrane alone ($k_{\text{eff},n} = m_{\text{eff},n} \cdot (2\pi f_{M,n})^2 \sim 360$ N/m, for $f_{M,n} \sim 10$ MHz, and $m_{\text{eff},n} \sim 90$ pg) [53]. The gas spring should therefore have an appreciable effect on the resonances of these membranes. The approximate areal spring was introduced into the COMSOL model as a spring distributed uniformly over the SiN membrane. By defining the spring as areal, it allows mode shapes of varying displacement profiles to be proportionally impacted. As a comparison to the simple gas spring approximation, a second set of simulations were run in which the areal spring was removed and instead a fluidic air layer was introduced below the membrane within COMSOL, which allows direct numerical analysis of the physics of the fluid layer below the membrane in a coupled ‘multi-physics’ simulation [79]. As shown in Fig. 3.23(a), both of these approaches provided very good fits between the predicted $(0,n)$ modal frequencies and the three strongest vibrational resonances in the experimental scan, consistent with the discussion above. In this figure, simulated resonance frequencies for the areal spring model (red dashed line) and the coupled membrane-fluid simulation (black dotted line) nearly overlap. The lower-amplitude resonances in the experimental scan can be associated with (m,n) vibrational modes where $m \neq 0$, all of which have a nodal point at the membrane centre and thus exhibit lower coupling with the fundamental optical mode. The COMSOL model reproduced the number of observed

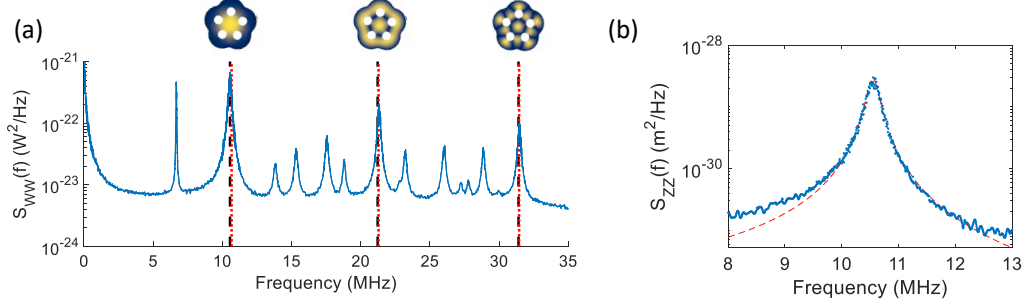


Figure 3.23: (a) Comparison between the experimental vibrational spectra and predicted (0,1), (0,2), and (0,3) resonance frequencies obtained using a finite-element numerical solver (COMSOL). The red dashed lines indicate frequencies predicted using an areal spring model for trapped gas beneath the membrane, and the black dotted lines indicate frequencies predicted using the COMSOL fluid physics modules for this trapped gas. The insets show the predicted vibrational mode shapes. (b) The curves show the results of a thermo-mechanical fitting procedure (see main text), performed on the experimental vibrational resonance at ~ 10.5 MHz (*i.e.* the (0,1) mode). The displacement spectral density extracted from the experimental data (blue solid line) is compared to that predicted for a damped harmonic oscillator subject to Brownian motion (red dashed line).

modes (for example, four additional modes between (0,1) and (0,2)) but there was less agreement between simulated and observed frequencies in these cases. We believe this is due to significant overlap between the etch holes and the vibrational profiles for these modes (*i.e.* these modes ‘live’ closer to the etch holes.) The simple COMSOL model employed does not capture possible local variations in the membrane stress near the etch holes [30], nor does it capture likely reductions in the local ‘gas spring’ effects near the etch holes due to a well-known gas escape mechanism [80]. Nevertheless, the excellent high-level agreement between the models and the experimental data suggests that the most critical details are understood.

Further insight into the mechanical vibrations of the membrane can be gained by performing a ‘thermomechanical calibration’ of the fundamental mechanical resonance [53]. As discussed in Chapter 2, the measured PSD may be modelled through Equation 2.52, using the displacement Lorentzian function (Equation 2.51.) Fitting Equations 2.51 and 2.52 to the data shown in Fig. 3.23(a), yielded $S_{WW}^W \sim 5 \times 10^{-24}$ W²/Hz, $f_{M,n} \sim 10.55$ MHz, $Q_{M,n} \sim$

48, and $\alpha \sim 2 \times 10^7 \text{ W}^2/\text{m}^2$. The result of this fit is shown in Fig. 3.23(b), where $S_{zz}(f)$ has been plotted, using the extracted S_{WW}^{W} and α to rescale the experimental data from Fig. 3.23(a). Note that the disagreement in the ‘tail’ regions can be attributed to adjacent resonance lines not accounted for by the simple Lorentzian line-shape function.

For our system specifically, α can also be predicted directly using Equation 2.47. For the measurements shown, the laser was adjusted to give maximum observed signal, and thus we assume $S \sim S_{\text{max}} \sim 1.3/(0.3 \text{ nm}) \sim 4.3 \times 10^9 \text{ m}^{-1}$. Furthermore, using a Gaussian beam profile for the fundamental optical mode (with $w_0 \sim 3.1 \text{ }\mu\text{m}$) and the numerically predicted mode-field profile for the (0,1) vibrational mode (see inset of Fig. 3.23(a)), we estimated $\eta \sim 0.98$. Combining these with the stated value of P_0 from above and the simulated $G \sim 0.185$ (Fig. 3.18), Equation 2.47 then predicts $\alpha \sim 3.9 \times 10^7 \text{ W}^2/\text{m}^2$, in very good agreement with the value extracted from the thermomechanical calibration above. Residual disagreement is likely due to uncertainty in S , since as mentioned our laser is not locked relative to the cavity resonance, and overestimation of η , since even for the ‘case study’ dome the membrane is somewhat misaligned with the cavity. For instance, calculations show using Equation 2.41 that a misalignment of $5 \text{ }\mu\text{m}$ in both the x and y directions can potentially cause η to drop down to ~ 0.81 .

The temperature of the vibrational mode can also be predicted from Fig. 3.23(b) as described in Ref. [3]:

$$T_{\text{eff}} = \frac{m_{\text{eff},n} \cdot \Omega_{\text{M},n}}{k_{\text{B}}} \int df S_{zz}(f), \quad (3.3)$$

using the ‘experimental’ S_{zz} curve. This yields $T_{\text{eff}} \sim 300 \text{ K}$, as expected for our room-temperature measurements given that optical sideband cooling effects are not expected to be appreciable with the current experimental setup.

From another point of view, the analysis above is essentially equivalent to comparing the experimentally observed and theoretically predicted optomechanical coupling coefficients (G) for our system. The reasonable agreement (*i.e.* within a factor of 2) obtained provides confidence that the observed vi-

brational features are in fact attributable to the motion of a suspended SiN membrane inside our buckled cavities, and that our models have captured the essential physical details of the system. This should in turn provide a good basis for the optimization of fabrication processes and cavity parameters in future work.

3.3.4 Discussion

These measurements provide evidence of the first truly integrated and scalable MIM systems in the literature. These devices also have the potential to exhibit values for g that are among some of the highest reported. This is primarily owing to the small cavity length of these systems. The trade-off however is that with small, integrated cavities it becomes extremely difficult to obtain cavity finesses that would allow operation in the fully resolved side band regime $\kappa \ll \Omega_M$, where the retarded optical cavity response relative to the membrane motion would enable the possibility of achieving ground-state cooling. The short length of our cavities places extreme requirements on the cavity finesse in order to achieve low κ . In the literature, $\mathcal{F} \sim 10^6$ has been achieved using high-period-count Ta₂O₅/SiO₂ mirrors, and our group has recently proposed that hydrogenated amorphous silicon might enable similar performance with far fewer layers [64]. Assuming $\mathcal{F} = 10^6$ as a best-case scenario projects to $Q_M^* \sim 5 \times 10^6$ and $\kappa^* \sim 40$ MHz for this cavity. That being said, alternative cooling schemes have also been proposed such as optical feedback cooling, which actually require the non-resolved side band regime and weak coupling conditions present here [81, 82, 83].

The effective temperature needed for ground-state cooling of a mechanical vibrational mode depends on frequency, as quantified by the time-averaged thermal occupation number, $N_{\text{TH}} = k_B T_{\text{eff}} / \hbar \Omega_M$, where T_{eff} is the effective temperature of the mechanical mode of interest. In order for quantum fluctuations to dominate over thermal noise, $N_{\text{TH}} \ll 1$ is required, and this is obviously achieved at higher T_{eff} for higher-frequency oscillators. Lithographic alignment enables us to embed a smaller membrane in our cavities than is typically used in hybrid systems [26]. Moreover, a membrane ‘pad’ [29, 84] as

small as $\sim 10 \mu\text{m}$ in diameter should be sufficiently large given the small waist radius of the optical modes in our cavities. This projects to fundamental mode resonance frequencies as high as $f_M^* = \Omega_M^*/2\pi \sim 50 \text{ MHz}$. Another advantage of this size reduction is the relative sparsity of vibrational modes, which should make it easier to isolate interactions between light and a particular vibrational mode of interest.

Regardless of the choice of cooling scheme, radiation-pressure back-action effects can be used to approach ground-state cooling in a room-temperature environment, provided that $Q_M > N_{\text{TH,room}} = k_B T_{\text{room}}/\hbar\Omega_M$ [4, 29, 45]. Using f_M^* from above implies that this condition would be satisfied for $Q_M > 10^5$. Projections about mechanical quality factor are admittedly speculative until measurements in vacuum can be completed. Nevertheless, we reiterate that similar SiN membranes have been reported to have room-temperature $Q_M > 10^6$ [3, 26, 32, 45]. Moreover, we believe that our surface-micromachining process, with judicious design of etch hole sizes and patterns, might be used to implement either trampoline-style [30, 29] or phononic-crystal-isolated [31, 84] central pads, and both of these strategies have reportedly enabled $Q_M > 10^7$. Thus, we conservatively project that $Q_M^* = 10^6$ and $\Gamma_M^*/2\pi = f_M^*/Q_M^* = 50 \text{ Hz}$ are feasible goals.

Some of the most powerful envisioned schemes in optomechanics are predicated on non-linear interactions at the single photon level [4], which become significant in the strong-coupling regime defined by $g_0 > \kappa, \Gamma_M$. Other interesting non-linear effects are manifested when $g_0 > \Omega_M$ [28]. For our cavities, assuming reductions in membrane size/mass as mentioned above, we can project that $x_{\text{ZPF}} \sim 5 \text{ fm}$ and $g_0^* \sim 0.5 \text{ MHz}$ are maximum feasible values. Thus, even in the most optimistic scenario, $g_0^* < \kappa^*, \Omega_M^*$. However, they are of similar magnitude, suggesting that the less-restrictive strong-coupling condition $g = g_0 \cdot n_{\text{cav}}^{1/2} > \kappa, \Gamma_M$, which can enable observation of hybrid optical-mechanical states and optomechanically induced transparency (OMIT) [4, 83], should be possible at relatively low drive powers (*i.e.* small time-averaged number of photons circulating in the cavity mode, n_{cav}). Finally, it is interesting to note that very high single-photon cooperativities, $C_0^* = 4 \cdot (g_0^*)/(\kappa^* \cdot \Gamma_M^*) > 100$

might be feasible for these micro-scale systems. C_0 is an important measure of the strength of interaction between the light and mechanics at the single-photon level [4, 29], which has recently been confirmed as a relevant figure of merit for quantum optomechanics [85].

Chapter 4

Conclusions and Future Work

4.1 Summary

In this work, an integrated membrane-in-the-middle Fabry-Pérot device was presented. The measured mechanical behaviour of the enclosed membrane was well predicted by the established theory, indicating a high degree of control over the fabrication process, and a promising future for these devices in the field of cavity optomechanics. The devices presented had membranes with effective masses ~ 90 pg and fundamental mechanical resonance frequencies of approximately 10 MHz. The mechanical quality factor in all cases was limited by viscous damping due to non-vacuum measurements, and were sensibly well modelled by viscous damping theory. The enclosing optical cavities had a reflection-limited finesse of approximately 10^3 and a linewidth of 0.3 nm. Although this is low relative to many hybrid-assembled cavities in the literature, it should be possible to significantly increase the finesse in future work through the use of improved Bragg mirrors. Moreover, the approach described results in monolithically fabricated, on-chip device arrays with potential densities of over a thousand devices per square centimetre. The work described here has potential to play a role in the future practical implementation of MIM optomechanical systems, where size reduction and system integration might be important factors.

4.2 Future Work

The most immediate improvement that must be made to properly characterize these devices is to repeat the measurements in a vacuum environment, and subsequently also a cryogenic environment. As discussed in Chapter 3, the membrane mechanical quality factor is currently limited by the viscous damping caused by the surrounding air environment. Removing this damping effect and characterizing more intrinsic forms of damping (such as mechanical losses intrinsic to the membrane material and its clamping to the surrounding structure,) will give insight into a more fundamental Q_M that can be compared to other works in the literature. Work is ongoing to design a vacuum chamber that can be suitably integrated into the experimental setup. After these ‘bench top’ measurements have been completed, it is foreseeable that optical measurements could even be moved to a dilution refrigerator allowing for sub-Kelvin cooling. Further experimental work could involve direct measurements of the optomechanical coupling rate g_0 , possibly similar to as is done in Refs. [86, 87] using the technique introduced by Ref. [88].

Even if the Q_M is still comparatively low after vacuum measurements have been conducted, careful engineering of the membrane structure itself can enable both drastic improvements in Q_M through the use of phononic crystals [31], and even room-temperature quantum effects, with a ‘trampoline’ resonator design [30]. Care will have to be taken during patterning to preserve the upper mirror structure, since any access holes designed for patterning of the membrane will have to be similarly etched through the upper mirror.

Fabricating higher quality membranes will likely also require more control over the release process. Ideally, in contrast to the ‘flower’ shaped membranes shown in this thesis, it would be preferable to realize circular membranes. A detailed analysis of membrane stresses could reveal that the contoured edge of the flower membranes introduces non-uniformities in stress, and thus potentially reduces the Q_M . A lateral, in-plane patterned etch stop could possibly be introduced adjacent to the sacrificial layer, to constrain the etching to a pre-defined pattern (*e.g.* a circle,) much as the thin layer of SiO_2 underneath

the sacrificial layer provides a floor for the etch. One implementation may be through the local oxidation of silicon (LOCOS,) which is a process that allows for the masked growth of SiO_2 within a layer of Si [89]. Thus, after sputtering, but before SiN deposition, an additional masking and thermal oxide growth step could be introduced. Modifications to the entire membrane release chemistry may also be worth exploring - a compatible process with a zero etch rate of silicon nitride (unlike XeF_2) would be most desirable.

Finally, once all complications associated with fabrication have been sorted out, these MIM devices may be coupled with accompanying on-chip microwave electrical circuitry, allowing for monolithic, integrated hybrid systems to be made. This would likely involve some sort of patterning of the membrane surface with a metal to allow for capacitive coupling to a microwave resonator, such as in Ref. [13]. Demonstrating bidirectional microwave to optical quantum transduction would make these devices of significant interest for future integrated, quantum technology.

References

- [1] LIGO Scientific Collaboration and Virgo Collaboration, “Observation of gravitational waves from a binary black hole merger,” *Physical Review Letters*, vol. 116, 6 Feb. 2016. DOI: 10.1103/PhysRevLett.116.061102.
- [2] J. Chan, T. P. Mayer Alegre, A. H. Safavi-Naeini, J. T. Hill, A. Krause, S. Groeblacher, M. Aspelmeyer, and O. Painter, “Laser cooling of a nanomechanical oscillator into its quantum ground state,” *Nature*, vol. 478, no. 7367, pp. 89–92, Oct. 2011. DOI: 10.1038/nature10461.
- [3] J. D. Thompson, B. M. Zwickl, A. M. Jayich, F. Marquardt, S. M. Girvin, and J. G. E. Harris, “Strong dispersive coupling of a high-finesse cavity to a micromechanical membrane,” *Nature*, vol. 452, no. 7183, pp. 72–75, Mar. 2008. DOI: 10.1038/nature06715.
- [4] M. Aspelmeyer, T. J. Kippenberg, and F. Marquardt, “Cavity optomechanics,” *Reviews of Modern Physics*, vol. 86, pp. 1391–1452, 4 Dec. 2014. DOI: 10.1103/RevModPhys.86.1391.
- [5] A. G. Krause, M. Winger, T. D. Blasius, Q. Lin, and O. Painter, “A high-resolution microchip optomechanical accelerometer,” *Nature Photonics*, vol. 6, no. 11, pp. 768–772, Nov. 2012. DOI: 10.1038/NPHOTON.2012.245.
- [6] M. Metcalfe, “Applications of cavity optomechanics,” *Applied Physics Reviews*, vol. 1, no. 3, Sep 2014. DOI: 10.1063/1.4896029.
- [7] K. Takeda, K. Nagasaka, A. Noguchi, R. Yamazaki, Y. Nakamura, E. Iwase, J. M. Taylor, and K. Usami, “Electro-mechano-optical detection of nuclear magnetic resonance,” *Optica*, vol. 5, no. 2, pp. 152–158, Feb. 2018. DOI: 10.1364/OPTICA.5.000152.
- [8] M. Karuza, C. Biancofiore, M. Bawaj, C. Molinelli, M. Galassi, R. Natali, P. Tombesi, G. Di Giuseppe, and D. Vitali, “Optomechanically induced transparency in a membrane-in-the-middle setup at room temperature,” *Physical Review A*, vol. 88, no. 1, Jul. 2013. DOI: 10.1103/PhysRevA.88.013804.
- [9] V. Fiore, Y. Yang, M. C. Kuzyk, R. Barbour, L. Tian, and H. Wang, “Storing Optical Information as a Mechanical Excitation in a Silica Optomechanical Resonator,” *Physical Review Letters*, vol. 107, no. 13, Sep. 2011. DOI: {10.1103/PhysRevLett.107.133601}.

- [10] J. T. Hill, A. H. Safavi-Naeini, J. Chan, and O. Painter, “Coherent optical wavelength conversion via cavity optomechanics,” *Nature Communications*, vol. 3, Nov. 2012. DOI: 10.1038/ncomms2201.
- [11] H. Ramp, T. J. Clark, B. D. Hauer, C. Doolin, K. C. Balram, K. Srinivasan, and J. P. Davis, “Wavelength transduction from a 3d microwave cavity to telecom using piezoelectric optomechanical crystals,” *Applied Physics Letters*, vol. 116, no. 17, Apr. 2020. DOI: 10.1063/5.0002160.
- [12] T. Bağcı, A. Simonsen, S. Schmid, L. G. Villanueva, E. Zeuthen, J. Appel, J. M. Taylor, A. Sorensen, K. Usami, A. Schliesser, and E. S. Polzik, “Optical detection of radio waves through a nanomechanical transducer,” *Nature*, vol. 507, no. 7490, pp. 81–85, Mar. 2014. DOI: 10.1038/nature13029.
- [13] R. W. Andrews, R. W. Peterson, T. P. Purdy, K. Cicak, R. W. Simmonds, C. A. Regal, and K. W. Lehnert, “Bidirectional and efficient conversion between microwave and optical light,” *Nature Physics*, vol. 10, no. 4, pp. 321–326, Apr. 2014. DOI: 10.1038/NPHYS2911.
- [14] M. Forsch, R. Stockill, A. Wallucks, I. Marinkovic, C. Gaertner, R. A. Norte, F. van Otten, A. Fiore, K. Srinivasan, and S. Groblacher, “Microwave-to-optics conversion using a mechanical oscillator in its quantum ground state,” *Nature Physics*, vol. 16, no. 1, pp. 69–74, Jan. 2020. DOI: 10.1038/s41567-019-0673-7.
- [15] J. Bochmann, A. Vainsencher, D. D. Awschalom, and A. N. Cleland, “Nanomechanical coupling between microwave and optical photons,” *Nature Physics*, vol. 9, no. 11, pp. 712–716, Nov. 2013. DOI: 10.1038/NPHYS2748.
- [16] J. D. Teufel, T. Donner, D. Li, J. W. Harlow, M. S. Allman, K. Cicak, A. J. Sirois, J. D. Whittaker, K. W. Lehnert, and R. W. Simmonds, “Sideband cooling of micromechanical motion to the quantum ground state,” *Nature*, vol. 475, no. 7356, pp. 359–363, Jul. 2011. DOI: 10.1038/nature10261.
- [17] E. Verhagen, S. Deleglise, S. Weis, A. Schliesser, and T. J. Kippenberg, “Quantum-coherent coupling of a mechanical oscillator to an optical cavity mode,” *Nature*, vol. 482, no. 7383, pp. 63–67, Feb. 2012. DOI: 10.1038/nature10787.
- [18] T. P. Purdy, R. W. Peterson, and C. A. Regal, “Observation of radiation pressure shot noise on a macroscopic object,” *Science*, vol. 339, no. 6121, pp. 801–804, 2013. DOI: 10.1126/science.1231282.
- [19] I. Marinković, A. Wallucks, R. Riedinger, S. Hong, M. Aspelmeyer, and S. Gröblacher, “Optomechanical bell test,” *Physical Review Letters*, vol. 121, p. 220404, 22 Nov. 2018. DOI: 10.1103/PhysRevLett.121.220404.

- [20] K. Jaehne, C. Genes, K. Hammerer, M. Wallquist, E. S. Polzik, and P. Zoller, “Cavity-assisted squeezing of a mechanical oscillator,” *Physical Review A*, vol. 79, no. 6, Jun. 2009. DOI: 10.1103/PhysRevA.79.063819.
- [21] T. P. Purdy, P.-L. Yu, R. W. Peterson, N. S. Kampel, and C. A. Regal, “Strong optomechanical squeezing of light,” *Physical Review X*, vol. 3, p. 031012, 3 Sep. 2013. DOI: 10.1103/PhysRevX.3.031012.
- [22] A. Heidmann, Y. Hadjar, and M. Pinard, “Quantum nondemolition measurement by optomechanical coupling,” *Applied Physics B - Lasers and Optics*, vol. 64, no. 2, pp. 173–180, Feb. 1997. DOI: 10.1007/s003400050162.
- [23] J. C. Sankey, C. Yang, B. M. Zwickl, M. Jayich, and J. G. E. Harris, “Strong and tunable nonlinear optomechanical coupling in a low-loss system,” *Nature Physics*, vol. 6, no. 9, pp. 707–712, Sep. 2010. DOI: 10.1038/NPHYS1707.
- [24] A. Nunnenkamp, K. Borkje, J. G. E. Harris, and S. M. Girvin, “Cooling and squeezing via quadratic optomechanical coupling,” *Physical Review A*, vol. 82, no. 2, Aug. 2010. DOI: 10.1103/PhysRevA.82.021806.
- [25] V. Dumont, S. Bernard, C. Reinhardt, A. Kato, M. Ruf, and J. C. Sankey, “Flexure-tuned membrane-at-the-edge optomechanical system,” *Optics Express*, vol. 27, no. 18, pp. 25731–25748, Sep. 2019. DOI: 10.1364/OE.27.025731.
- [26] N. E. Flowers-Jacobs, S. W. Hoch, J. C. Sankey, A. Kashkanova, A. M. Jayich, C. Deutsch, J. Reichel, and J. G. E. Harris, “Fiber-cavity-based optomechanical device,” *Applied Physics Letters*, vol. 101, no. 22, Nov. 2012. DOI: 10.1063/1.4768779.
- [27] T. P. Purdy, R. W. Peterson, P.-L. Yu, and C. A. Regal, “Cavity optomechanics with Si_3N_4 membranes at cryogenic temperatures,” *New Journal of Physics*, vol. 14, Nov. 2012. DOI: 10.1088/1367-2630/14/11/115021.
- [28] F. Fogliano, B. Besga, A. Reigue, P. Heringlake, L. M. de Lépinay, C. Vaneph, J. Reichel, B. Pigeau, and O. Arcizet, *Cavity nano-optomechanics in the ultrastrong coupling regime with ultrasensitive force sensors*, 2019. eprint: 1904.01140 (quant-ph).
- [29] C. Reinhardt, T. Mueller, A. Bourassa, and J. C. Sankey, “Ultralow-noise sin trampoline resonators for sensing and optomechanics,” *Physical Review X*, vol. 6, no. 2, Apr. 2016. DOI: 10.1103/PhysRevX.6.021001.
- [30] R. A. Norte, J. P. Moura, and S. Groblacher, “Mechanical resonators for quantum optomechanics experiments at room temperature,” *Physical Review Letters*, vol. 116, p. 147202, 14 Apr. 2016. DOI: 10.1103/PhysRevLett.116.147202.

- [31] C. Reetz, R. Fischer, G. G. T. Assumpcao, D. P. McNally, P. S. Burns, J. C. Sankey, and C. A. Regal, “Analysis of membrane phononic crystals with wide band gaps and low-mass defects,” *Physical Review Applied*, vol. 12, no. 4, Oct. 2019. DOI: 10.1103/PhysRevApplied.12.044027.
- [32] V. P. Adiga, B. Ilic, R. A. Barton, I. Wilson-Rae, H. G. Craighead, and J. M. Parpia, “Approaching intrinsic performance in ultra-thin silicon nitride drum resonators,” *Journal of Applied Physics*, vol. 112, no. 6, Sep. 2012. DOI: 10.1063/1.4754576.
- [33] T. W. Allen, J. Silverstone, N. Ponnampalam, T. Olsen, A. Meldrum, and R. G. DeCorby, “High-finesse cavities fabricated by buckling self-assembly of a-si/sio2 multilayers,” *Optics Express*, vol. 19, no. 20, pp. 18 903–18 909, Sep. 2011. DOI: 10.1364/OE.19.018903.
- [34] C. A. Potts, A. Melnyk, H. Ramp, M. H. Bitarafan, D. Vick, L. J. LeBlanc, J. P. Davis, and R. G. DeCorby, “Tunable open-access microcavities for on-chip cavity quantum electrodynamics,” *Applied Physics Letters*, vol. 108, no. 4, p. 041 103, 2016. DOI: 10.1063/1.4940715.
- [35] G. J. Hornig, T. R. Harrison, L. Bu, S. Azmayesh-Fard, and R. G. DeCorby, “Wavemeter based on dispersion and speckle in a tapered hollow waveguide,” *OSA Continuum*, vol. 2, no. 2, pp. 495–506, Feb. 2019. DOI: 10.1364/OSAC.2.000495.
- [36] M. H. Bitarafan and R. G. DeCorby, “Small-mode-volume, channel-connected fabry-perot microcavities on a chip,” *Applied Optics*, vol. 56, no. 36, pp. 9992–9997, Dec. 2017. DOI: 10.1364/AO.56.009992.
- [37] —, “On-Chip High-Finesse Fabry-Perot Microcavities for Optical Sensing and Quantum Information,” *Sensors*, vol. 17, no. 8, AUG 2017. DOI: 10.3390/s17081748.
- [38] S. Al-Sumaidae, M. H. Bitarafan, C. A. Potts, J. P. Davis, and R. G. DeCorby, “Cooperativity enhancement in buckled-dome microcavities with omnidirectional claddings,” *Optics Express*, vol. 26, no. 9, pp. 11 201–11 212, Apr. 2018. DOI: 10.1364/OE.26.011201.
- [39] J. Maldaner, S. Al-Sumaidae, G. J. Hornig, L. J. LeBlanc, and R. G. DeCorby, “Liquid infiltration of monolithic open-access fabry-perot microcavities,” *Applied Optics*, vol. 59, no. 23, pp. 7125–7130, Aug. 2020. DOI: 10.1364/AO.399601.
- [40] H. A. Macleod, *Thin-film optical filters*. CRC Press, 2018.
- [41] C. Hood, H. Kimble, and J. Ye, “Characterization of high-finesse mirrors: Loss, phase shifts, and mode structure in an optical cavity,” *Physical Review A*, vol. 64, no. 3, Sep. 2001. DOI: 10.1103/PhysRevA.64.033804.

- [42] L. Brovelli and U. Keller, “Simple analytical expressions for the reflectivity and the penetration depth of a bragg mirror between arbitrary media,” *Optics Communications*, vol. 116, no. 4, pp. 343–350, 1995. DOI: [https://doi.org/10.1016/0030-4018\(95\)00084-L](https://doi.org/10.1016/0030-4018(95)00084-L).
- [43] E. Hecht, *Optics*. Addison-Wesley, 1998.
- [44] W. Shen, X. Liu, B. Huang, Y. Zhu, and P. Gu, “The effects of reflection phase shift on the optical properties of a micro-opto-electro-mechanical system fabry-perot tunable filter,” *Journal of Optics A: Pure and Applied Optics*, vol. 6, no. 9, pp. 853–858, Aug. 2004. DOI: 10.1088/1464-4258/6/9/006.
- [45] D. J. Wilson, C. A. Regal, S. B. Papp, and H. J. Kimble, “Cavity optomechanics with stoichiometric sin films,” *Physical Review Letters*, vol. 103, no. 20, Nov. 2009. DOI: 10.1103/PhysRevLett.103.207204.
- [46] A. Yariv and P. Yeh, *Photonics*. Oxford University Press, 2007.
- [47] J. Hutchinson, M. Thouless, and E. Liniger, “Growth and configurational stability of circular, buckling-driven film delaminations,” *Acta Metallurgica et Materialia*, vol. 40, no. 2, pp. 295–308, 1992. DOI: [https://doi.org/10.1016/0956-7151\(92\)90304-W](https://doi.org/10.1016/0956-7151(92)90304-W).
- [48] M. H. Bitarafan, H. Ramp, T. W. Allen, C. Potts, X. Rojas, A. J. R. MacDonald, J. P. Davis, and R. G. DeCorby, “Thermomechanical characterization of on-chip buckled dome fabry-perot microcavities,” *Journal of the Optical Society of America B - Optical Physics*, vol. 32, no. 6, pp. 1214–1220, Jun. 2015. DOI: 10.1364/JOSAB.32.001214.
- [49] L. B. Freund and S. Suresh, *Thin Film Materials : Stress, Defect Formation and Surface Evolution*. Cambridge University Press, 2003, ISBN: 9780521822817.
- [50] J. J. Talghader, “Thermal and mechanical phenomena in micromechanical optics,” *Journal of Physics D: Applied Physics*, vol. 37, no. 10, pp. 109–122, Apr. 2004. DOI: 10.1088/0022-3727/37/10/r01.
- [51] D. R. M. Crooks, G. Cagnoli, M. M. Fejer, G. Harry, J. Hough, B. T. Khuri-Yakub, S. Penn, R. Route, S. Rowan, P. H. Sneddon, I. O. Wygant, and G. G. Yaralioglu, “Experimental measurements of mechanical dissipation associated with dielectric coatings formed using sio₂, ta₂o₅ and al₂o₃,” *Classical and Quantum Gravity*, vol. 23, no. 15, pp. 4953–4965, Aug. 2006. DOI: 10.1088/0264-9381/23/15/014.
- [52] C. Y. Wang and C. Wang, *Structural vibration: exact solutions for strings, membranes, beams, and plates*. CRC Press, 2013.
- [53] B. Hauer, C. Doolin, K. Beach, and J. Davis, “A general procedure for thermomechanical calibration of nano/micro-mechanical resonators,” *Annals of Physics*, vol. 339, pp. 181–207, 2013. DOI: <https://doi.org/10.1016/j.aop.2013.08.003>.

- [54] D. R. Southworth, H. G. Craighead, and J. M. Parpia, “Pressure dependent resonant frequency of micromechanical drumhead resonators,” *Applied Physics Letters*, vol. 94, no. 21, May 2009. DOI: 10.1063/1.3141731.
- [55] A. H. Ghadimi, D. J. Wilson, and T. J. Kippenberet, “Radiation and internal loss engineering of high-stress silicon nitride nanobeams,” *Nano Letters*, vol. 17, no. 6, pp. 3501–3505, Jun. 2017. DOI: 10.1021/acs.nanolett.7b00573.
- [56] L. G. Villanueva and S. Schmid, “Evidence of surface loss as ubiquitous limiting damping mechanism in sin micro- and nanomechanical resonators,” *Physical Review Letters*, vol. 113, pp. 227 201–227 207, 22 Nov. 2014. DOI: 10.1103/PhysRevLett.113.227201.
- [57] A. M. Jayich, J. C. Sankey, B. M. Zwickl, C. Yang, J. D. Thompson, S. M. Girvin, A. A. Clerk, F. Marquardt, and J. G. E. Harris, “Dispersive optomechanics: A membrane inside a cavity,” *New Journal of Physics*, vol. 10, Sep. 2008. DOI: 10.1088/1367-2630/10/9/095008.
- [58] G. J. Hornig, S. Al-Sumaidae, J. Maldaner, L. Bu, and R. G. DeCorby, “Monolithically integrated membrane-in-the-middle cavity optomechanical systems,” *Optics Express*, 2020.
- [59] S. A. Campbell, *The science and engineering of microelectronic fabrication*, 2nd ed. Oxford University Press, 2001.
- [60] A. Stoffel, A. Kovacs, W. Kronast, and B. Muller, “Lpcvd against pecvd for micromechanical applications,” *Journal of Micromechanics and Microengineering*, vol. 6, no. 1, pp. 1–13, Mar. 1996, 6th European Workshop on Micromechanics (MME 95), Copenhagen, Denmark. DOI: 10.1088/0960-1317/6/1/001.
- [61] A. Bagolini, A. Picciotto, M. Crivellari, P. Conci, and P. Bellutti, “Pecvd silicon-rich nitride and low stress nitride films mechanical characterization using membrane point load deflection,” *Journal of Micromechanics and Microengineering*, vol. 26, no. 2, Feb. 2016. DOI: 10.1088/0960-1317/26/2/025004.
- [62] M. Ghaderi and R. F. Wolffenbutte, “Design and fabrication of ultrathin silicon-nitride membranes for use in uv-visible airgap-based mems optical filters,” in *27th Micromechanics and Microsystems Europe Workshop (MME 2016)*, ser. Journal of Physics Conference Series, vol. 757, IOP Publishing Ltd., 2016. DOI: 10.1088/1742-6596/757/1/012032.
- [63] P. Froeter, X. Yu, W. Huang, F. Du, M. Li, I. Chun, S. H. Kim, K. J. Hsia, J. A. Rogers, and X. Li, “3d hierarchical architectures based on self-rolled-up silicon nitride membranes,” *Nanotechnology*, vol. 24, no. 47, Nov. 2013. DOI: 10.1088/0957-4484/24/47/475301.

- [64] J. Maldaner, S. Al-Sumaidae, and R. G. DeCorby, “Theoretical study of silicon-based bragg mirrors for cavity qed applications,” *Journal of the Optical Society of America B*, 2020.
- [65] G. Caliano, F. Galanello, A. Caronti, R. Carotenuto, M. Pappalardo, V. Foglietti, and N. Lamberti, “Micromachined ultrasonic transducers using silicon nitride membrane fabricated in pecvd technology,” in *2000 IEEE Ultrasonics Symposium. Proceedings. An International Symposium (Cat. No.00CH37121)*, vol. 1, 2000, pp. 963–968.
- [66] J. Maldaner, *Towards a Monolithic Process for Open-Access Fabry-Perot Etalon Cavities*. University of Alberta, Aug. 2020, MSc. Thesis.
- [67] K. Williams, K. Gupta, and M. Wasilik, “Etch rates for micromachining processing - part ii,” *Journal of Microelectromechanical Systems*, vol. 12, no. 6, pp. 761–778, Dec. 2003. DOI: 10.1109/JMEMS.2003.820936.
- [68] K. Williams and R. Muller, “Etch rates for micromachining processing,” *Journal of Microelectromechanical Systems*, vol. 5, no. 4, pp. 256–269, Dec. 1996. DOI: 10.1109/84.546406.
- [69] O. Paul, D. Westberg, M. Hornung, V. Ziebart, and H. Baltes, “Sacrificial aluminum etching for cmos microstructures,” in *Proceedings IEEE The Tenth Annual International Workshop on Micro Electro Mechanical Systems. An Investigation of Micro Structures, Sensors, Actuators, Machines and Robots*, 1997, pp. 523–528.
- [70] S. Kuiper, M. de Boer, C. van Rijn, W. Nijdam, G. Krijnen, and M. Elwenspoek, “Wet and dry etching techniques for the release of sub-micrometre perforated membranes,” *Journal of Micromechanics and Microengineering*, vol. 10, no. 2, pp. 171–174, Jun. 2000. DOI: 10.1088/0960-1317/10/2/312.
- [71] S. Soulimane, F. Casset, P.-L. Charvet, C. Maeder, and M. Aid, “Planarization of photoresist sacrificial layer for mems fabrication,” *Microelectronic Engineering*, vol. 84, no. 5, pp. 1398–1400, 2007, Proceedings of the 32nd International Conference on Micro- and Nano-Engineering. DOI: <https://doi.org/10.1016/j.mee.2007.01.220>.
- [72] N. Ferrell, J. Woodard, and D. Hansford, “Fabrication of polymer microstructures for mems: Sacrificial layer micromolding and patterned substrate micromolding,” *Biomedical Microdevices*, vol. 9, no. 6, pp. 815–821, 2007.
- [73] C. Wu, L. Erdmann, K. J. Gabriel, and S. C. Yao, “Fabrication of silicon sidewall profiles for fluidic applications using modified advanced silicon etching,” in *MEMS Design, Fabrication, Characterization, and Packaging*, vol. 4407, SPIE, 2001, pp. 100–108. DOI: 10.1117/12.425289.

- [74] Y. Xie, N. Banerjee, and C. H. Mastrangelo, “Microfabricated spherical pressure sensing particles for pressure and flow mapping,” in *2013 Transducers Eurosensors XXVII: The 17th International Conference on Solid-State Sensors, Actuators and Microsystems (TRANSDUCERS EUROSENSORS XXVII)*, 2013, pp. 1771–1774.
- [75] J. Soon, P. Neuzil, C. Fang, J. Reboud, C. Wong, and L. Kao, “High throughput melting curve analysis in monolithic silicon-based microfluidic device,” *microTAS*, 2010.
- [76] J. D. Zahn, K. J. Gabriel, and G. K. Fedder, “A direct plasma etch approach to high aspect ratio polymer micromachining with applications in biomems and cmos-mems,” in *Technical Digest. MEMS 2002 IEEE International Conference. Fifteenth IEEE International Conference on Micro Electro Mechanical Systems (Cat. No.02CH37266)*, 2002, pp. 137–140.
- [77] P. B. Chu, J. T. Chen, R. Yeh, G. Lin, J. C. P. Huang, B. A. Warneke, and S. J. Pister, “Controlled pulse-etching with xenon difluoride,” in *Proceedings of International Solid State Sensors and Actuators Conference (Transducers '97)*, vol. 1, 1997, 665–668 vol.1.
- [78] T. Capelle, Y. Tsaturyan, A. Barg, and A. Schliesser, “Polarimetric analysis of stress anisotropy in nanomechanical silicon nitride resonators,” *Applied Physics Letters*, vol. 110, no. 18, May 2017. DOI: 10.1063/1.4982876.
- [79] C. Fairclough, “Analyzing the viscous and thermal damping of a mems micromirrors,” *COMSOL Blog*, 2018. [Online]. Available: <https://www.comsol.com/blogs/analyzing-the-viscous-and-thermal-damping-of-a-mems-micromirror/> (visited on 08/11/2020).
- [80] M. A. J. Suijlen, J. J. Koning, M. A. J. van Gils, and H. C. W. Beijerinck, “Squeeze film damping in the free molecular flow regime with full thermal accommodation,” *Sensors and Actuators A - Physical*, vol. 156, no. 1, pp. 171–179, Nov. 2009. DOI: 10.1016/j.sna.2009.03.025.
- [81] M. Poggio, C. L. Degen, H. J. Mamin, and D. Rugar, “Feedback cooling of a cantilever’s fundamental mode below 5 mk,” *Physical Review Letters*, vol. 99, p. 017201, 1 Jul. 2007. DOI: 10.1103/PhysRevLett.99.017201.
- [82] F. Elste, S. M. Girvin, and A. A. Clerk, “Quantum noise interference and backaction cooling in cavity nanomechanics,” *Physical Review Letters*, vol. 102, pp. 207–209, 20 May 2009. DOI: 10.1103/PhysRevLett.102.207209.
- [83] T. Ojanen and K. Borkje, “Ground-state cooling of mechanical motion in the unresolved sideband regime by use of optomechanically induced transparency,” *Physical Review A*, vol. 90, no. 1, Jul. 2014. DOI: 10.1103/PhysRevA.90.013824.

- [84] M. Rossi, D. Mason, J. Chen, Y. Tsaturyan, and A. Schliesser, “Measurement based quantum control of mechanical motion,” *Nature*, vol. 563, no. 7729, pp. 53+, Nov. 2018. DOI: 10.1038/s41586-018-0643-8.
- [85] K. Børkje, “Critical quantum fluctuations and photon antibunching in optomechanical systems with large single-photon cooperativity,” *Physical Review A*, vol. 101, p. 053 833, 5 May 2020. DOI: 10.1103/PhysRevA.101.053833.
- [86] A. J. R. MacDonald, B. D. Hauer, X. Rojas, P. H. Kim, G. G. Popowich, and J. P. Davis, “Optomechanics and thermometry of cryogenic silica microresonators,” *Physical Review A*, vol. 93, p. 013 836, 1 Jan. 2016. DOI: 10.1103/PhysRevA.93.013836.
- [87] H. Ramp, B. D. Hauer, K. C. Balram, T. J. Clark, K. Srinivasan, and J. P. Davis, “Elimination of thermomechanical noise in piezoelectric optomechanical crystals,” *Physical Review Letters*, vol. 123, p. 093 603, 9 Aug. 2019. DOI: 10.1103/PhysRevLett.123.093603.
- [88] M. L. Gorodetsky, A. Schliesser, G. Anetsberger, S. Deleglise, and T. J. Kippenberg, “Determination of the vacuum optomechanical coupling rate using frequency noise calibration,” *Optics Express*, vol. 18, no. 22, pp. 23 236–23 246, Oct. 2010. DOI: 10.1364/OE.18.023236.
- [89] A. Naiman, B. Desiatov, L. Stern, N. Mazurski, J. Shappir, and U. Levy, “Ultrahigh-Q silicon resonators in a planarized local oxidation of silicon platform,” *Optics Letters*, vol. 40, no. 9, 1892–1895, May 2015. DOI: 10.1364/OL.40.001892.

Appendix A

Wafer Survey

While a reasonable yield of good devices was obtained for the process, many failed for a variety of reasons. Microscope images of these are given in Fig. A.1 for future reference and process development.

With reference to Figure A.1:

- a) Dome popped off prior to etching.
- b) A popped off dome.
- c) An intact dome, but with multilayer etching visible in the rings around the holes.
- d) An extremely short, non symmetrically buckled dome.
- e) A dome half broken after etching, showing the holes in the membrane surface below
- f) An unbuckled dome.
- g) Dome with flakes visible on the top surface.
- h) 3 micron sized etch holes, visible but not going completely through the top mirror.
- i) Unbuckled dome, but with internal sacrificial layer removal. This was unusual, most unbuckled domes showed no activity during XeF_2 etching.
- j) A short dome with no etch holes.
- k) A dome with only some of it's holes etched, showing the sacrificial layer only being removed from one side.
- l) A typical otherwise good device with a huge overbuckling. Original teflon location is visible in the lower right. side.
- m) An island of sacrificial Si in the centre of the etch holes that was note removed.

- n) A larger island, similar to m.
- o) Multilayer etching
- p) Multilayer etching with an island.
- q) A popped off dome with 3 micron holes, showing the etching hasn't touched the membrane surface inside.
- r) Etched holes that were patterned too close to one another, causing them to bridge and the center of the dome to collapse.




Figure A.1: An overview of various different types of devices visible on the wafer. See text for descriptions.

Appendix B

Extended Theory

B.1 Optical Loss Expressions

For a Bragg mirror of large N such that $R \rightarrow 1$, the absorption can be approximated with the following equations:

In the case that the first layer is a high index layer:

$$A \approx 2\pi n_{\text{in}} \frac{\kappa_{\text{H}} + \kappa_{\text{L}}}{n_{\text{H}}^2 - n_{\text{L}}^2}. \quad (\text{B.1})$$

Or in the case that the first layer is a low index layer

$$A \approx \frac{2\pi}{n_{\text{in}}} \frac{n_{\text{L}}^2 \kappa_{\text{H}} + n_{\text{H}}^2 \kappa_{\text{L}}}{n_{\text{H}}^2 - n_{\text{L}}^2}. \quad (\text{B.2})$$

The scattering due to a rough interface can be approximated from the following expression

$$S \approx 8\pi^2 \frac{n_{\text{in}}}{n_{\text{L}}} \frac{4n_{\text{H}}n_{\text{L}}}{(n_{\text{H}} + n_{\text{L}})^2} (n_{\text{H}}^2 - n_{\text{L}}^2) \left(\frac{\sigma}{\lambda_{\text{Bragg}}} \right)^2, \quad (\text{B.3})$$

where σ is the RMS roughness of the planar interfaces within the thin-film stack.

B.2 Optical Transfer Matrices

Simulating the optical characteristics of complex thin film structures is extremely useful in their design and characterization. Most commonly used in this work is a one-dimensional transfer matrix model, that allows for for one to




Figure B.1: A diagram of the electric fields for a general transfer matrix problem, corresponding to Equation B.6.

simulate the transmission, reflection, and absorption of an arbitrary stack of thin-films at any wavelength. The internal electric field may also be mapped inside of the structures, which is particularly useful for studying cavities on resonance.

Consider a monochromatic plane wave incident on a thin film. The electric field takes the form:

$$E = E(z) \exp [i(\omega t - \beta z)], \quad (\text{B.4})$$

where $E(z)$ represents the transverse electric field profile, $\omega = 2\pi c/\lambda$ is the angular frequency of the electric field, and β is the spatial frequency in the direction of z such that:

$$\beta = n_q \left(\frac{\omega}{c} \right) \sin \theta_q, \quad (\text{B.5})$$

where q denotes a particular layer, n_q is the complex refractive index, and θ_q is the complex angle of propagation in that layer according to Snells' law. For any multilayer structure as described above, the complex electric field amplitudes of the incident (from both sides), reflected, and transmitted can be related through a simple transfer matrix [46]:

$$\begin{bmatrix} A_1 \\ B_1 \end{bmatrix} = \begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix} \begin{bmatrix} A_2 \\ B_2 \end{bmatrix}, \quad (\text{B.6})$$

where A and B represent the electric field amplitudes in directions as shown in Figure B.1 on either side of the thin film. This transfer matrix M can be used to describe any multilayer structure. The structure itself can also be represented as a series of these matrices, and they can be multiplied together to obtain the overall matrix. If a multilayer structure is broken up into N layers, the overall matrix is represented by

$$M = M_N M_{N-1} \dots M_2 M_1 \quad (\text{B.7})$$

Two types of matrices are required to define a structure: one concerning the interface between two different materials, and one that deals with the propagation through a homogeneous material. As the electric field propagates through a length of a homogeneous material a phase is incurred according to its spatial frequency, in addition to an amplitude loss from any absorption and/or scattering within the material. Both of these effects can be described by the materials complex index. The propagation matrix M_{prop} for the layer is given by [46]:

$$M_{\text{prop}} = \begin{bmatrix} e^{ik_z d} & 0 \\ 0 & e^{-ik_z d} \end{bmatrix}, \quad (\text{B.8})$$

where k_z is the component of the complex wavevector in the material projected along the optical axis, and d is the distance of propagation within the layer. Upon an interface with a different material, the behaviour of light is polarization dependent and dictated by Fresnel's equations. By considering Fresnel's equations for both polarizations, and the movement of light incident from both the left and right, an interface transfer matrix M_{int} can be created [46]:

$$M_{\text{int}} = \frac{1}{t_{12}} \begin{bmatrix} 1 & r_{12} \\ r_{12} & 1 \end{bmatrix} = \begin{cases} \begin{bmatrix} \frac{1}{2}(1 + \frac{k_{2z}}{k_{1z}}) & \frac{1}{2}(1 - \frac{k_{2z}}{k_{1z}}) \\ \frac{1}{2}(1 - \frac{k_{2z}}{k_{1z}}) & \frac{1}{2}(1 + \frac{k_{2z}}{k_{1z}}) \end{bmatrix}, & \text{TE} \\ \begin{bmatrix} \frac{1}{2}(1 + \frac{n_2^2 k_{1z}}{n_1^2 k_{2z}}) & \frac{1}{2}(1 - \frac{n_2^2 k_{1z}}{n_1^2 k_{2z}}) \\ \frac{1}{2}(1 - \frac{n_2^2 k_{1z}}{n_1^2 k_{2z}}) & \frac{1}{2}(1 + \frac{n_2^2 k_{1z}}{n_1^2 k_{2z}}) \end{bmatrix}, & \text{TM} \end{cases} \quad (\text{B.9})$$

$$k_{qz} = \frac{\omega}{c} n_q \cos \theta_q, \quad (\text{B.10})$$

where k_{qz} is the complex wavevector along the optical axis for material q (either 1 or 2).

These two types of matrices may be multiplied in succession to form a complete description of the multilayer system. This is done from right to left, with the last layer being the leftmost matrix. The total matrix for a system of N layers would be [46]:

$$M_{\text{sys}} = M_{\text{prop},N} M_{\text{int},N,N-1} M_{\text{prop},N-1} \dots M_{\text{prop},1} = \begin{bmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{bmatrix}. \quad (\text{B.11})$$

If the input medium is assumed to be lossless, the reflectance can be found through [46]:

$$R = |r|^2 = \left| \frac{m_{21}}{m_{11}} \right|^2. \quad (\text{B.12})$$

If the exit medium is also lossless, and the propagation angle in every layer is real (*i.e.* the wave actually propagates in each layer and does not form an evanescent wave) the transmittance can be found through [46]:

$$T = |t|^2 = \frac{n_{\text{out}} \cos \theta_{\text{out}}}{n_{\text{in}} \cos \theta_{\text{in}}} \left| \frac{1}{m_{11}} \right|^2. \quad (\text{B.13})$$

Hecht Form

An alternative transfer matrix formulation has also been presented in Ref. [43] by Hecht which tracks the tangential electric field (E) and the tangential magnetic field (H) field at each interface. Instead of having separate propagation and interface matrices as in the previously described formulation from Ref.

[46], it combines these into one for each layer allowing each layer to be represented succinctly by a single matrix. For a layer q the Hecht transfer matrix $M_{H,q}$ takes the form [43]:

$$M_{H,q} = \begin{bmatrix} \cos k_0 h & i \sin k_0 h / X_q \\ i X_q \sin k_0 h & \cos k_0 h \end{bmatrix}, \quad (\text{B.14})$$

where

$$X_q = \begin{cases} Y_q = \sqrt{\frac{\epsilon_0}{\mu_0}} n_q \cos \theta_q, & \text{TE} \\ Z_q = \sqrt{\frac{\mu_0}{\epsilon_0}} n_q \cos \theta_q, & \text{TM} \end{cases}. \quad (\text{B.15})$$

In this description, Y_q is the tilted optical admittance for the layer q , and Z_q is the tilted optical impedance for the layer q . $k_0 = 2\pi/\lambda$ is the spatial frequency in a vacuum, and $h = n_q d_q \cos \theta_q$ is the optical path length along the axis of propagation in layer q .

The overall propagation matrix is formed by multiplying these matrices left to right, in similar fashion to Equation B.11. For a system of N layers the overall transfer matrix M_H is [43]:

$$M_H = M_{H,1} M_{H,2} \dots M_{H,N} = \begin{bmatrix} m_{H11} & m_{H12} \\ m_{H21} & m_{H22} \end{bmatrix}. \quad (\text{B.16})$$

The reflection and transmission coefficients can be found then with the following formula [43]:

$$r_{TE} = \frac{Y_{\text{in}} m_{H11} + Y_{\text{in}} Y_{\text{out}} m_{H12} - m_{H21} - Y_{\text{out}} m_{H22}}{Y_{\text{in}} m_{H11} + Y_{\text{in}} Y_{\text{out}} m_{H12} + m_{H21} + Y_{\text{out}} m_{H22}}, \quad (\text{B.17})$$

$$t_{TE} = \frac{2Y_{\text{in}}}{Y_{\text{in}} m_{H11} + Y_{\text{in}} Y_{\text{out}} m_{H12} + m_{H21} + Y_{\text{out}} m_{H22}}, \quad (\text{B.18})$$

$$r_{TM} = \frac{Z_{\text{in}} m_{H11} + Z_{\text{in}} Z_{\text{out}} m_{H12} - m_{H21} - Z_{\text{out}} m_{H22}}{Z_{\text{in}} m_{H11} + Z_{\text{in}} Z_{\text{out}} m_{H12} + m_{H21} + Z_{\text{out}} m_{H22}}, \quad (\text{B.19})$$

$$t_{TM} = \frac{n_{\text{in}}}{n_{\text{out}}} \cdot \frac{2Z_{\text{in}}}{Z_{\text{in}} m_{H11} + Z_{\text{in}} Z_{\text{out}} m_{H12} + m_{H21} + Z_{\text{out}} m_{H22}}, \quad (\text{B.20})$$

where Y and Z still represent the tilted optical admittances and impedances, but for the input and output media. Note the symmetry between the TE and

TM expressions except for the additional fraction in front of the t_{TM} expression. The transmittance and reflectance may then be found through the magnitude squared of the respective expressions.

Appendix C

Norcada Inc. Membrane Development

Prior to the fabrication approach described in the main text, initial work was done in collaboration with local Edmonton company Norcada. Norcada is well respected in the optics community for their production of high quality micro membranes available for commercial purchase. Their devices are featured in a number of high profile publications. For this reason, as well as our groups initial unfamiliarity with mechanical membranes, the project began as an effort to take their high quality mechanical systems and attempt to directly integrate them with our groups optical systems. Norcada has developed a process for producing individual membranes as standalone devices mounted on frames. LPCVD is used to place a stoichiometric high tensile stressed silicon nitride layer on the top face of an Si wafer, and then a masked KOH etch is applied on the wafer back side. KOH characteristically etches the silicon crystalline $\langle 100 \rangle$ plane is etched at a 54.7° angle relative to the $\langle 111 \rangle$ plane, which when etched from the back creates sloped inward side walls. The etch reaches the topside membrane layer, and the membranes become released and free. Once the process is completed, the wafers is diced smaller into square frames.

C.1 Process Design

The initial approach to integrating these membranes within our cavities was through a ‘sandwich’ process, outlined in Figure C.1. The ‘off-the-shelf’ pre-




Figure C.1: The proposed process for integrating our mirror dome cavity technology with Norcada’s high quality mechanical membranes. (a) A fully fabricated membrane chip is used as the starting substrate, already released from the silicon below. (b) The PTFE is patterned, top mirror deposited, and buckled as per the standard process. (c) The substrate is flipped to the other side, and a sacrificial layer is filled to the level of the substrate on the other side. The bottom mirror is deposited on top of this, and then access holes are etched. (d) A final etching step removes the internal sacrificial layer and the device is finished.

fabricated Norcada membrane would be used as a starting substrate for our process. We would start by patterning the Teflon low adhesion layer on top, depositing a mirror, and buckling the devices as is usually done for the top mirror on a wafer. The substrate would then be flipped, and a bottom mirror would be placed on the backside, offset from the membrane back surface by some sort of sacrificial layer. The sacrificial layer would then be removed through some form of access holes, presumably straight through the back mirror as to leave the membrane untouched. The process could be in principle scaled from an individual frame to a full production wafer.

Norcada provided a set of individually packaged membranes on frames for process development. The geometries of the various devices are given in Table C.1. Initial success was achieved, albeit with great difficulty, producing devices that had frames with intact membranes and fully buckled domes on

	Set 1	Set 2	Set 3	Set 4
Membrane Width (μm)	10	20	40	100
Membrane Thickness (nm)	150	150	150	11
Frame Width (mm)	5	5	5	7.5
Amount	5	5	5	10

Table C.1: The membranes provided by Norcada. All membranes and frames are square. The 100 μm wide membranes are actually nine per frame, configured as a closely spaced 3x3 grid.

top. However, practical handling of the individual membranes made processing extremely difficult. Due to the small size of the frames, many devices were either damaged or destroyed. The processing of substrates smaller than a standard wafer comes with it many complications in the use of equipment, since most tools in the nanoFab are equipped to process standard 4 inch or 6 inch diameter substrates. In most cases smaller pieces such as the membrane frames have to be mounted to a 4 inch carrier wafer during all processing.¹ We chose to bond the membrane frames to a 4 inch wafer through the use of double sided, high temperature 170 °C heat release tape (Semiconductor Equipment Corp., REVALPHA Heat Release No.3195V). This pink coloured tape is notably distinct from the blue coloured double sided tape used in the nanoFab, since the blue tape has a lower release temperature and thus would potentially fail during the 150 °C sputtering.

A major issue during processing was that due to the sensitive nature of these membranes, sonication was not possible during lift-off of the PTFE layer. Simply soaking the devices in acetone did not produce a high quality lift-off result, and so many different methods were tested in an attempt to produce a better quality lift-off. This included agitation with a mixing hotplate, a different resist chemistry (LOR, lift-off-resist), and a heated lift-off (with PG remover in the place of acetone.) The poor lift-off from these samples mo-

¹There was a lot of work done here in parallel with another project relating to the placement of domes on top of fibre ferrules. Ferrules small size and unusual format, and being made of zirconia (ZrO_2) which is a notoriously difficult surface to adhere to, presented similar challenges.




Figure C.2: Examples of typical poor liftoff of PTFE found during processing. (a) A folded over flake (b) A Hanging Chad (c) Rough Edges (on a different mask used for testing lift-off).

tivated the group to pursue an effort in lift-off optimization. Poor lift-off is characterized by ‘hanging chads’ (Fig. C.2.a), folded over flakes (Fig. C.2.b), and rough edges (Fig. C.2.c). Initially, removal of these flakes and chads on the dome patterns was attempted by scraping with both the Alphastep profilometer needle tip and a micrometer controlled four point probe station needle, but this only resulted in scratching the PTFE completely off of the substrate.

LOR bilayer, or lift-off resist, was a different photoresist approach that was experimented with extensively. A commonly documented problem in literature with lift-off is the collapsing of side walls onto the feature. When the material is deposited on top of the mask isotropically as the Teflon is, it coats both the bottom surface as well as the sides of the mask. A bilayer LOR resist works to fix this problem by using instead two different photoresists, each with a different solvent. LOR is placed as the first resist on the bottom, and then a standard resist such as HPR504 or AZ1512 is placed on top. During lift-off, the device is placed in a chemical (MF-319) which selectively dissolves the bottom LOR resist. This causes the top layer of resist to physically fall off of the wafer in one piece, and this breaking action should in principle prevent the side walls from falling back onto the device. Further, the different developers

used by each resist allowed for a selective over development of the lower resist, which allowed for a sloped sidewall to be created, which are known to further mitigate poor lift-off. Mixed and inconsistent results were obtained with the use of the bilayer LOR procedure. Most work in literature is concerned with the lift-off of metals, which typically have thicknesses approaching 1 μm . The lift-off of PTFE is different from the typical lift-off of metals since the thickness of our layers is comparatively thin (~ 20 nm.) Evidence in the literature concerning the lift-off of metals suggests that the mask should be on the order of the thickness of the layer one wishes to lift-off. Effective masks this thin are not realistic with the equipment available to the process. Further, the polymer nature of PTFE means it potentially consists of long distance molecular chains, which interfere with how the material is broken. This is in comparison to an amorphous metal that doesn't have extensive long range structural order, and is more likely to break locally. One other advantage of the LOR chemistry was that its solvent (Remover PG) could be safely heated, unlike acetone. Heat can provide a substitute for sonication in some instances. Trials with Remover PG heated to ~ 60 $^{\circ}\text{C}$ yielded positive results, compared to room temperature lift-off trials.

Despite our best efforts, this approach still produced significant challenges during lithography. The frames were too small to use with the smaller spinner chucks, and since they are vacuum mounted from the bottom there was some fear the pressure could damage the membrane. They were instead mounted on a 4 inch wafer and the photoresist was spun onto the wafer as for a regular process. The height differential between the mounted frames and the wafer caused large non-uniformities in the photoresist (Figure C.3,) indicated by the multicolour 'streaking' visible on the surface. Mask alignment and exposure was extremely difficult on the mounted frames. A mask was designed that had a grid of frames and domes patterned in a grid so that when multiple frames were placed on a single host, the mask could be aligned to find at least one that lined up with the frame allowing them to be exposed one at a time. Exposing the devices one at a time was done by placing a sheet of aluminum foil with a small cutout in the centre overtop of the mask, so that light could only reach

the single frame. Even with this mask, we often could not properly align a frame placed on the host. This whole process was extremely laborious and time consuming. Small differences in the height of the surface of the various frames placed on the surface meant that the mask was not always consistently close to the surface during exposure. Further, each time the wafer was approached for exposure any contact of the mask with the devices could result in disturbing the photoresist layer, as well as damaging the sensitive membrane inside. The lithography difficulties with the small sample could potentially be mitigated with the use of a newly installed mask-less lithography system in the nanoFab. While we still may have difficulties with the placement of resist, membranes could be individually patterned without much difficulty since alignment can be done digitally for each frame.

Since lift-off itself seemed to be the issue, a masked etch was also attempted instead. This would involve depositing the PTFE first, placing the patterned mask on top of this layer (of reverse polarity,) and then etching the PTFE below with an O_2 plasma. However, this was found to be impractical due to the high hydrophobicity of PTFE. Before spin coating could even begin, any photoresist placed on the frame surface would be wicked off the edges and onto the host substrate. This was made worse by the small surface of the device frame sitting on its host, since the liquid was wicked up by the lower, more hydrophilic wafer surface below. Even if this did produce a thin layer of photoresist, there may be further problems down the process during buckling since the required HMDS at the beginning of lithography could chemically modify the low adhesion, hydrophobic properties of the PTFE which makes it so valuable to our process.

C.2 Fabrication

During the testing phase for the various lift-off procedures, a layer of silicon nitride was deposited onto an Si wafer and diced into 5 mm x 5 mm sections. These ‘dummy’ frames allowed for trials to be run without fear of damaging a membrane inside, which is only present on the actual devices. Despite the




Figure C.3: A photograph of the host wafer with taped on devices, after the photoresist was spun on. The multicolour streaks are indicative of nonuniform photoresist.

various difficulties discussed above, a good lift-off was achieved with varying consistency on a few dummy samples over a series of lithography trials using the standard photoresist process, in addition to the LOR process. The best candidates from these trials were selected for a bilayer top mirror to be deposited on them. A 4.5 period 1550 nm equivalent Si/SiO₂ bilayer was deposited on these dummies, all mounted on a single substrate with the double sided tape and placed in Doug. The samples were then buckled in the RTA with a 10 second ramp to 400 °C recipe. A high quality buckle was achieved on a couple samples, shown in Figure C.4(a) and C.4(b). A Zygo scan was done on both, which revealed the sample in Fig. C.4.a buckled $\sim 4.0 \mu\text{m}$ tall and was 100 μm in diameter. The sample shown in Fig. C.4(b) buckled about $\sim 3.8 \mu\text{m}$ tall and was also 100 μm in diameter. This indicated that the process was in principle possible with all of the same materials as the real devices, and the only additional difficulty now would be in carefully preserving the integrity of the membrane.

For the actual membrane devices, the LOR bilayer procedure was chosen due to some evidence for improved lift-off. Inconsistent lithography was continually a problem, which meant that they had to be ‘cleaned’ multiple times




Figure C.4: Successfully buckled dummy devices with a 4.5 period 1550 nm equivalent Si/SiO₂ bilayer on top.

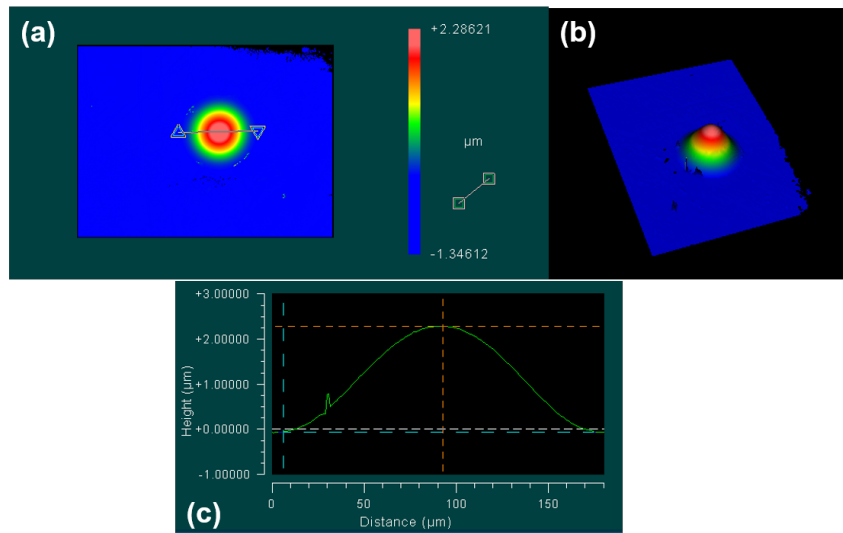


Figure C.5: Zygo data from membrane a. (a) A 2D heightmap (b) 3D model (c) Height profile, along the line between the two triangles in (a).

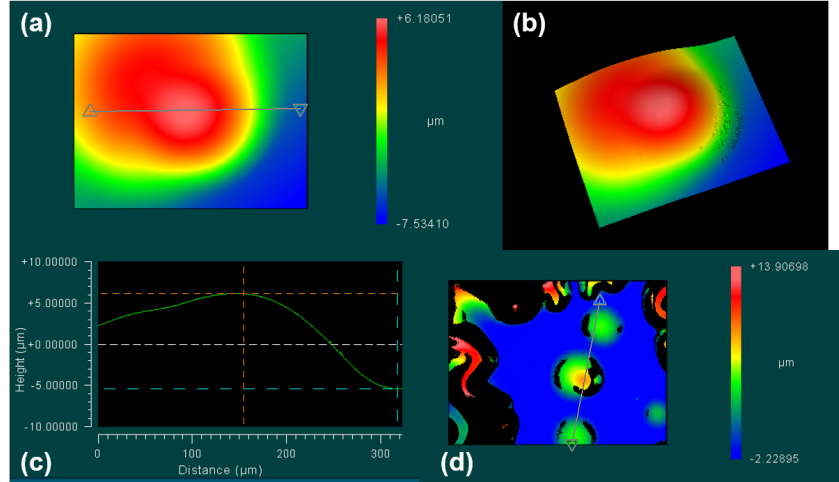


Figure C.6: Zygo data from membrane g. (a) A 2D heightmap (b) 3D model (c) Height profile, along the line between the two triangles in (a). (d) An additional heightmap of the sample at a lower zoom showing a broader area. Note the two other ‘bubbles’ of undesired buckling that occurred above and below the actual dome in the center.

both through stripping the mask chemically, and/or removing poorly patterned PTFE through an O_2 etch. This continued handling resulted in the loss of integrity of many of the membrane samples. With the membranes still mounted on a single wafer, they were lifted off upside down in a $60\text{ }^\circ\text{C}$ PG remover bath for 15 minutes. Photographs post lift-off, for a variety of membranes some good and some bad, are shown in Figure C.7. Multiple wafers were processed with a similar lift-off procedure, and the best were then chosen to receive a 4 period 1550 nm Si/SiO₂ mirror. These were mounted in Doug the same as the dummies previously. The devices were inspected under the microscope post deposition, shown in Figure C.8. Of note are the membranes labelled (c) and (d) in Figure C.8, which clearly have a broken membrane which has caused the mirror to not be able to form there. In all cases, the membrane is not visible underneath the mirror. These samples were then removed from their host substrate by removing the tape at $190\text{ }^\circ\text{C}$ on a hotplate, and reinspected to verify that no buckling had occurred at this stage accidentally. Two of these samples, (a) and (g), were individually buckled on the RTA with a 10 second to $400\text{ }^\circ\text{C}$ recipe. Photographs post buckling are shown in Figure C.9. The height profiles of these two devices were taken on the Zygo, and height maps




Figure C.7: Microscope images of various membrane devices, with patterned PTFE after liftoff as described in the text. The labels on these membranes correspond to those in Figure C.8 where the same ones are shown after the top mirror has been deposited.

of these are given in Figure C.5 and C.6. Sample a was $\sim 2.3\mu\text{m}$ tall, and $100\mu\text{m}$ in diameter. Sample g was measured at $\sim 11.5\mu\text{m}$ at its peak, but the erratic buckling made this measurement difficult.

At the time, I believed that the over buckling that occurred in sample g, which based on its recorded history as being an otherwise good sample, was a result of an unrefined RTA buckling recipe, causing the buckling to happen in such a fashion that caused it to happen outside of its PTFE pattern. For this reason, I held off on buckling further membrane samples and did some more tests on previously unbuckled dummy pieces with the bilayer on top of them. Integrity of the membrane was further verified using on the buckled samples using Zygo profilometer scans from the backside. If the membrane is intact, a square smooth and flat surface from the back should be observed. Zygo scans from these two samples are shown in Figure C.10 comparing side-by-side sample a and sample g. Clearly, sample g had a flat surface indicating the membrane is likely not broken, unlike sample a.

With this work here, we verified that it was indeed possible to buckle our




Figure C.8: Microscope images showing the membrane devices after a 4 period 1550 nm Si/SiO₂ mirror was deposited on top of them. The labels on these membranes correspond to those in Figure C.7 where the same ones are shown before the mirror was deposited.




Figure C.9: Microscope images showing the membrane devices (a) and (g) after buckling. (a) looks to have buckled well, but buckling of (g) was bubbly, and outside of the teflon outline which is visible underneath and indicated with the overlaid red arrow.

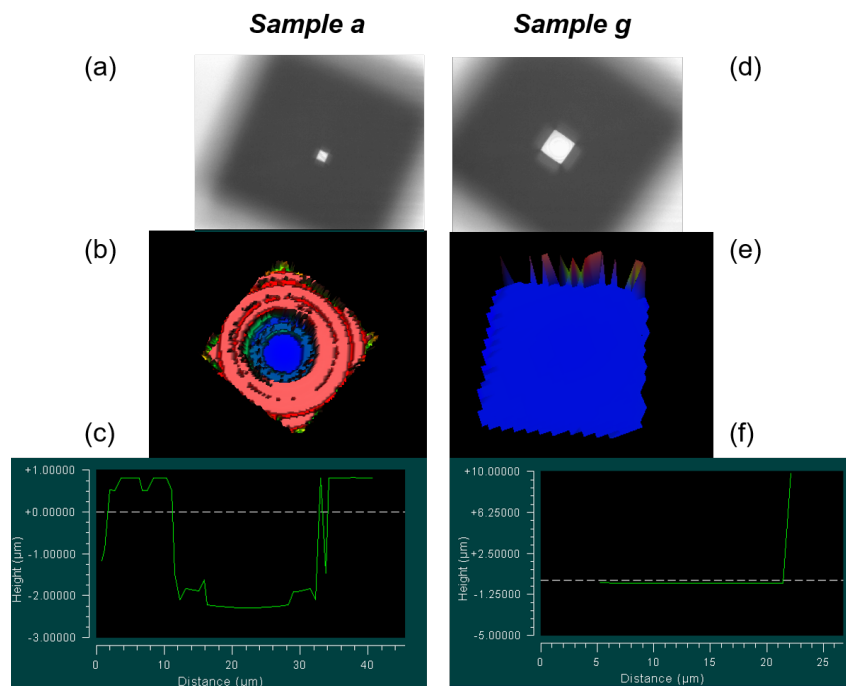


Figure C.10: Zygo data from imaging the reverse side of samples a and b (data for which is in the left and right columns, respectively.) (a) Microscope image from Zygo. (b) 3D map of bottom. (c) Profile map across bottom. Note in sample a a curved surface is visible, likely the inside of the dome, while in sample g the bottom is flat, which is likely the membrane.

domes on top of prebuilt membranes. However, there was a lot of difficulty in processing the small chips that caused many devices to fail from unintentional contamination and physical damage. Proceeding with the second half of the system development, which would involve all backside processing, would be tedious and impractical if this was to be done on single chips. Furthermore, placing the chips upside down is problematic since the device already formed on the backside can be damaged due to scratching or being torn off. Further development on this front would require a full, undiced wafer of membrane devices from Norcada, along with a matching mask of our own with alignment marks corresponding to the Norcada mask to allow for proper alignment of the domes to the membranes. It may also require additional steps taken to protect the frontside from damage during backside processing, such as a removable polymer coating or another wafer temporarily bonded to the frontside. Due to these concerns, it was decided that adopting a two-sided approach was just too difficult, especially since we were working with extremely delicate membranes. This approach was abandoned in February 2019, and instead work was started on the top down single-sided approach discussed in the main text.

Appendix D

Further Process Details And Operating Procedures

D.1 Process Details

D.1.1 Photolithography

In order to create any sort of features on the surface of a substrate, we require a way to transfer designs created on the computer into actual shapes on the wafer. First, we use a CAD program (Tanner EDA L-Edit) to design what we would like on the wafer in the form of a two dimensional map. The files from these designs are then sent to a mask writer. A 5 inch by 5 inch square piece of glass, 0.090 inches thick and coated on one side with a 100 nm thick layer of chrome, is placed within the mask writer. The mask writer uses an automated laser to then selectively remove chrome from the glass in the shapes and patterns designated by the mask file. The final product is a piece of glass with the exact patterns from the CAD written in chrome on the surface of the glass. There are a few technical points which should be outlined as to how the mask is constructed. From the patterns on the CAD files, the masks are written 'non-inverted' such that the dark parts of the mask (polygons, shapes etc.) become glass on the mask. All areas without shapes remain chrome. The patterns are transferred according to the RRCD (right read, chrome down) orientation. This means that when the mask is viewed from the glass side, the pattern will be as designed and not mirrored. This is done for the mask to be appropriate for frontside patterning.

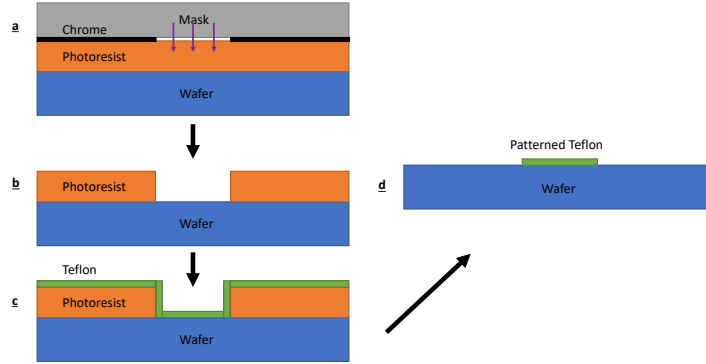


Figure D.1: A demonstration of how a mask is used for the patterning of a material through liftoff. (a) A mask exposes the photoresist, where the patterning is defined in terms of the photoresist portions that received light is to be removed. (b) The photoresist is developed and the pattern of the mask emerges. (c) PTFE is deposited ovetop of the entire wafer, overlapping both where the photoresist remains and where it was removed. (d) The wafer is submerged in acetone, which dissolves the photoresist, leaving behind patterned PTFE.

Two different kinds of photoresists are used in the fabrication of these devices. They are designated by model numbers from their manufacturers. The details of each is shown in Table D.1. It should be noted that a standard SOP supplied by the nanoFab is used for all of the resists except for AZ4620. Our group has developed a more expedient process for patterning, and has achieved acceptable results with it.

The lithography process starts with the placement of the wafers into the HMDS furnace. HMDS (hexamethyl-disilazane) is a chemical surface treatment that is commonly used in lithography in order to promote the adhesion of the photoresist onto the surface below. The wafer is heated up to 150 °C in a furnace and exposed to a vapour form of HMDS. The use of HMDS is more necessary when patterning on certain materials (such as silicon nitride,) but our group has found over time that it's usage results in more consistent lithography. In addition to the surface treatment, the HMDS process also acts as a heated dehydration step. The removal of all moisture from the surface can be critical to a good photoresist process. The full HMDS process takes 45 minutes.

Table D.1: Photoresist used in the fabrication process, with their corresponding recipes.

	AZ1512	AZ4620
Thickness	$\sim 1.1 \mu\text{m}$	$\sim 12.5 \mu\text{m}$
Exposure Dosage	$120 \text{ mJ}/\text{cm}^2$	$750 \text{ mJ}/\text{cm}^2$
Bake Time	100 °C for 60 s	1. Post Spin: 90 s N ₂ on + 60 s vac on. @ 100 °C 2. Re-bake after 15 minutes rehydration: 90s vac @ 115 °C, followed by additional 15 min rehydration
Developer	Diluted 1:4 AZ400K	Diluted 1:4 AZ400K
Development Time	45 s	80 s - 110 s

The wafers are then placed onto a spinner (Cee 200CB Coat-Bake System) and fixed from below to a rotating chuck with a vacuum. Approximately 5 mL of photoresist is poured onto the surface, and the spinner is started. All photoresists are applied in a two cycle process, known as ‘spread’ and then ‘spin’. Each cycle has a specified duration and rotational speed. These are specified in Table D.1. After the spinning process is complete, the wafer is immediately placed onto a hotplate for a ‘bake’. Baking is done at a set temperature for a specific duration. Depending on the recipe, it may or may not involve nitrogen being flowed from below the substrate. The baking hotplate is a part of the aforementioned spinner system.

A mask aligner system is now used to align the wafer to the mask. For the sake of simplicity where precise alignment is not required the older manual alignment system (‘Bert’) is used. When precise alignment is required, the more advanced computer controlled system (SUSS MicroTec MA/BA6) is chosen instead. While using the computer controlled system in principle produces better results, it’s complexity has more often than not resulted in more difficulty during processing as a result of both user error and maintenance issues. The principle of both alignment systems is the same. The mask is placed in a holder above the substrate and held in place with clamps and a vacuum. The substrate is placed below in a holder with three axis XYZ con-

trol, and held in place with a vacuum. With the aid of microscope cameras, the wafer is slowly brought into the correct position below the mask and into close contact below. Once alignment is satisfactory, the user sets in a specific exposure time for a UV lamp that then shines light down through the mask and onto the substrate. The UV lamps drift over time in their power output, due to environmental conditions and age, and so the lamp power is regularly recalibrated and the exposure time must be recalculated each time to reflect the most recent power measurements. The resists used here are sensitive to UV at both 365 nm and 405 nm. The lamps used in the nanoFab produce light around these two wavelengths, and thus the total combined power is used when determining the total dosage. In recipes, dosages are given in terms of mJ/cm^2 . Knowing the lamp irradiance in mW/cm^2 , the exposure time can be simply calculated. The dosage used in the recipes here differs from the standard nanoFab dosage, and has been chosen to suit our mask and features.

When working with a set of masks, they include corresponding patterns known as ‘alignment marks’ in order to assist in the placement of the masks in the correct position relative to one another. Careful placement of subsequent masks so that they line up relies on the careful alignment of the mask with multiple sets of alignment marks, on either side of the wafer. With the photoresist exposed to the UV light, they are then taken to be submerged into a developer solution. The developer solution is selectively soluble to portions that have been exposed by UV light. The solution used for all resists here is the diluted 1:4 AZ400K/DI Water solution. It is a solution of proprietary composition provided by the photoresist manufacturer for use with their specific resist. These solutions are typically basic. Development times are typically on the order of 30 seconds, and so exposure to the developer must be carefully controlled. The typical procedure is to pour the developer into a glass casserole dish, and then place the wafer inside. The dish is rocked vigorously throughout the process in order to ensure an even distribution of reactants in the solution. Once the time is up, the wafer is removed with tweezers and vigorously sprayed with water in order to halt the reaction from any residual developer on the surface. The wafer is then dried with nitrogen, and inspected

under the microscope. If the features look well defined, then the lithography is complete. If it appears that there is still some resist present where it should be removed, then the wafer is placed back into the developer solution for further development and is reinspected. Evidence for more development being needed may include a multicolour thin film pattern present on areas where there should be no photoresist (indicating residual photoresist,) or features (such as corners in the mask pattern) that are not sharply defined.

D.1.2 Sputtering

Operating Description

When using the Doug sputtering system, the procedure will be typically as follows. The chamber is opened up and targets of the users choice are placed in the guns. There are three guns, but we will typically just use guns 1 and 2. The metal components associated with the targets, including the dark space shields, the spacer rings, and the ‘Y separator’, are all removed and cleaned as often as possible. Our group has our own set of hardware which is sandblasted regularly to remove debris, and then wiped with IPA before being put into the chamber. The ‘cleanliness’ of the hardware is verified by testing the surface resistance with a multimeter. Little to no resistance indicates that the metal surfaces are free from contaminants, while resistance approaching $M\Omega$ usually indicates significant thin-film build up. This can cause irregularities in the electric field within the chamber during the sputtering process, and manifests as ‘arcing.’ Arcing, which looks like small sporadic lightning bolts between the surface of the target and the chamber walls, should be avoided as it causes damage to the target and unpredictable sputtering.

The targets placed in the guns are owned by our group. The life and quality of them is closely monitored. Typically, a target near the end of its life gives a lower than usual voltage during deposition as well as a slower deposition rate. Targets should be replaced well in advance, or else the material may be completely etched through and the target itself destroyed. A picture of this is in Figure D.2. We use separate and dedicated targets for both reactive




Figure D.2: A Si target that has been used too long, and broken. The inner annular pattern from the magnetron configuration has gone too deep, causing the centre to be carved out.

and non-reactive sputtering. It is important to avoid target ‘poisoning’, which occurs when a non reactive target such as Si is exposed to oxygen, and thus the chemistry of its sputtered film is impacted.

The targets we use are ordered from manufacturer Kurt J. Lesker. They are 3 inches in diameter and are bonded on their backs to a copper puck, in order to increase thermal conductivity with the gun and allow for better cooling. We currently use P-doped targets for both reactive and non reactive sputtering of Si. Doped targets are lower cost than pure intrinsic targets, and small amounts of dopants allow for a more reliable and repeatable conductivity. Although the dopant concentration is too low to affect any of the optical or physical properties of the films for our applications, it does impact the conductivity of the target itself and thus the power supply parameters used during the deposition.

The substrate is mounted upside down to the substrate holder. Small clips are used, which are adjusted such that they keep the wafer in place while blocking as little of the wafer surface as possible.

D.1.3 Piranha Cleaning

Piranha is a standard solution of acids, commonly used in microfabrication for the removal of organic contaminants and surface preparation. It is composed of a 3:1 volumetric ratio of 96% concentration sulphuric acid (H_2SO_4) to 30% concentration hydrogen peroxide (H_2O_2). Once the chemicals are mixed, a vigorous exothermic reaction begins and the wafers to be cleaned are submerged into the solution, usually in a PTFE holder. The wafers are removed after 15 minutes, after which the bath begins to lose its effectiveness. The wafers are immediately placed into a dump rinser for 5 cycles, and then into an SRD washer/dryer for a standard cycle. Wafers are cleaned with Piranha whenever possible before starting a process, since it ensures the surfaces are free from any contaminants and typically promotes adhesion of whatever layer is to follow (photoresist, sputtered film etc.) It should be noted that while bulk PTFE holders are resistant to Piranha, our patterned PTFE is not compatible and members of our group have observed erosion. Note that it may be possible to clean patterned PTFE safely with a cold Piranha solution.

D.1.4 Wet Deck Processing: Lift off and Mask Removal

At various steps throughout the process, the submersion of the wafer in different solvents is required. In addition to the previously described cleaning done with Piranha, acetone is commonly used to dissolve and remove photoresist masks. When appropriate, an electro-sonic cleaner ('sonicator') is filled with acetone and the wafer to be cleaned is placed inside a Teflon holder and submerged in the solution. The sonicator pulses high frequency vibrations through the liquid and substrate. In addition to the acetone chemically dissolving the photoresist, this sonication adds a mechanical effect that allows for a more thorough clean to be achieved. The sonicator is left on for a long period of time, typically one hour. Sonication is not done when buckled domes exist on the wafer, since this causes them to be ripped off and destroyed. In this case, the wafer is simply submerged in acetone and left to sit for a period of time. Regardless if sonication is used or not, the wafer is typically cleaned

after this process with a manual spray of acetone, isopropyl alcohol, deionized water, and then dried with clean nitrogen (in that order.) During lift off, if possible, the wafer should be placed such that the mask side is down. Gravity helps facilitate the migration of particulates away from the surface of the wafer, and reduces the chance any stray pieces of PTFE sticking back onto the wafer surface.

D.1.5 Zygo Optical Profilometer

The Zygo optical profilometer is a device located in the nanoFab's characterization lab. It is a microscope with an integrated interferometer. In general, when light is split into two paths and then later recombined, small variations in the path length manifests in the form of an interference pattern, or interference rings. These rings form a periodic pattern of intensity that will shift in phase as the path length of one of the arms is varied. While interferometers of this type are typically used to simply sense the change in distance of a beam path, when integrated into a microscope, they can be used to carefully map a surface provided it is sufficiently reflective. While the Zygo system excels at providing a map of a surface that is homogenous in it's depth, it begins to fail when surfaces used are not highly reflective, and when subsurface or buried features are present (such as the cavities in our domes.) Subsurface features confuse the system because interference patterns may look similar when coming from underneath as they do when coming from the top surface. Artifacts and false measurements have been observed, such as inversions (raised domes appearing as recessed pits,) or extremely high noise. It is important then that measurements taken with the Zygo system are not accepted at face value. However, the ease of use of the Zygo system compared to other more advanced measurement techniques (such as SEM) makes it appealing to use.

D.1.6 Confocal Laser Scanning Microscope (CLSM)

When samples are not highly reflective, the nanoFab's confocal laser scanning microscope (CLSM) is used in order to characterize the surface. Confocal microscopes are similar to regular microscopes, except that the optics used

to collect light from the sample have a very small depth of field. Traditional microscopes collect light from a large vertical area, while confocal microscopes have the ability to selectively view light from a select vertical slice of the sample planer. By collecting and analysing the images, detailed three-dimensional information can be obtained from the sample. By comparing the focus height with the amount of light collected, the amount of content of the sample at that particular height can be measured. A CLSM takes this same principle, but instead uses a computer controlled, tightly focused laser to rasterize the sample quickly and obtain height information at a very high number of points on the sample assembling a complete three-dimensional sample map. The CLSM system allows for improved imaging of transparent samples over the Zygo system because it measures the scattered light from a surface, instead of relying on coherently reflected light. Further, the ability to tightly focus on a specific height in the sample allows for the mapping of subsurface features, provided the user is roughly aware of what their device looks like. The system is not perfect though, and still tends to get confused at times when mapping our cavities. I believe the resonance of our cavities causes the laser to get captured and resonate internally during measurements, confusing the CLSM system and software.

D.1.7 Alphastep Step Profilometer

While optical techniques for measuring surface profiles can provide a lot of information at a high resolution, it has been noted that these can sometimes be subject to illusions and artifacts, due to the complex interactions of light with our multilayered surfaces. For this reason, it can be useful to corroborate potentially ambiguous optical measurements of surface features with a physical measurement. The principle of operation of a step profilometer is very simple. A very sharp needle is dragged across the surface of a sample, and based on how far down the needle falls before it hits the surface, height information about the sample can be obtained. By scanning the needle, one-dimensional profiles can be obtained. This technique gives a direct and high precision physical measurement of the surface, that is not prone to optical illusions like other

techniques. Unlike optical measurements however, this technique is invasive and involves touching the needle probe to the surface. However, the force with which the needle pushes down is highly controlled and on most surfaces there is typically no damage. The lateral resolution of this technique is limited by the sharpness of the needle. Also, it is highly susceptible to vibrations, and is unfortunately located in a very high traffic area of the clean room. The limitations with the system in the nanoFab also makes it difficult to align and make measurements with small features such as domes. It is for this reason it is mostly used as a quick, corroborative tool for measuring larger, micron order step edges.

D.1.8 Flexus

The stress of thin-films is extremely critical in the fabrication of our devices, since the buckling phenomena of the top mirror is highly dependent on a certain amount of stress in the top mirror. Measuring stress, compared to other physical properties, is a difficult process and is usually never measured directly. The Flexus system measures stress through determining the ‘bow’ of a wafer with a stressed film on top, and comparing this to a baseline of the wafer at a previous time when the layer was not present. The system takes a laser, and maps the height of the wafer surface across its diameter in a line. It saves this height profile for this particular wafer. At a later time, when a thin-film has been deposited on the wafer, the wafer is returned to the Flexus system and an identical measurement is made. If the film is under any stress, it should warp the profile of the whole wafer itself by bending or bowing the wafer in a particular direction. By comparing the profiles before and after, the stress can be estimated by knowing the thickness of the thin-film applied. The Flexus system in the nanoFab’s clean room is an old system, and does not allow for the removal of data from its computer. While the baselines are stored, the only information extracted from the computer is the stress measured after the software has done the fit itself. There is sometimes difficulty in accurately measuring the stress for multilayer mirror structures, due to the optical nature of the measurements. Testing should be done with the different wavelength

laser in the Flexus to provide more accurate results if need be.

D.2 Standard Operating Procedures (SOP)

D.2.1 Doug

The document in this section is taken from our group and is the product of efforts from past and current members. It is **not** strictly my work, but is included here for future reference and documentation.

OPENING AND LOADING

1. Press "emis" on the multigauge controller to turn OFF the ion gauge. Press "channel" until TC1 is shown and nothing else in the bottom right
2. Close the Baratron valve.
3. Close the cryo valve. Make sure the cryo temp is $\leq 20\text{K}$ before proceeding
4. Flip up the chamber vent switch. The pulsed DC power supply will read "Interlock User" once atmosphere is reached ($\approx 760\text{Torr}$). Close chamber vent.
5. Press the up arrow button on the hoist control to lift the chamber. Swing chamber top over to the left. Open the shutter on the desired gun(s).
6. Put on gloves.
7. Remove the Y insert separating the guns. Remove darkspace shields (DSS) and clamping rings for the desired gun(s).
8. Check the inside chamber for debris, vacuum if necessary.
9. Gather clean clamping rings and DSSs for each target (if necessary) - wipe with IPA.
10. Grab desired target(s) and inspect their condition. Blow off target(s) with nitrogen and load into desired guns.
11. When loading Si targets: Ensure you are using targets elastomer bonded to copper back plates. Place target in gun, followed by aluminum protector ring (cleaned with IPA). Place clamping ring on top and screw in place. Stop tightening as soon as resistance is felt - clamping the Si targets too tight can lead to target fracture.
12. Place clean DSS over gun(s). For $1/8''$ targets the DSS should be on the first notch. For $1/4''$ targets the DSS should NOT be on any notches.
13. Load substrate on holder - make sure there is a sharpie mark if alphastep measurement is needed.
14. Replace the glass viewport if opaque.
15. Check target spacing with voltmeter if using a metal target. For Si targets, target spacing can be checked by using the center pin in the antenna RF coax connector. You should have an open circuit with resistance in the M Ω range.

16. Replace Y insert
17. Swing substrate holder over and lower into chamber.

PUMP DOWN

1. Open the chamber roughing valve 1 full turn. Wait until pressure reaches 350 Torr then open the valve all the way.
2. When the pressure reaches 300 mTorr (3×10^{-1} Torr) close the roughing valve.
3. Slowly open the cryo gate valve.
4. Open the Baratron valve.
5. Press "channel" on the multi-gauge controller until BA1 is shown in the lower right of the display, then press "emis" - you should see the ion gauge turn on.
6. If using heater: wait until pressure is in the low 10-5 range, then turn on heater power.

DEPOSITION

1. Connect the appropriate power supply cable to the desired gun. The pulsed power supply (cable 3) is used for all reactive sputtering. Cables 1 and 2 are for DC sputtering of metals (eg: Ag). The external pulsed power supply can be used for reactive sputtering of metals only.
2. Press "emis" on the multi-gauge to turn off the ion gauge filament. Press "channel" until aux1 is shown in the lower section of the display - this is the reading for the Baratron gauge
3. Input parameters into the pulsed DC power supply (frequency, off time, power, ramp time)
4. Input flow parameters for gases on the MFC. Ar flow is typically set to 50sccm. After setting gas flow parameters, move cursor with left/right arrows off of the channel you are using (cursor should go white when editing is finished). If cursor is still "dark", you are still in edit mode and when turned on the MFC will use the max setpoint for gas flow.
5. On the MFC press "on" then "1" to activate Ar flow.
6. Turn on substrate rotation. 3 is the typical setting.
7. Close the cryo gate valve to reach the ramp pressure of the target. 3.75 turns to reach 7 mTorr, 3.875 turns to reach 10mTorr.
8. Double check power supply settings and that the guns are closed.
9. Press "on" on the power supply and allow target to ramp up to desired power.

10. If sputtering a target used for reactive sputtering, allow at least 2 minutes to burn the oxide off OR until voltage becomes stable (at the ramp pressure).
11. Turn the cryo gate valve counter clockwise to reduce to deposition pressure.
12. If reactive sputtering turn on oxygen flow by pressing "on" then "3".
13. Allow voltage and pressure to stabilize with oxygen flow on - usually 1-2 min.
14. Open shutter and time deposition to desired thickness
15. Once time has been reached close shutter
16. Turn off oxygen flow if reactive sputtering
17. Turn cryo gate valve clockwise to increase pressure to ramp pressure
18. On the pulsed power supply (while the supply is still on) change the power to 0W. The supply will ramp down at the same rate it ramped up. If processing the last layer turn off heater while power supply is ramping down.
19. Once the power ramps down to 30W turn off the power supply.
20. Repeat steps 7 through 17 for each layer
21. Last layer: turn off Ar flow by pressing "off" then "1"
22. Turn off substrate rotation.
23. Close the Baratron valve
24. Press "Channel" on the multigauge controller until TC1 is shown and nothing else in the bottom right.
25. Allow temperature to reach 10C less than processing temperature (if heater was used)
26. Close cryo valve
27. Flip the chamber vent switch up for a 1/2 second to add 1-2Torr of N2 into chamber
28. Wait 5 minutes.
29. Follow steps 3 through 5 of opening and loading section.
30. Remove heat shield, target(s), substrate and be sure to put power supply cables back to their appropriate gun.
31. Take care to inspect target condition and make notes of any cracks, chips, pitting, excess oxidizing, etc.
32. Follow the instructions in the pump down section to finish. Note during step 1, you can open the roughing pump all the way without waiting to get down to 350 Torr.

TROUBLESHOOTING

- Plasma won't ignite:
 - Wait 20-30s. If plasma still won't ignite increase pressure up to 10 mTorr
 - If at 10 mTorr plasma still won't ignite turn power supply off then back on
 - “RF trick” can be used to spark the plasma
 - * Turn on both “light switches” for RF supply (SW3,4)
 - * Turn Auto-enable Off and then On
 - * Quickly press RF On, then RF Off
 - If you are processing with another metal target (ex: Ti), you may need to ignite that target first to “spark” the ignition of the target
- Power supply reads “short circuit”
 - Do you have the right cable hooked up to the right gun?
 - If yes then the DSS is likely too close or touching the gun. You will need to vent and readjust the DSS
- Target is sparking:
 - The target is likely overheating, did you clamp the metal target tightly?
- Target is arcing:
 - There is oxide build up on the target surface, DSS and/or clamping ring.
 - If arcing is infrequent (1-2 times per layer) continue processing, but take note of when the arcs occur
 - If arcing is frequent (several times per minute) or the voltage is spiking, ramp the target down and stop the deposition. If arcing during the ramp down turn off the power supply. Continued arcing will damage and/or fracture the target.

D.2.2 Coupling Station

The following procedure should be used when lining up devices for measurement using the system described in the main text.

1. Sample is placed on the stage, overtop of the hole in the plate. Typically it is placed so that only a portion of the sample covers the hole.
2. Laser and camera are powered on, and camera is opened up on the computer for viewing. Auto Exposure is enabled at this point.
3. With the camera viewing a position through the sample plate hole where there is no sample, the laser is moved into view so that it is centered on the screen. It should appear as a roughly circular dot. Spherical aberrations and other interferences patterns are visible during this process, but it should very clearly appear as a dot.
4. Adjust laser height in order to make the spot from the laser as small as possible. If a good image of the laser is no obtainable at this stage, something is wrong and good coupling and measurements will be impossible.
5. With the camera and laser roughly aligned to one another in XY, turn on the lamp to illuminate the sample surface and move the sample stage in XY until the sample becomes visible on the camera. To make everything easier, one should start with a lower magnification objective. With the edge of the sample visible, adjust the sample stage height until the sample comes into focus.
6. The edge of the sample, in focus, should now be visible in the frame of the camera alongside the laser. Move the sample until the laser now passes through the sample, and adjust the laser height slightly in order to create the smallest spot possible on the sample surface. If the laser is not visible, you may have to either turn up the laser power, or change the wavelength. Due to the nature of the mirrors, certain wavelengths can be much more transmittive than others.
7. With a focused and crisp spot visible on the wafer (at a location with no devices present), the sample stage can now be moved until the laser is passing through a dome.
8. With the laser overtop of a dome, begin coarsely adjusting the wavelength while watching the transmission through the dome. The camera may not initially be in focus with the mode so that a clear mode is visible, but a good indicator of a mode being present is as wavelength is varied there is a significant change in brightness. Be aware of the expected mode spacing and width.
9. Having identified a wavelength that showed more light being transmitted through, move the laser in XY slightly to maximize this brightness and get the best coupling. Change the focus of the camera on the sample (the sample Z height) slightly to see if a mode of a recognizable shape can be obtained.

10. With a low Q sample or one with relatively low reflectance mirrors, the laser power must be turned down and the attenuation turned on in order to reduce the brightness on the camera so that a clear image of the mode can be found. Use this to further tune the laser XYZ, as well as the sample focus. It can sometimes be better to do focusing on the second order (two lobed mode) since it has a more distinct shape than the fundamental Gaussian mode.
11. Coupling and wavelength should be adjusted in order to maximize the transmission through the device on resonance.

If completing spectral scans with the camera:

1. Adjust the collection stage XY until the dome moves into the centre of the camera frame
2. Identify the portion of the frame that contains the modes, and excludes any noise from scattered light etc.
3. Input these pixels into the Python script to allow it to crop this portion of the image in real time.
4. With the laser power set suitably, turn off auto gain on the camera and set the digital gain to 0 dB. Adjust the exposure until a good SNR is achieved without saturating the image. Input this exposure setting into the Python program.
5. Input the desired wavelength range and resolution into the Python script.
6. Close the XCap Camera Viewing program, and run the python script.

If coupling to the photodetector:

1. Connect the photodetector and open up the provided software 'Ptero-DACtyl'
2. There is always a misalignment between the exact position of the fibre, and the centre of the frame shown on the camera. The difficult part in this step is to move the collection assembly in XY until the laser shines into the multimode fibre. At the time of writing, the fiber is 'located' just out of frame in the top right of the camera frame. This may change if it is removed and replaced. It is a very narrow region where light is coupled into the fibre.
3. Set the laser to a relatively high power, and change the collection stage XY while monitoring the photodetector power on the computer. Maximize this power.
4. FFT's can now be viewed directly in the software GUI, and this data can be saved and exported for further analysis. Typically averaging of i 30 consecutive FFT samples must be done to obtain a good SNR.

Appendix E

Mask Details

A set of three masks is used in the fabrication of these devices. The first contains alignment marks, the second contains dome patterns, and the third contains etch hole patterns. The domes are separated into dies according to the various configurations in Table E.1, as labelled in Figure E.1. Each die, regardless of its dimensions, contains a 32x16 array of all of the hole configurations specified in Table E.2. Each of the 8 hole layouts is placed twice within each row, and there are 32 rows.

Newer designs of this mask include alignment marks in more central positions, since the alignment marks on previous iterations were too close to the periphery and were not etched properly by the Cobra system since it has clamps that cover and block features on the very edge of the wafer. To further aid in proper alignment, some cells also have an inverse pattern of the etch holes on the bottom row - meaning that instead of the holes being open (as they should be for proper patterning) they are opaque on the mask with chrome. Although these devices won't form properly, they allow one to further inspect the centring of the etch holes on the domes in addition to regular alignment. This is not possible with etch holes usually because besides the holes themselves, the rest of the mask is opaque, blocking vision of the buckled dome underneath.




Figure E.1: The layout of a dome with etch holes on the mask. Labelled dimensions correspond to Table E.1.

Config.	1	2	3	4	5	6	7	8	9	10	11	12	13	14
D (μm)	50	50	50	50	50	50	50	100	100	100	100	100	100	100
H (μm)	10	10	20	20	20	30	30	20	20	40	40	40	60	60
E (μm)	3	4	3	4	5	3	4	3	4	3	4	5	3	4

Table E.1: Geometry configurations of the dies used in this mask.

Config.	A	B	C	D	E	F	G	H	I
Number of Holes	0	3	4	5	6	8	9	10	12
Angle Between Holes	-	120°	90°	72°	60°	45°	40°	36°	30°

Table E.2: Hole configurations present on each die.

Appendix F

Material Dispersion Models

When necessary, the full wavelength dependent (dispersive) optical coefficients are used instead of a constant value as is sometimes done for convenience. Materials, when possible, are fit using an ellipsometer to a Cauchy model. Ellipsometry is a technique used to characterize the optical properties of a material. It is used by our group primarily to experimentally measure the refractive index of a thin film with respect to wavelength, obtain reflection and transmission information of our mirrors with respect to wavelength, and measure thin-film thickness. The (three term) Cauchy model allows for the refractive index $n(\lambda)$ to be well approximated for dielectric materials in this wavelength range according to the following formula [43]:

$$n(\lambda) = A + \frac{B}{\lambda^2} + \frac{C}{\lambda^4} + \dots, \quad (\text{F.1})$$

where λ is the wavelength in microns, and A , B , and C are the Cauchy parameters for the particular fits. For applications where a high precision fit is required over a larger range of wavelengths, the similar but more complicated Sellmeier equations can provide more accuracy.

For the imaginary portion of the refractive index, an exponential model is used here to fit the loss. As specified by the ellipsometer software:

$$k(\lambda) = \alpha \exp\left(1.24 \mu\text{m} \cdot \beta \cdot \left(\frac{1}{\lambda} - \frac{1}{\gamma}\right)\right), \quad (\text{F.2})$$

where α and β are fitting parameters, and γ is the fitted bandgap wavelength of the material in microns. There is some ambiguity in the fit of this model,

	A	B	C	α	β	γ	Sample
Si	3.7093	-0.39745	0.42191	0.04168	0.093852	0.4	270220
SiO ₂	1.478	2.56x10 ⁻³	3.3506x10 ⁻⁴	4.3883x10 ⁻⁴	7.889	0.4	030620
Ta ₂ O ₅	2.0727	0.012393	2.138x10 ⁻³	0	0	0.4	0622020-2
Si ₃ N ₄	1.9862	0.02512	-3.3627x10 ⁻⁵	0.024544	1.3355	0.4	110320C

Table F.1: Fitting parameters for the dispersion models of relevant thin film materials. Here γ , B , and C are in units of μm , μm^2 , and μm^4 respectively. All other parameters are unitless.

since knowledge of the material bandgap parameter must be used to obtain an accurate fit. For fitting purposes, the assumed value for γ will be given along with the fits where applicable. If a more precise model for these parameters is required, tabulated values can be extracted directly from the ellipsometer and used instead. Table F.1 provides fit parameters used for the materials of interest in this work, and these are plotted in the figures in this section.




Figure F.1: Dispersion model of sputtered thin-film a-Si.




Figure F.2: Dispersion model of sputtered thin-film SiO₂.




Figure F.3: Dispersion model of sputtered thin-film Ta₂O₅. Loss is assumed to be negligible for this model.




Figure F.4: Dispersion model of PECVD thin-film Si₃N₄.

Appendix G

Thin Film Simulator GUI

Optical transfer matrix simulations will typically be performed in MATLAB, using ‘one off’ code separate for each project. While this makes it easy to segregate projects into individual MATLAB files, there is often difficulty in comparing results between individuals, and preserving a ‘master’ simulation. Often times, small modifications are made within the code to accommodate a special requirement of a project, and then if not removed later can cause unforeseen problems. Further, it is often difficult for new group members or those not completely comfortable with programming to create and share their own simulations.

To solve these problems and attempt to establish a shared ‘standard’ simulation among the group, a graphical user interface (GUI) was created in MATLAB and released for use among the research group. This program can solve arbitrary optical structures, with any materials one wishes to use provided their Cauchy and loss models. Figure G.1 shows an overview of one version of the program. Users input their structures graphically as a ‘list’ into the central interface, and then specify dielectric material modes from either a provided database or their own measurements. Structures may be saved for later and exported as files, or else imported from previous sessions in order to save time. The ‘Hecht’ and ‘Yeh’ formulations of the transfer matrix theory have been reduced to functions (also included as .m files,) which accept a standardized input and can also be used separately from this GUI in users own code for more advanced simulations (*e.g.* optomechanics.) More features may

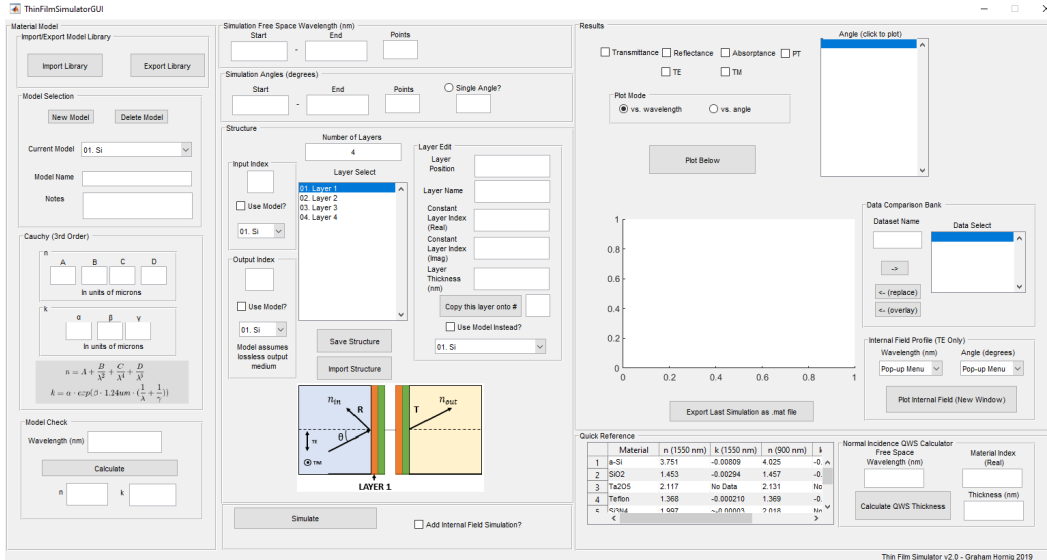


Figure G.1: A screenshot of the MATLAB Thin Film Simulator GUI version 2.1. Materials models can be selected from the list, or inputted from data (left column.) Arbitrary wavelengths and angles may be simulated (top centre,) and layer structure can be created graphically in the centre by clicking layers in the list and changing their individual properties such as name, thickness, and material model (centre.) Results are displayed in the right, with options to plot both linear polarizations of transmittance, reflectance, absorbance, and potential transmittance with respect to both angle and wavelength.

be added in the future such as parameter sweeps and dispersive metal models, as well as moving the program to a stand-alone executable format that would remove the MATLAB dependency and make it accessible on any computer.